

AN EPISTEMIC STRUCTURALIST ACCOUNT
OF MATHEMATICAL KNOWLEDGE

by
Lisa Lehrer Dive

Thesis submitted for the degree of
Doctor of Philosophy

2003

Department of Philosophy, University of Sydney

ABSTRACT

This thesis aims to explain the nature and justification of mathematical knowledge using an epistemic version of mathematical structuralism, that is a hybrid of Aristotelian structuralism and Hellman's modal structuralism.

Structuralism, the theory that mathematical entities are recurring structures or patterns, has become an increasingly prominent theory of mathematical ontology in the later decades of the twentieth century. The epistemically driven version of structuralism that is advocated in this thesis takes structures to be primarily physical, rather than Platonically abstract entities. A fundamental benefit of epistemic structuralism is that this account, unlike other accounts, can be integrated into a naturalistic epistemology, as well as being congruent with mathematical practice.

In justifying mathematical knowledge, two levels of abstraction are introduced. Abstraction by simplification is how we extract mathematical structures from our experience of the physical world. Then, abstraction by extension, simplification or recombination are used to acquire concepts of derivative mathematical structures.

It is argued that mathematical theories, like all other formal systems, do not completely capture everything about those aspects of the world they describe. This is made evident by exploring the implications of Skolem's paradox, Gödel's second incompleteness theorem and other limitative results. It is argued that these results demonstrate the relativity and theory-dependence of mathematical truths, rather than posing a serious threat to moderate realism.

Since mathematics studies structures that originate in the physical world, mathematical knowledge is not significantly distinct from other kinds of scientific knowledge. A consequence of this view about mathematical knowledge is that we can never have absolute certainty, even in mathematics. Even so, by refining and improving mathematical concepts, our knowledge of mathematics becomes increasingly powerful and accurate.

TABLE OF CONTENTS

Acknowledgements	v
Introduction.....	1
0.1 Mathematical Structuralism	2
0.2 Realist Assumptions.....	6
0.3 Outline of Epistemic Structuralism.....	11
0.4 Incompleteness.....	15
0.5 Definition of Terms.....	18
Chapter 1: Structuralism	29
1.1 Numbers as Sets or Objects.....	31
1.2 Set-theoretic Structuralism.....	38
1.3 <i>Sui Generis</i> Structuralism	41
1.4 Eliminative Structuralism.....	46
1.5 Modal Structuralism.....	47
1.6 Epistemic Structuralism	50
Chapter 2: Platonism and Structuralist Ontology.....	58
2.1 Platonism and Realism.....	58
2.2 Structures are in the World	63
2.3 Epistemic Structuralist Ontology	65
2.4 Mathematical Objects as Universals	70
Chapter 3: Abstraction and Mathematical Knowledge Acquisition.....	74
3.1 Structuralist Epistemology	74
3.2 Platonist Account	77
3.3 Abstraction: How it Works	85
3.4 Types of Abstraction.....	91
3.5 Abstraction and Epistemic Structuralism: Why it Works.....	99
3.6 Iterativism	103
Chapter 4: Basic and Derivative Structures.....	109
4.1 Two Levels of Mathematical Structures	109
4.2 From Basic to Derivative Structures	113
4.3 Ontological Status of Derivative Structures.....	115
4.4 Instrumental Use of Derivative Mathematical Concepts	122
4.5 Example: Number Systems	126
4.6 Infinity.....	129
4.7 Truth in Derivative Structures.....	135

Chapter 5: Skolem’s Paradox	144
5.1 Benacerraf’s Dilemma	144
5.2 The Skolem-Löwenheim Theorem and Skolem’s Paradox	147
5.3 Putnam’s Argument Against Moderate Realism.....	149
5.4 Truth Value of Mathematical Claims.....	152
5.5 Zermelo’s Refutation of Skolem’s Paradox.....	154
5.6 Implications of Skolem’s Paradox: Inadequacy of Formal Systems	156
Chapter 6: Mathematics as Science.....	163
6.1 Mathematics as Quasi-Empirical	164
6.2 A Posteriori Confirmation of Mathematical Truths	169
6.3 The Shift in Meaning Argument	174
6.4 Role of Proof.....	177
6.5 Formal Sciences Bridge the Gap.....	183
6.6 Computer Proofs	189
6.7 Formal Systems.....	195
Chapter 7: Mathematical Concepts	198
7.1 Justification of Mathematical Knowledge	198
7.2 Importance of Mathematical Concepts	204
7.3 Brown’s Argument from Notation	210
7.4 Theory Dependence of Mathematical Concepts	213
7.5 Incompleteness of Reference	216
Chapter 8: Certainty, Necessity and Fallibility.....	220
8.1 Certainty.....	220
8.2 Physical Necessity.....	223
8.3 Necessity of Mathematics	226
8.4 Fallibility of Mathematics	231
Conclusion	236
Bibliography	242

ACKNOWLEDGEMENTS

I wish to thank various members of the Departments of Philosophy at both the University of Sydney and the University of British Columbia, the graduate students of both these departments, and Mark Steiner, all of whom have given me valuable comments and feedback during the course of writing this thesis. I am deeply grateful to my supervisors, John Bacon and Andrew Irvine, who have both been outstanding in providing me with insightful comments, encouragement, and their own time. Finally I wish to thank Nanna and Gus, and Hugh, for all their support and encouragement over the years.

INTRODUCTION

To account for the indubitability, objectivity and timelessness of mathematical results, we are tempted to regard them as true descriptions of a Platonic world outside of space-time. This leaves us with the problem of explaining how human beings can make contact with this reality. Alternatively, we could abandon the idea of a Platonic realm and view mathematics as simply a game played with formal symbols. This would explain how human beings can do mathematics, since we are game players *par excellence*, but it leaves us with the task of specifying the rules of the game and explaining why the mathematical game is so useful - we don't ask chess players for help in designing bridges.¹

How do we come to know the truths of mathematics? What is the nature of mathematical knowledge? These two epistemological questions form the primary focus of this thesis. To answer these questions, I have developed a version of mathematical structuralism called *epistemic structuralism*. An important way in which this theory differs from other versions of structuralism is that its approach is motivated by epistemological concerns rather than purely ontological questions. This thesis describes in detail the nature of mathematical entities² according to epistemic structuralism, and considers it as an alternative to Platonism with respect to questions of mathematical ontology. An account of abstraction is developed in order to explain the processes by which we acquire mathematical knowledge. One of the advantages of this account is that it facilitates the integration of mathematical knowledge into a unified naturalistic epistemology, rather than requiring a distinct epistemology to account for our knowledge of mathematics. Abstraction grounds mathematical knowledge in our experience of the physical world. Within a framework of natural realism, epistemic structuralism accounts for both our knowledge of mathematics and

¹ Tymoczko (ed.) [1998], p. *xiii*.

² By "mathematical objects" or "mathematical entities" I mean nothing more than the things which one refers to in doing mathematics, its subject matter. In using the terms "objects" or "entities" I do not intend to ascribe to them any sort of independent existence as objects in their own right.

the applicability of mathematics in the physical world. It is argued that mathematical knowledge is just like any other kind of knowledge that we have about the world and that there are no significant distinctions between mathematical knowledge and scientific knowledge. This thesis examines the implications of such an account of mathematical knowledge by considering the implications of limitative results, particularly Skolem's paradox. A fallibilistic view of mathematical knowledge is defended. Mathematical theories are examples of formal systems by which we attempt to capture an aspect of reality, and there is evidence to support the inability of formal systems to capture any aspect of reality in its entirety. However the fallibility of mathematics is an optimistic fallibilism, as it does not preclude mathematics from having a strong explanatory power that is constantly being improved.

0.1 Mathematical Structuralism

The origins of mathematical structuralism are found in Benacerraf's famous paper "What Numbers Could Not Be"³ in which he argues that numbers could be neither sets nor objects, so the only essential properties they possess are the relational properties they hold with other elements of the structure (i.e. other numbers). Numbers have no internal defining properties, only structural properties. Resnik and Shapiro are also pioneers of structuralism, with their work on the subject culminating (to date) in their books *Mathematics as a Science of Patterns*⁴ and *Philosophy of Mathematics*⁵ respectively. Hellman champions a modal variety of mathematical structuralism as put forward in his book *Mathematics without Numbers*⁶, and also some more recent papers.⁷ Although Shapiro's and Resnik's accounts of mathematical

³ Benacerraf [1965].

⁴ Resnik [1997].

⁵ Shapiro [1997].

⁶ Hellman [1989].

⁷ For example, see Hellman [1990] and Hellman [2001].

structuralism are often considered representative of the theory, there exist several different varieties. The following chapter describes some of the versions of mathematical structuralism, and introduces the epistemic structuralism that I am advocating.

Many versions of structuralism, including epistemic structuralism, fall into the realist category of theories of mathematical ontology, since most varieties of structuralism (although notably, not Hellman's) hold that mathematical entities exist independently of mind, language and all other human attributes and constructions. There are three main arguments in favour of mathematical realism. The first is the Quine-Putnam indispensability argument,⁸ which claims that since mathematical entities are indispensable in science we must be ontologically committed to them. The second strong argument for mathematical realism is the argument from uniform semantics. This argument is based on the desirable goal of providing a uniform semantics for all discourse, both mathematical and non-mathematical. In trying to determine how we can refer to mathematical entities and say true things about them, it is a strong advantage if our theory of semantics works in the same way with respect to mathematical statements as it does for all other statements. The third argument for mathematical realism is Brown's argument from notation, which is discussed in greater detail in Chapter 7.

Platonism satisfies the desideratum of uniform semantics, since the statements "3 is less than 5" and "the book is on the desk" have the same semantics; "3," "5," "the book" and "the desk" are all independently existing entities, and any statements that we make about them are true in virtue of facts about these items. However a

⁸ There is some question as to whether this thesis is appropriately attributed to Quine. Although Quine infers the truth of mathematical claims from their necessity within physical theories, he nowhere refers to explanatory power. Even so, the label is now in common use. See Colyvan [2001] § 1.2.2 for an account of the development of this theory, and how it acquired this term of reference.

serious problem for Platonism is what Brown refers to as the “Problem of Access.”⁹ This is the problem of explaining how we can have knowledge about Platonically abstract entities that exist independently of space and time. This problem was identified in Benacerraf’s [1973], in which he explicates one of the main challenges in the philosophy of mathematics, namely that of reconciling an account of mathematical objects as abstract entities with a classical theory of truth and reference. A variety of structuralism such as Resnik’s or Shapiro’s, which take structures to be abstract in a fundamental ontological sense (making them similar to Platonic entities), faces this problem too. Since mathematical structures are taken to be primarily abstract, they exist outside of space and time, so it is difficult to explain how we have epistemic access to them.

The theory I am advocating, epistemic structuralism, is different from this variety of structuralism. I consider mathematical structures to be fully present in their physical instantiations. The structures themselves do not have an independent existence over and above their physical instantiations. They are not external to the natural world; when we refer to structures independently of their instantiations, these are simply ways of thinking or speaking about structures. This view has the advantage of explaining both how we can acquire mathematical knowledge, and why mathematics describes the physical world as well as it does. The natural numbers are the structure that all concrete sequences of discrete items have in common and, as such, they are located in the physical world. The structure of the natural numbers is not a substance in itself. It is a pattern that is repeated in the physical world. It is useful to consider the numbers independently of any particular instantiation, since it lets us reason about them effectively. Mathematical structures that exist as patterns in physical systems are what mathematical entities are identified with, in a fundamental

⁹ Brown [1999].

ontological sense, and our abstract concepts that describe mathematical structures are nothing more than mental reflections of an aspect of the physical world.¹⁰

The variety of structuralism which I advocate holds the physical instantiations of mathematical structures to be both epistemically and ontologically prior to the formal systems by which we attempt to capture them. When we observe a mathematical structure instantiated in a physical system, we form a concept of that structure. This idea does not constitute knowledge of, or acquaintance with, an abstract mathematical realm in the Platonic sense, since it is an idea of a physical thing, rather than an abstract thing in itself. When we have an idea of a mathematical entity, such as an initial sub-segment of the natural number structure or an element thereof, all this means is that we are holding in our mind a perceived feature of the world, in isolation from all its separate instantiations. We are thinking of the features that all concrete instances of a particular mathematical system have in common. By further abstraction we acquire concepts of other mathematical structures which are more deeply hidden in the world, and some that possibly are not even a part of the physical world at all.¹¹ The reason we can still consider statements about such structures to be objectively true is that they all have their origins in the physical world, and thus they are a part of objective reality. Statements about derivative structures incorporate elements taken from our experience of the world, and those structures that are not found in the world are still defined by characteristics that let us derive objective truths about them. Grounding mathematical structures in the physical world in this way solves the problem of access (or interaction) which is such a serious problem for most realist epistemologies of mathematics.

¹⁰ In providing definitions, I shall shortly distinguish between basic mathematical structures, which are patterns in physical systems, and derivative mathematical structures, which are not directly observable in the physical world, but are contained in and derived from our experience of the world.

¹¹ Chapter 4 provides an account of the various types of abstraction that yield derivative structures, and a thorough description of this type of mathematical structure.

Epistemic structuralism is not subject to the problem of access faced by Platonism and Platonistic varieties of structuralism, since the structures and relations that exist as a part of the physical world are ontologically prior to those that exist only in our minds. This means that the physical structures are also epistemically prior to our abstract ideas about them, because it is only by first interacting with the concrete instantiations, and then abstracting from the structures thus learned about, that we can acquire knowledge about mathematical structures. Thus epistemic structuralism is the theory that comes closest to explaining how it is that the nature of mathematical entities enables us to become acquainted with them. This is a distinct advantage over other forms of mathematical realism that take mathematical entities to be fundamentally abstract, including the current dominant variety of mathematical structuralism, that of Shapiro and Resnik.

0.2 Realist Assumptions

Since this is a thesis about mathematical knowledge, rather than about knowledge in general, my focus will be on specifying the nature of mathematical knowledge and addressing ontological questions that illuminate the nature of mathematical knowledge. Thus it is not my goal to defend a particular epistemological scheme. However, much debate about mathematical knowledge assumes a *causal theory of knowledge*. This is largely an effect of Benacerraf's influential paper "Mathematical Truth"¹² which brings out one of the most serious problems in the philosophy of mathematics: the difficulty of reconciling a classical theory of truth and reference with an account of mathematical knowledge. The difficulty stems from the incompatibility of the traditional Platonist position with a causal theory of knowledge. If mathematical entities are supposed to be Platonic, that means they exist outside of

¹² Benacerraf, [1973].

space and time. Hence we have no causal access to them. This makes an explanation of mathematical knowledge difficult, since even if Goldman's causal theory of knowledge is no longer the dominant epistemology, it is nonetheless difficult to account for knowledge without causation playing some role.¹³

Maddy¹⁴ argues that Benacerraf's argument fails since Goldman's causal theory of knowledge¹⁵ is best replaced by *reliabilism*. However, I will accept a broadly causal reliabilist epistemology, for two reasons. First, reliable methods of knowledge-justification are often causal. Second, it is difficult to show how there can be knowledge of a physical phenomenon without causal interaction.¹⁶ Thus I am not making a controversial assumption when working with the notion that in order to acquire genuine, objective knowledge there must be some degree of causal interaction.

This assumption is closely related to the aim of maintaining classical notions about truth and reference. Since the primary concern of this thesis is not semantic, I shall attempt to stay as neutral as possible in this regard. The account of mathematical knowledge is given without engaging in arguments contrary to classical notions of truth and reference, since epistemic structuralism should not have to depend upon particular mechanisms for truth and reference. This is in order to avoid one of the most serious problems for Platonists, namely that of accounting for how we have epistemic access to Platonic forms. This is the dilemma brought out in Benacerraf's [1973], and it forces Platonists either to abandon their semantics and ontology, or else to develop an epistemology specific to mathematical knowledge. Gödel¹⁷ and

¹³ This claim is defended in more detail in Chapter 3 of the thesis.

¹⁴ Maddy [1984].

¹⁵ Goldman [1973].

¹⁶ For example, this last point is brought out in the attempt in Brown [1999] to defend Platonism by finding an example of scientific knowledge that was obtained by non-causal means. This attempt, which is discussed in a later chapter, is not successful.

¹⁷ Gödel [1947].

Maddy¹⁸ both provide alternative explanations of how we can perceive mathematical entities that are Platonically abstract. However, it is an advantage of a philosophy of mathematics if it allows mathematics to be integrated into a dominant epistemology that is already established, and this is the aim of epistemic structuralism.

Another assumption this thesis makes is of the viability of an Aristotelian conception of universals. This thesis aims to explicate epistemic structuralism, and to defend Aristotelian universals is beyond the scope of the current project. However, it is acknowledged that a complete argument in favour of epistemic structuralism would include a defence of Aristotle's position on universals. While providing such an argument is beyond the scope of this thesis, it is hoped that the arguments for this version of mathematical structuralism provided may nevertheless hold some interest to those who do not favour Aristotle's approach to the problem of universals.

This thesis also assumes a broadly Quinean *naturalism*.¹⁹ Naturalism as explained by Quine in "Five Milestones of Empiricism"²⁰ is the doctrine that philosophy is neither privileged over, nor prior to, natural science. It is a view that accepts the scientific realisation that our knowledge of the world is limited, and sets forth to determine how we can know what we do. The naturalistic philosopher also recognises that some parts, although we do not know which, of our current theory about the world, must be wrong. He or she "tries to improve, clarify, and understand the system from within."²¹ This is in marked contrast to the notion of a first philosophy, which stands outside science and describes it from a clear, unobstructed perspective.

¹⁸ Maddy [1980].

¹⁹ By assuming a broadly Quinean naturalism, I am adopting a straightforward, common sense view of our relationship with the world that our scientific (and mathematical) theories describe. This view is based on Quine's philosophy of science, but in adopting this approach I do not intend that this thesis assumes all of Quine's philosophy.

²⁰ Quine [1981], pp. 67-72.

²¹ Quine [1981], p. 72.

Steiner²² makes a useful distinction between the notions of ontological reality, and epistemic reality, which echoes his analysis of Platonism into ontological and epistemic varieties. Following Quine's criterion of existence²³ he defines "to be real" as "to exist independently." However, 'independence' may be either ontological or epistemic. One entity is ontologically independent of another if the former does not depend on the latter for its existence. Steiner gives the example²⁴ of a hole in a piece of cheese as an entity that is ontologically dependent on another entity (the piece of cheese) for its existence. This means that even though the hole is existent it is not ontologically real, since it does indeed appear to exist, but it depends for its existence on the existence of the cheese (if the cheese were to disappear then the hole would no longer exist). Steiner contrasts the notion of ontological reality with that of epistemic reality. Epistemic reality has to do with an entity being independent of the conceptual scheme by which we discover and describe it. As Steiner puts it, "to be independent of our conceptual scheme is to be epistemically real."²⁵ If an entity is epistemically real, it can be described in at least two independent conceptual schemes.

Insofar as this project assumes a natural form of realism with respect to mathematical entities, it is the ontological form of realism to which I refer. Steiner argues that the indispensability argument of Quine and Putnam is unable to prove the epistemic reality of mathematical entities. The epistemic reality of an entity is usually demonstrated by "show[ing] that a theoretical entity is indispensable in explaining some new phenomenon,"²⁶ thereby acquiring new descriptions of the entity. However if mathematical entities are indispensable in the natural sciences, then no natural phenomena can be described without reference to some mathematical entities. This

²² Steiner [1983].

²³ Quine [1948].

²⁴ Steiner [1983], p. 363.

²⁵ Steiner [1983], p. 369.

²⁶ Steiner [1983], p. 373.

deprives mathematical entities of their explanatory power, since there can be no phenomenon that is describable without reference to mathematical entities, which is then explained if we assume their existence. If mathematical entities have no explanatory power, then they cannot be proven to be epistemically real. How can they be independent of our conceptual scheme, if they are indispensable to that scheme? Hence, according to Steiner, the indispensability argument not only fails to prove the epistemic reality of mathematical entities, it actually prevents the possibility of constructing such an argument.

It is not my intention to refute the claim that the Quine-Putnam indispensability argument proves only the ontological reality of mathematical entities. Following Putnam, I take the view that (ontological) reality of mathematical entities is plausible given the fruitfulness of mathematics, its success in describing the world. However as Steiner points out: “Putnam’s view is in dire need of an account of mathematical discovery, or coming-to-know.”²⁷ He gives the example of Gödel’s view as one epistemological scheme that could be used to complete Putnam’s view, but allows that some other view might also be able to account for the relationship between (real) mathematical entities and ourselves. The project of this thesis is to explain how the process of mathematical knowledge acquisition ensures the fruitfulness of mathematical knowledge. The epistemic structuralist account of mathematical entities aims to integrate mathematical knowledge into a natural realist epistemology, along with the rest of our knowledge about the world. The process by which we acquire knowledge of mathematical entities, understood in the epistemic structuralist sense, should demonstrate why mathematics is so successful in describing the world. However, it is not the project of this thesis to delineate this demonstration, only to put forward the epistemic structuralist account of mathematical knowledge

²⁷ Steiner [1983], p. 374.

and to argue for the integration of mathematical knowledge into a unified naturalistic epistemology.

0.3 Outline of Epistemic Structuralism

Given how well mathematics is able to describe the world, it seems natural to locate its subject matter within the world rather than in a Platonically abstract realm, human minds or a set of rules for the manipulation of symbols. In this thesis I advocate epistemic structuralism as a theory that gives a workable ontology of mathematical entities, as well as fitting in with an epistemology of mathematical knowledge. The variety of structuralism that I am advocating in this thesis can be summarised as follows:

- (i) Basic mathematical entities are real and exist independently of human thought;
- (ii) These entities exist in space and time and are part of objective reality;
- (iii) Mathematical entities are structures, which are patterns that may be displayed by various physical systems;
- (iv) We can observe many mathematical structures via abstraction from sensory perceptions;
- (v) We can refer to mathematical structures independently of their physical instantiations using a process of abstraction, but this does not make them Platonically abstract;
- (vi) Statements of mathematics possess objective truth values, independent of our ability or inability to obtain knowledge of them;
- (vii) Such statements obtain their truth-values as a result of properties of mathematical structures;
- (viii) Derivative mathematical structures are not directly instantiated in the physical world, these are abstractions from basic structures (those that are part of the physical world);
- (ix) Claims about non-instantiated mathematical structures have truth values because the structures are defined in terms of concepts that are abstracted directly from the physical world;
- (x) Our mathematical knowledge is fallible;

- (xi) Many mathematical truths are physically necessary (although not logically necessary).

The first three points specify the nature of mathematical entities, they explain what structures are like. Point (iv) tells us how we can grasp or understand mathematical structures and some of their properties. Points (v), (vi) and (vii) refer to the objectivity of mathematical statements, their truth values and the semantics of what makes them true. Point (viii) has to do with the objectivity of mathematical claims that seem to be entirely fictional, explaining their grounding in the physical world. Points (viii) and (ix) both specify the nature of mathematical structures that we cannot perceive, and that may not exist in the physical world. The final two points are epistemological, they have to do with the degree of certainty we attribute to mathematical knowledge. The notion of physical necessity as distinct from logical necessity is explicated in Chapter 8.

The following example will illustrate these points, although there are many more examples throughout the thesis. A sub-segment of a Euclidean (flat) plane is an example of a basic mathematical structure, which means it is a pattern displayed in various physical systems. The top of my desk, a blackboard and the surface of a smoothly frozen lake are all concrete instances of this structure.²⁸ It is important to note that, in a fundamental ontological sense, these concrete instances are the Euclidean plane. Without these instantiations of the structure we would have no idea of a Euclidian plane, and it is through its instantiations that we acquire knowledge of the plane. Our idea of the structure in isolation from any particular concrete instance is only a copy of the real, physical structure (more will be said on this point later). Since the structure is the pattern that these physical systems have in common, it obviously exists independently of human thought, within space and time and as a part

²⁸ These examples are all approximations to a flat Euclidean plane. Chapter 3 explains how we can acquire the idea of a flat Euclidean plane from an imperfectly flat surface. As well as abstracting away what the plane is made of and the colour of the plane, we can abstract away any surface irregularities.

of objective reality. We perceive the structure whenever we look at the desk, the blackboard, or the frozen lake. Noticing that they have something in common, we may want to refer to the properties that arise from this common structure, so we can speak of the structure as a thing in itself, independent of any particular instantiations. We do this by focusing on the planar properties of each physical system and ignoring all the irrelevant aspects. Referring to the structure itself, the concept of which we obtain in this way, is a useful way to express truths about mathematical structures, truths that will hold for any physical system that instantiates that structure.

For example, we notice that if we carefully construct two parallel line segments on the blackboard, the desk or the lake surface, the distance between them remains constant so they will never meet. This common property can be expressed by the mathematical statement: “parallel lines in a flat plane will never meet, and the distance between them remains constant.” This statement has an objective truth value, which is obtained in virtue of a fact about Euclidean planes. This fact is entirely independent of our discovering it, and would still be a fact if we had failed to discover it, or even if we never recognised flat planes as a recurring structure in the physical world.

In the same way that we obtain knowledge of a two-dimensional Euclidean plane, we can investigate three-dimensional space and even one-dimensional space. Having grasped these structures, we can identify their relations to each other, and thus we acquire the notion of dimensionality. If we abstract dimensionality from a particular space we have moved to the next level of generality, and by varying the dimensionality of a space we acquire the concepts of spaces that we have not experienced. Such spaces (for example a fifteen-dimensional space) are *derivative* structures that are not directly instantiated in the world, as far as we know. Any statements we make about the geometry of such spaces have truth values in virtue of being defined in terms of concepts that we already have from our experience of the world. This grounding in the physical world is what makes statements about

derivative structures objective. The value of derivative structures varies; some may have instrumental value in telling us something about the physical world, some may be interesting in their own right, and some may yield inconsistent systems and be of no use at all. We choose to investigate the derivative structures that have some value to us.

Let us return to our investigation of parallel lines on the surface of a frozen lake. Perhaps I am not entirely convinced that the parallel lines will never meet, so I ask a friend to drag a sled across the lake for several hundred metres, leaving a trail of two parallel lines and taking great care to keep them straight. If the lake is very long I may find that if I look along the lines they do in fact appear to approach each other, and perhaps if they are extended for a few kilometres they will appear to meet at or near their ends. In this case my limited perceptual faculties are unable to provide me with a clear view of the entire plane that I am observing, so I am misled into doubting the fact that parallel lines in a flat plane will never meet. This shows that it is possible to doubt a mathematical claim, and as we study more complex facts about mathematical structures it is less obvious which ones may be doubted, as the mathematical facts no longer correspond to observable mathematical facts about the physical world. Note that this is only one sense in which a mathematical statement can be fallible; Chapter 8 discusses the fallibility of mathematical claims in greater detail.

A final important feature to note about the claim that parallel lines drawn on a flat plane never meet is that this truth is a physical necessity rather than a logical necessity. The notion of physical necessity, explained fully in Chapter 8, means that a statement is always true, given that our world has the laws of nature that it does. It is logically possible that the statement is false, but it is not physically possible in our world. Given the laws of nature of our world, the statement is logically necessary, however it is not truly logically necessary since it would not be logically necessary in another world with different laws of nature. Hence it is a physical, rather than a logical necessity. It is a feature of our world (at least on a small, local scale) that if we

draw two parallel lines on a flat surface they will remain a constant distance apart and will never meet. It is conceivable that there exist a world in which seemingly flat planes behaved as our curved planes do, so that two parallel lines would eventually meet each other if extended far enough. Indeed this is how parallel lines behave in our world if a large enough scale is used. Euclidean geometry is only true of small, localised physical systems. It is also logically possible that there could exist a world consisting only of two atoms, so any lines in that world would have to have these two atoms as its endpoints and the notion of two lines having distance between them or touching would be meaningless. It is very difficult for us to imagine what it would be like in a world with different geometry to ours, since the underlying physical features of our world are so fundamental to our perceptions. Indeed, this is why mathematical truths often seem to us to be logically necessary. They are aspects of our experience of the world, however that does not entail their logical necessity. It accounts for why we believe the truths of mathematics with such a high degree of certainty, but that does not rule out the possibility of an entirely different physical world with different mathematical structures to ours. Thus the truths of mathematics are physically necessary, rather than logically necessary.

0.4 Incompleteness

Additionally to putting forward an account of epistemic structuralism and arguing for its benefits over other philosophies of mathematics, this thesis considers a potential problem faced by the theory. Putnam²⁹ argues that Skolem's paradox threatens a moderate realist approach to mathematics. The paradox demonstrates that axiomatic set theory cannot capture our intuitive notion of set. This result could be used to argue for Platonism, since a faculty of mathematical intuition could account

²⁹ Putnam [1983].

for our understanding of set. Putnam argues that if Platonism is not adopted, classical theories of truth and reference must be abandoned. This thesis argues against Putnam's claim, maintaining instead that the Skolem result reveals that our concepts or ideas of mathematical entities may not capture mathematical reality completely.

Gödel's first incompleteness theorem³⁰ says that any axiomatisable theory that is adequate for arithmetic is incomplete. This implies that an attempt to capture some arithmetical aspect of the world using a formal system will not be able to prove all the truths about that aspect of the world. This is not so surprising a fact as it might at first seem. Skolem's paradox shows that no formal axiomatised system can capture our intuitive notion of set. This is analogous to the incompleteness of arithmetic, since our formal axiomatisation of arithmetic cannot capture everything about the mathematical system itself. These phenomena are examples of a wider aspect of our interaction with the world and our attempts to formalise our experiences: formal systems that we construct generally cannot completely capture the aspect of reality that they describe. However, as we see in the physical sciences and in language, this does not preclude them from being useful. Although our abstract structure may not reveal *everything* about mathematical reality, it can still give us a lot of valuable insight.

My argument, which will be explicated in greater detail later in the thesis, is that Skolem's paradox shows us that set theoretical results are relative to their context, since they are about derivative structures rather than basic structures. The claim that all sets are constructible may be true or false, depending on whether it is an assumption built into the theory in question. Further, it may be true in some models and false in others, without any contradiction. Set theoretical results may not have absolute truth values. This claim may seem to indicate that mathematics is not objective, that our mathematical knowledge is entirely (or even partially)

³⁰ See Boolos & Jeffrey [1974] for details of Gödel's Incompleteness Theorems.

discretionary, based on the assumptions we choose. This is not the case. There is a high degree of objectivity in our mathematical knowledge, and this comes from its origins in our experience of the physical world. However, in more complex derivative structures we make some assumptions without knowing whether they reflect mathematical reality. These assumptions then contribute to determining the truths within that mathematical system. Every mathematical theory contains elements abstracted from the physical world, and many mathematical theories also have additional assumptions, which may or may not be true. It may be the case that the assumptions are more complex than mathematical reality, so there *is* no objective fact as to whether the assumption is true or false.

It is argued in Chapter 6 that mathematical knowledge is just like any other kind of scientific knowledge: we have theories about some aspect of the world that are incomplete, but this does not undermine their explanatory power or render them useless. We constantly refine and improve our theories and their concepts, and may even have competing theories. Generally, as our theories are refined they get closer and closer to capturing mathematical reality.³¹ This notion introduces a degree of fallibility into mathematics, however like any other scientific theory, it is still a powerful tool for describing the physical world. The primary distinction between mathematical knowledge and other kinds of knowledge about the world is the subject matter: mathematical structures are metaphysically very simple compared to the structures studied by other sciences.

This thesis argues against the view that mathematics exists in its own epistemological category, entirely certain and privileged over all other kinds of knowledge. The epistemic mechanisms involved in acquiring mathematical knowledge are not significantly different to those by which we acquire other kinds of

³¹ This is not necessarily true of all mathematical theories; some are not attempts to capture an aspect of the world, but rather are studied as interesting subjects in their own right.

knowledge about the world. This has the desirable result of unifying all our knowledge of the world into a single naturalistic epistemology, and although mathematical knowledge is no longer considered to be certain, it is still highly reliable. Fallibility does not render mathematical knowledge useless. On the contrary, the integration of mathematical knowledge into a unified epistemology removes the mystery surrounding its nature and justification.

0.5 Definition of Terms

It is worthwhile at this point to define some terms that I will be using throughout the thesis. By both *mathematical entities* and *mathematical objects* I mean the entities to which we refer when doing mathematics, namely the subject matter of mathematics. In terms of a definition this statement is somewhat tautologous, but it is intended only to show how I will be using the terms. The nature of mathematical entities is explicated in detail in Chapter 2. As mentioned in an earlier note, by referring to the subject matter of mathematics in this way I do not mean to ascribe to them any attributes, such as being independently existing abstract objects, or particulars existing in space and time and possessing various properties. When I use the term *entity* or *object* in this context I mean to imply only a very general ontological status. A mathematical entity is something that we can refer to and make claims about, and nothing more. Some formalists have argued³² that mathematics has no subject matter, that it consists solely of the manipulation of symbols and formal systems. However the striking explanatory power of mathematics in the physical world and its applicability to such a variety of non-abstract situations provides convincing evidence against such a formalist view.³³ Hence I shall work with the

³² For example Hilbert, see Detlefsen [1986].

³³ The Quine-Putnam indispensability argument is one forceful argument in favour of the existence of the subject matter of mathematics, see Colyvan [2001]. Another is the argument from uniform

assumption that mathematics has a subject matter, and proceed to determine the nature of the subject matter and how we obtain knowledge of it.

Eventually I shall argue that mathematical entities are *structures*, rather than objects in their own right. A structure is a recurring pattern exhibited by a system of objects. These could be either physical objects, or mental objects (concepts, which are explained below). *Mathematical systems* (also explained below) are used to describe structures. Structures are in the same ontological category as Aristotelian universals, in that they are not substances in their own right and have no independent existence over and above their instantiations. Although we may consider them in isolation from their physical instantiations, and speak about them as if they are abstract entities, they are not truly abstract, in the sense of being distinct from the physical world. Their physical instantiations are both ontologically and epistemically prior to any conception of them that we may have. Physically instantiated structures are objective features of the physical world, and do not depend on us for their existence. A *physical system* is defined as a collection of physical objects with relations holding between them.³⁴ An *exemplification* or an *instantiation* of a mathematical structure is any physical (or, sometimes, non-physical) system of which the structure holds. For example the contents of a (full) carton of eggs is an exemplification or an instantiation of the number twelve. Such instantiations are also sometimes referred to as *concrete instances* of the number twelve. Another example is the Zermelo-Fraenkel hierarchy of sets, which is (likely) a non-physical instantiation of the natural number structure (more about this later).

A *node* or an *element* in a structure is a component part that stands in a particular relation to other parts of the structure. The patterns that hold between the

semantics, which is based on the desirability of providing a single theory of semantics that applies to all discourse, both mathematical and non-mathematical.

³⁴ Shapiro [2000], p. 259.

nodes of a structure distinguish each particular structure. Nodes can be thought of as empty places in a structure that may be filled by objects, in the same way that a position in a sporting team can be filled by a person. When we speak mathematically we refer to the nodes as things in themselves and make claims about their relations to each other. We speak the same way when we refer to a position in a team, such as “the inside centre,” as if it was a person. If we say “the inside centre flanks the fly half” we are making a general claim that is true of any person who fills that position. A node in a structure is dependent on other elements of the structure for its definition and meaning. Just as “the inside centre” makes little sense if there are no other rugby players around, so the number 3 makes little sense if there is no number 2. Because nodes are typically defined by their position within a given structure, I may refer to a node in a structure as a ‘structure’. For example I may refer to a particular number as a structure. The number 3 is a structure in that it is a pattern that is repeated throughout the physical world, and for epistemological purposes there is no significant difference between treating 3 as a structure or as a node in a structure, even though this remains an important ontological distinction. Chapter 2 discusses this in more detail, however since my primary concern is epistemological I do not dwell on this distinction elsewhere.

A concrete instance or exemplification is distinct from an *abstract structure*, which is a *conception* or an *idea* of a mathematical entity.³⁵ The concrete instance is prior, both ontologically and epistemologically, to the abstract version.³⁶ Both conceptions and ideas of mathematical structures are concepts (defined below), with

³⁵ Abstract structures in this sense are not Platonically abstract. Conceptions and ideas of mathematical structures are concepts (defined below) which do not exist independently of us.

³⁶ This is more strictly true of basic structures than of derivative structures, which will soon be defined. For derivative structures we usually have access only to the idea, and not the concrete instance, since derivative structures generally do not have concrete instances (at least, not to our knowledge). In this case, the concrete instances of the basic structures from which they were abstracted are epistemically and ontologically prior to the resulting idea of a derivative structure.

conceptions being fully realised concepts of the structure, and ideas only partially realised. We may think of a structure independently of all its various instantiations using a process of abstraction (described in Chapter 3), and this allows us to consider the structure in isolation, independently of all the non-mathematical aspects of the system that instantiates the structure. A conception of a mathematical structure is thus normally a copy or an image in our minds of the (usually physical) structure, and is ontologically secondary to the structure in all its concrete instances. For example, three is the numerical aspect of the pens on my desk, while our conception of “three” is a mental reflection of the structure, which I get when I consider the numerical aspect of the pens independently of all other aspects such as shape and colour. If I imagine that there were a million times as many pens, I would have an idea of the natural number three million. This is no longer a conception, because three million is too big a number for me fully to realise it mentally. However I have an idea of it because I know what it is, and can partially realise it mentally.

A useful distinction to make with respect to conceptions and ideas of mathematical structures is between *conceiving* and *imagining* a structure. When we conceive of a particular structure, we realise it in full detail in the mind. In other words, we have a conception of the structure. When we imagine a structure, we realise it only partially. We have an idea of the structure. Most basic structures can be both conceived and imagined. For example, a triangle can be conceived, since we can hold all its details in our mind. In our conception of a triangle, the mathematical structure that it reflects is fully captured. We can also imagine a triangle, by thinking of aspects of it without fully conceiving of it. In the case of a one-million-sided polygon, we can imagine it (have an idea of it) because we know what some of its features are. However we cannot conceive of such a polygon, since we are unable to realise all its details in our minds.

These ways of referring to structures independently of their instantiations should not lead to their being confused with *Platonically abstract* (or *Platonic*)

structures, which are thought to exist independent of the physical world, outside of space and time. Platonically abstract entities are considered substances in their own right, as per the definition of Platonism given in Chapter 2. Conceptions or ideas of mathematical structures are concepts (defined below), and exist only in our minds.

A *physical entity* is any entity that is a part of the physical world, made up of matter and energy and existing in space and time. Physical entities include both observable (empirical or perceptual) entities such as trees, chairs, stars and human beings, as well as those that we cannot directly perceive, such as microwaves and protons (and most likely more entities that we have not discovered). For the physicalist, they also include all mental entities – entities which are either identical with, or in some other way reducible to, states of the brain and central nervous system.

Mathematical structures are physical entities (although not objects in the sense of substances) because they are a part of the physical world. Many mathematical structures (basic structures, defined below) exist in the non-mental physical world, and our concepts of them are reflections in our minds of a physical thing. If I look at the tree outside my window and then consider that tree in isolation, independently of the other trees around it, this does not make it a non-physical thing. I am simply thinking of a certain part or aspect of what I have just perceived and leaving out everything else. This is what we do when we consider basic mathematical entities in isolation. For mathematical structures that we do not know to exist in the non-mental physical world (derivative structures, defined below), our concepts of them may not correspond directly to independently existing non-mental structures. However, they are constructed out of basic structures that are independently existing physical entities. Additionally, they have objective properties because they are grounded (definition to follow) in structures that exist in the non-mental physical world.

Some mathematical structures are directly available to us through perception. We may infer the existence of others, and there are most likely many more

mathematical structures of which we have no knowledge. From the structures that we come to know through perception, we may derive concepts of other mathematical structures by *processes of abstraction* consisting of *recombination*, *simplification* and *extension* of the concepts that we acquire from the structures we already know about. The result is *derivative structures*, structures that are not found in the non-mental physical world. These are formal constructions, tools that often but not always give us greater insight into physical or mathematical systems. Structures that we know to exist in the physical world are referred to in this thesis, by contrast, as *basic structures*. It may be the case that a structure we originally thought of as derivative is later found in the physical world, and hence turns out to be a basic structure. Derivative structures that do not, as far as we know, exist in the (non-mental) physical world are considered to be *fictional* or *hypothetical entities*.

While derivative structures exist in the mind, statements about them are *objective*. Following Putnam,³⁷ I will take it that to claim that a statement is objective means that there is a fact of the matter as to whether it is true or false and, further, that what makes it true or false is external to the processes by which we discovered it. This means that statements about derivative structures are true independently of human cognition and epistemic processes.

The reason that we can make objective statements about structures that exist in our mind is because they are *grounded* in basic structures by a process of abstraction. This is to say that although derivative structures are formal constructions, they are composed of elements taken from basic structures, structures that exist independently of us in the physical world. For one structure (D) to be grounded in another structure, or collection of structures (B_1, B_2, \dots, B_n , where $n \geq 1$) means that the structure D does not possess any features that are not found in any of the structures B_i . All the features

³⁷ Putnam [1975], p. 57.

of D may be found in physically instantiated structures, although in D they may be extended, simplified or recombined. For instance, if D is a space it may take the feature of dimensionality from our space and extend it to 8 rather than 3. It is grounded in existent entities, since it is an abstraction from the concepts associated with them. If a structure D does not really exist, we may still wish to distinguish between claims that would be true of D if it was ontologically real, and claims that would be false of D if it was ontologically real. If D is grounded in structures that are ontologically real (B_1, B_2, \dots, B_n), then the grounding will allow us to make some conditional claims about D that are objectively true or objectively false, due to the properties of the B_i 's.³⁸ Of course, grounding a derivative structure in basic structures does not determine all the properties of the derivative structure. This is the reason that in more highly abstract mathematical systems, there can be undecidable claims.

It can sometimes be unclear where exactly to draw the line between basic and derivative structures: what does it mean for a structure to be *in* the physical world? Some seemingly non-existent or derivative structures may be instrumentally useful for describing aspects of the physical world, and many derivative structures have explanatory power for physical phenomena. Additionally, we might choose to draw the line by either epistemological criteria (basic structures are those that *we know* to be in the world) or by ontological criteria (basic structures are those that *in fact are* a part of the world, in which case it may even turn out that there are no derivative structures at all). However, where we mark the dividing point between basic and derivative structures is not of philosophical importance: what matters is that there are some structures that are clearly basic (such as a square) and some that are clearly, as far as we know, derivative (such as a 28-dimensional cube), and that the process of abstraction grounds aspects of derivative structures in our experience of the world so

³⁸ The notion of grounding is explained in greater detail in Chapter 4, and Hinckfuss' bridge laws show us how the resulting claims about derivative structures can be considered objective.

that they are objectively defined entities and we may say true things about them. Chapter 4 gives a full account of basic and derivative structures.

Conceptions and ideas of mathematical structures are concepts. *Concepts* are mental entities. This means they exist only in our minds.³⁹ However as pointed out by Burgess and Rosen,⁴⁰ concepts are concrete rather than abstract, since they exist in our minds, and many consider concepts to be in some sense reducible to states of the brain and central nervous system. Some concepts describe empirical or non-empirical physical entities, for example my concept of a chair is based on all my perceptual experiences of chairs. I also have a concept of an electron, based on the physics I learned in high school, but a theoretical physicist would have a more highly developed concept of an electron than I. We can also have concepts that do not correspond to any real entities at all, which we derive by combining elements of concepts that we already have. Generally these are fictional concepts, which form a subset of concepts more generally. For example most of us have a concept of a unicorn, which we obtain by combining the concept of a horse with the concept of an animal that has a horn on its forehead. Derivative structures that do not correspond to anything in the physical world are in the same ontological category as unicorns, for example our concept of a 28-dimensional cube or (arguably) infinity. Sometimes we cannot know if a concept corresponds to an entity in the physical world. Irrespective of whether a concept corresponds to an independently existing physical entity, the concept itself exists as part of the physical world, since it is reducible to or identical with particular states of the brain and central nervous system.

An *instrument* is a type of concept that does not correspond to any (non-mental) physical entities. The reason that instruments are useful is that they are

³⁹ In this context, I see no useful reason to distinguish terminologically between a ‘mental entity’ or ‘mental event,’ and an entity or event that is ‘in the mind.’ I shall use these interchangeably.

⁴⁰ Burgess & Rosen [1997], p. 15.

defined in terms of concepts that correspond to physical entities, and they can be used to reveal more information about them. Some derivative structures, although not all, are in this category.

Mathematical statements about basic structures can be true or false outright, namely true or false of the physical world. Statements about derivative structures are true only if the assumptions made in deriving the structure are true, that is, only if they are true in the conceptual system (or model) that contains those assumptions. This is a kind of hypothetical truth, whereas truths about basic mathematical structures are objectively or absolutely true. I will be using Hinckfuss' bridge laws to show how we can say true things about derivative structures without committing ourselves to their existence (again, see Chapter 4 for details).

A *mathematical system* or *model* is a conceptual framework that we use to capture truths about mathematical structures. Mathematical systems that describe basic structures are composed of basic mathematical concepts (concepts that describe basic mathematical structures), which correspond directly to basic mathematical structures in the same way that my concept of "chair" corresponds to various instances of chairs. Derivative structures are described by formal systems that are abstracted via extension, simplification and recombination from the models that capture basic structures. Statements in these systems or models have objective truth values in virtue of the grounding of their subject matter in basic mathematical systems.

A *privileged model* is one that corresponds to the physical world, so that statements that are true in this model are true in the world. We cannot know with complete certainty which is the privileged model, although we have several contending models that are useful in describing the world. There is nothing to say that we will not one day discover a complete privileged model (i.e. provide a complete formalisation of mathematical reality), but this is unlikely since we have only limited perceptual access to mathematical reality.

The claim that the truths of mathematics are *objective* means that what makes them true is something external to the process by which we discover them, something external to us. *Subjective* claims, by contrast, depend for their truth value on the knowing subject, or on the process by which they came to be known (or both). It is not the aim of this thesis to address the issue of whether mathematical realism implies that mathematical truths are objective. Therefore, I shall follow Putnam⁴¹ in associating the objectivity of mathematical truths with mathematical realism, meaning that mathematical entities are real if and only if statements about them are true independently of how we come to know them. This is a fairly natural position to take with respect to mathematical entities and our knowledge of them, since if mathematical entities are real then it is facts about them that are the truth-makers of mathematical claims, and these facts are independent of whether we know them. Similarly if mathematical claims are made true by factors that are independent of the processes by which we come to know them, then the most likely explanation is that they are about mathematical entities that are real and independent of us.⁴²

In a few places this thesis refers to the *applicability* of mathematics to the physical world. It is claimed that epistemic structuralism accounts for the applicability of mathematics, because it takes the subject matter of mathematics to be physically-based structures that are given to us in experience. This may be taken to mean the descriptive applicability of mathematics, defined by Steiner as “the appropriateness of (specific) mathematical concepts in describing and lawfully predicting physical phenomena.”⁴³ However, the primary focus of this thesis is to put forward the position

⁴¹ Putnam [1975].

⁴² The equation of objectivity of truth with realism of subject matter seems fairly natural in the case of mathematics, but it is not conclusive for all types of knowledge. For instance, it has been argued that moral claims are objectively true (rather than culturally relative), and yet moral values are not universally agreed to be independently existing real objects. In mathematics the claim is not as controversial, but to argue this point is not the purpose of this thesis.

⁴³ Steiner [1998], p. 25, italics removed.

of epistemic structuralism. This theory may form a basis from which to investigate questions of how mathematical methods of investigation can lead to physical discoveries, since the structures that mathematics studies find their origins in the physical world. Nevertheless it is not within the scope of this thesis adequately to address descriptive and metaphysical questions of applicability, such as those raised by Steiner.⁴⁴ It is hoped that an epistemic structuralist approach is consistent with and may support efforts in this area.

⁴⁴ Steiner [1998].

CHAPTER 1: STRUCTURALISM

Mathematical structuralism originated in the second half of the twentieth century. Benacerraf's famous paper "What Numbers Could Not Be"¹ demonstrated the problems involved with characterising mathematical objects as sets, or identifying them with any other kinds of objects. These difficulties led to mathematical structuralism, the theory that mathematical entities are recurring structures or patterns, with no internal defining characteristics. This chapter begins with a detailed analysis of Benacerraf's argument which provided the original motivation for mathematical structuralism. Since any identification of numbers with some other kind of entity leads to contradictory claims, he finds that the only distinguishing features that numbers have are the relations in which they stand to other numbers. An individual number makes no sense on its own, its identity being essentially dependent on where it stands in relation to other numbers in the natural number sequence. This is why any system that has the form of an omega sequence is an instantiation of the natural numbers. The natural numbers are identified with the structure that all such systems have in common. However saying that the numbers (or any mathematical entities) are structures is only a vague claim. During the decades since the notion of mathematical structuralism first arose, several varieties of the theory have emerged. Several prominent varieties of structuralism are examined; this analysis is not exhaustive, but it includes the dominant varieties of mathematical structuralism.

Shapiro² advocates *ante rem structuralism*, which takes structures to be abstract entities that exist as things in themselves. In this theory structures are taken to be *ante rem* universals, or Platonic Forms. Here I shall follow Hellman's lead and refer to Shapiro's theory as *sui generis structuralism*, so called because *ante rem*

¹ Benacerraf [1965].

² Shapiro [1997].

universals are independently existing things in themselves. Resnik's³ version of mathematical structuralism is also a Platonic variety. In order to consider only the main varieties of mathematical structuralism I will not establish a separate category for Resnik's theory, and (again, following Hellman) Shapiro's theory will be taken as the paradigm of *sui generis* structuralism. Another option is to take structures to be *in re* universals, or Aristotelian universals, which do not exist over and above their exemplifications. Shapiro refers to such an approach as *eliminative structuralism*, since structures are not taken to be things in themselves, but rather common (Aristotelian) structures that are present in each of their instantiations. Another variety of structuralism identifies structures with sets. I shall follow Hellman in calling this third theory *set-theoretic structuralism*. Hellman himself advocates yet a fourth theory, *modal structuralism*, in which mathematical entities are considered to be merely possible structures.

The variety of structuralism that I am advocating may be seen as a hybrid between Hellman's modal structuralism and the view mentioned above which takes structures to be Aristotelian universals. I shall refer to this hybrid theory as *epistemic structuralism*, since our mathematical knowledge is taken to originate in perceptions of structures in the physical world. This theory is distinct from the other types of structuralism in part due to its motivation. It is a theory that arose from an epistemological examination of mathematics, rather than from questions of mathematical ontology. This epistemic origin gives rise to a structuralism that places emphasis on our interaction with the physical world, but has significant parallels to Hellman's modal structuralism.

³ Resnik [1997].

1.1 Numbers as Sets or Objects

While the structuralist account of the nature of mathematical objects seems an attractive idea, traditionally it has not been a well-established view. Numbers have more often been treated as things-in-themselves. In mathematical logic, the natural numbers have often been characterised as sets in order to explain numerical statements in logical terms. Many mathematical logicians hold that numbers *are* in fact sets, in a fundamental ontological sense.

In his paper “What Numbers Could Not Be,”⁴ Benacerraf responds to this view, giving good reasons why numbers cannot be either sets or objects. He begins by exploring various ways of describing numbers in set-theoretical terms, and concludes that numbers cannot be sets. His argument is based on the fact that there are infinitely many different ways to represent numbers as sets, and that these different representations are incompatible with one another. The example Benacerraf uses is that of two boys who are both taught set theory prior to learning any arithmetic. Each boy is taught a different set-theoretical equivalent of the natural numbers, and once they have a full grounding in this theory they are given the vocabulary of arithmetic. The natural number sequence as conceived by each boy is:

- (i) $[\emptyset], [\emptyset, [\emptyset]], [\emptyset, [\emptyset], [\emptyset, [\emptyset]]], \dots$
- (ii) $[\emptyset], [[\emptyset]], [[[\emptyset]]], \dots$ ⁵

This means that both boys can converse and reason about the natural numbers just like the rest of us, but each has a different conception of the sets which correspond to each number. Each boy is capable of knowing, proving and deriving all the theorems of the natural numbers using his set-theoretical concept of the numbers. However, there are certain propositions on which they disagree. For example, the boy who thinks of the

⁴ Benacerraf [1965], pp. 272-294.

⁵ These are set-theoretical conceptions of the natural numbers which were first posited by Von Neumann and Zermelo respectively.

natural numbers as the sequence (i) would claim that 3 belongs to 17, since it is less than 17 and hence a member of 17. However, the other boy would disagree with this claim, since a number x belongs to another number y if and only if y is the immediate successor of x . Thus the statement “3 belongs to 17” would be true under one set-theoretical conception of the numbers, and false under another. The question we then face is that, if the numbers are indeed sets in an ontologically significant way, how do we find out which is the correct set-theoretical formulation? For example, if the number 7 *is* in fact a set of some sort, then there must be a unique set with which it is identified. However, we are unable to reason that any one set has a greater degree of identity with the number 7 than another; it makes no sense to choose a particular set arbitrarily and declare that it is identical with the number 7, since the truth of statements such as “3 belongs to 17” would be completely contingent on the set-theoretical formulation of the numbers that we chose. Since there is no unique best way to conceive of numbers as sets, and different representations can lead to contradictory statements about the numbers, Benacerraf concludes that numbers cannot be sets in any real, ontological sense.

This conclusion leads to the question of what sort of objects numbers could be, given that they cannot be sets. Benacerraf puts forward a second argument, which concludes that numbers cannot be objects of any kind. This is based on his finding that “*any* system of objects, sets or not, that forms a recursive progression must be adequate.”⁶ Since any recursive sequence will do the job of the numbers, the only distinguishing features that numbers possess are the relations in which they stand to other numbers. An individual number is not useful on its own, its identity being essentially dependent on the place which it occupies in relation to other numbers in the natural number sequence. Thus Benacerraf is led to the structuralist view of

⁶ Benacerraf [1965], pp. 289-290.

numbers, concluding that numbers cannot be identified with any individual set of objects.

Benacerraf claims that numbers constitute an abstract structure, and “the ‘elements’ of the structure have no properties other than those relating them to other ‘elements’ of the same structure.”⁷ This is not strictly correct, for as Wetzel⁸ points out numbers can have both structural and non-structural properties. For example, the number 7 has properties such as coming after 6 and before 8, and being half of 14, which are structural properties (defined as properties that relate the number 7 to other elements of the natural number structure, in virtue of their position in the structure). It also possesses many non-structural properties such as numbering the days of the week and numbering Snow White’s dwarfs. However, it is still the case that the structural properties of numbers are those which are essential and defining. These properties are essential, since if 7 ceased to possess them it would no longer be 7. The non-structural properties do not define the number (although they can specify it), and if one of the dwarfs passed away or the calendar was changed this would have no effect on the number 7, except that it would lose one of its non-essential properties.

Wetzel extends Benacerraf’s argument that numbers could not be objects of any kind, claiming that this extension is necessary in order to make his argument conclusive. She does this by constructing reductionist arguments analogous to the one that shows why numbers cannot be sets, using other arbitrary entities in place of sets, for example physical objects, expressions, or subsets of the natural numbers. Thus she shows that numbers are neither physical objects, nor subsets of the natural numbers, nor expressions, and so on. Her basic argument runs as follows: in order to show that numbers are not objects of any kind, it would be necessary to produce such an argument for each kind of thing that exists in the universe. Wetzel claims that in order

⁷ Benacerraf [1965], p. 291.

⁸ Wetzel [1989].

to argue successfully that numbers cannot be objects, Benacerraf first should have given an initial list of ontological categories, and then shown that numbers could not be any of these kinds of things. Furthermore, she questions why numbers themselves could not constitute such an ontological category. In other words, she argues that in order to show that numbers do not exist (as Benacerraf attempts to do) it is necessary to produce a reductionist argument (as he did for sets) for each category of existent entities, yet even this would not be sufficient to show that numbers do not form such a category themselves.

It can be argued that Benacerraf's argument that numbers are not objects is not in fact a parallel argument to the one that shows numbers not to be sets, but rather the same argument only more generally construed. It is not the case that Benacerraf first takes the ontological category of 'sets' and proves that numbers are not sets, and then moves on to the ontological category of 'objects' and does the same thing. Rather he demonstrates that numbers are not sets because there are so many different ways to define numbers in terms of sets and no way of choosing between them. This is a good way to begin because set theory contains clear examples of statements that can be true or false based on which set theoretical conception of numbers is chosen. However rather than moving on to a different type of thing, namely 'objects' instead of 'sets', he broadens his argument by pointing out that we may construe numbers as things other than sets. For example we could take numbers to be people, assigning each person a numerical value. This could lead to absurd claims such as "24 is the father of 71," whose truth value may be true or false depending on which assignment of numbers to people we use. If we take the numbers to be some other kinds of object, this statement would seem ridiculous. Since there is an indefinitely large variety of objects that we may identify with the numbers, any of the properties of numbers that are dependent on such an identification cannot be essential numerical properties. Benacerraf's argument thus applies to all possible identifications of numbers with objects, not only sets. He begins with the example of competing set theoretical

conceptions of numbers only because they provide a clear and undisputable example of the point he makes, which generalises beyond set theory to any identification of the numbers with objects that have incidental properties beyond their numerical ones. Thus, in an ontologically significant sense, numbers cannot be objects.

Even so, while Benacerraf's argument successfully shows that numbers cannot be sets, and that their only essential or defining properties are structural ones, this is not sufficient to show that numbers do not exist at all. As Wetzel shows in her analysis of Benacerraf's structuralist argument, this argument does not entail his final conclusion, namely that "there are no such things as numbers."⁹ Even if Benacerraf has successfully demonstrated that numbers cannot be identified with any other objects that have non-numerical properties over and above their numerical properties, this does not prove their non-existence. It means that numbers may be specified by their structural properties, and that these are their only essential properties. In her analysis of Benacerraf's "structuralist argument,"¹⁰ which is intended to show that numbers are not objects, Wetzel argues against Benacerraf's claim that numbers have only structural properties.

Benacerraf's structuralist argument against numbers being objects is based on the premises that numbers have only structural properties (namely, properties that arise from the relations they bear to one another, in virtue of the structure of which they are a part), and that any genuine objects may be individuated independently of their structural properties. Since numbers can be individuated only by their structural properties, Benacerraf concludes that they are not genuine objects. Wetzel argues that this is an unsound argument for two reasons. First, numbers do have non-structural properties. She gives the example of 9 having the property of "being the number of

⁹ Benacerraf [1965], p. 294.

¹⁰ Wetzel [1989], p.278.

the planets.”¹¹ This is a non-structural property because it makes no reference to the structural context of 9. 9 does not have this property in virtue of its position in the natural number structure, and the property of numbering the planets does not relate 9 to other elements of the structure. Nevertheless, as Wetzel points out, the property of numbering the planets certainly is a property of the number 9. The second reason for the unsoundness of Benacerraf’s argument that Wetzel identifies is his claim that non-structural properties cannot be used to individuate numbers. However the property of “being the number of the planets” may be used successfully to individuate 9, for when we refer to the number possessing that property it is evident that we are referring to 9, not 10 and not 8 (etc.).

Wetzel correctly points out that numbers have non-structural as well as structural properties, which Benacerraf does not acknowledge, although I think that this is not the most important point. She notes that as well as its structural properties, the number 10 has many non-structural properties such as numbering my fingers. However the critical point with respect to Benacerraf’s aim concerns the essential defining properties of numbers. The reason for determining whether numbers have both structural and non-structural properties is that we are concerned with finding their essential properties. Benacerraf is investigating what sorts of things numbers might be if they are not sets or objects, and in attempting to define the numbers, and find out what sort of entities they are, we must seek out their essential properties. He determines that the only *essential* properties that numbers have are those they hold in relation to the natural number structure, namely their structural properties. As mentioned earlier, the non-structural properties of numbers are not their defining properties. If I lost a finger then 10 would no longer number my fingers, but 9 would. This does not change the numbers 9 and 10 in any way, it simply alters their non-

¹¹ Wetzel [1989], p.280.

structural properties. However if we could conceive of 9 as no longer being the square of 3, then it would no longer be 9, because being the square of 3 is an essential and defining structural property of the number 9.

Recall the distinction between essential and non-essential properties, and also between structural and non-structural ones. The reason for clearly bringing out these distinctions is to clarify Benacerraf's hidden premise, as identified by Wetzel. The hidden premise is that "Any genuine object has essential properties that suffice to individuate it."¹² This means that if numbers are in fact independently existing objects that do not depend ontologically on the other elements in their structure, then they should be capable of being individuated purely by their non-structural properties. While Benacerraf takes this as an assumption, it is not clear that it is a correct one. If I refer to "the number immediately after 3" then I have successfully specified the number 4 using a structural property. Perhaps I have not 'individuated' it in the strict sense since I have used a structural property and have thus specified it with reference to other numbers, but since numbers are structural entities, this is what we should expect. Benacerraf's assumption that if numbers are genuine objects they must have non-structural, essential properties is unfounded, because we would expect structural objects to have structural properties as their essential properties, and this provides no grounds on which to deny them genuine objecthood.

If we discard Benacerraf's hidden premise then we are no longer bound to the conclusion that numbers do not exist. Just because a number is dependent on the rest of the number structure for its essential, defining characteristics, this does not entail that it is non-existent. The fact that mathematical structures can be exhibited by physical systems seems to show that they are real, if not independently existing. It is preferable to accept Benacerraf's structuralism about numbers but reject his hidden

¹² Wetzel [1989], p. 282.

premise, thus taking the position that mathematical entities such as numbers do exist since they are structures instantiated by physical objects. They are not physical objects in themselves, but are recurring patterns in the physical world, or structures instantiated in physical systems. The recurrence of these patterns in different physical systems allows us to recognise common properties, which are the essential properties of each mathematical structure. If we claim that mathematics is the study of *structures* that are found recurring in the physical world, it makes sense that individual mathematical entities cannot be separated from their structural context. The structural nature of mathematical objects is the essence of what they are, and asking for a definition that is independent of their structural context makes no sense. This does not prevent us from thinking of and referring to numbers as if they were independently existent objects, and it is frequently useful for us to be able to do so, but this does not entail that individual numbers have an independent existence distinct from the natural number structure as a whole.

1.2 Set-theoretic Structuralism

A seemingly natural way to conceive of structures is in terms of set theory or model theory. From a mathematical perspective, this is the ideal way to pin down structures and to provide exact characterisations of relations between structures. However from a philosophical perspective, there are several problems with this method. Hellman¹³ provides a good account of many difficulties with set-theoretical structuralism; here I shall examine a few of them.

The presupposition of a standard fixed ontology of sets is problematic. As Hellman points out “[t]here are problems both at the ‘bottom’ and the ‘top.’”¹⁴ The

¹³ Hellman [2001], §2.

¹⁴ Hellman [2001], p. 185.

problems at the bottom pertain to the null set: it is conceded by many¹⁵ to be a fiction, and yet it is the basis for all of classical set theory unless urelements are used. For example, Zermelo's set theoretical conception of the natural numbers begins with the unit set of the null set, and the successor of any number is defined as the unit set of that number. So 1, 2, 3, ... is expressed as $[\emptyset]$, $[[\emptyset]]$, $[[[\emptyset]]]$, Another example is von Neumann's set theoretic representation of the natural numbers, which also starts with the unit set of the null set, but defines the successor to any number as the set consisting of that number and all its elements. Under this definition, the natural number sequence begins $[\emptyset]$, $[\emptyset, [\emptyset]]$, $[\emptyset, [\emptyset, [\emptyset, [\emptyset]]]]$, The null set is used as a basis for set theory because the only other alternative is to use urelements. If urelements are used then set theory (and, under set-theoretic structuralism, mathematics too) becomes dependent on contingent objects. This is an undesirable basis for the truths of set theory and pure mathematics, which are not supposed to be contingent truths.¹⁶ However the only other option is to use a mysterious, possibly fictional entity as the basis for all mathematical knowledge. Neither alternative is acceptable as a way of characterising structures that are the subject matter of mathematical knowledge.

At the top end of set theoretical hierarchies there are difficulties too. There is the question of the existence of a unique maximal totality of sets: does it even make sense to speak of such a set? The other alternative is to accept a kind of "indeterminacy or open-endedness"¹⁷ at the top of the set-theoretical hierarchy. Both

¹⁵ Hellman [2001], on p. 185 cites Zermelo, Gödel, Fraenkel, Bar-Hillel and Levy as among the distinguished set theorists who have made this concession.

¹⁶ Later in this thesis it is argued that mathematical truths are contingent on the physical world, rather than being necessary truths. Thus my own position is not opposed to the truths of mathematics being contingent. However the contingency that arises from identifying mathematical structures with sets built from urelements is a different kind of contingency to the contingency described as physical necessity in Chapter 8. Reducing mathematics to set theory, but then using urelements rather than the null set, undermines the necessity that the grounding in set theory is designed to achieve.

¹⁷ Hellman [2001], p. 186.

of these options pose serious ontological questions that set-theoretic structuralism must resolve.

The plurality of set theories poses a serious problem for set-theoretic structuralism, as illustrated in Benacerraf's paper "What Numbers Could Not Be,"¹⁸ discussed in the preceding section. Since there are many competing set theories, how do we know which one represents the 'actual' numbers? The classic example is in trying to decide between Zermelo's and von Neumann's set-theoretic conceptions of the natural numbers. These two alternatives were set out above, and each is a legitimate characterisation of the natural numbers. The problem is that set-theoretic structuralism is committed to choosing between the two, since they are incompatible with each other. The natural numbers cannot be identical to *both* set-theoretical hierarchies; the number 3 cannot be $[[[\emptyset]]]$ and also be $[\emptyset, [\emptyset], [\emptyset, [\emptyset]]]$, because $[[[\emptyset]]]$ is not identical to $[\emptyset, [\emptyset], [\emptyset, [\emptyset]]]$, they are different sets. Further, claims about numbers may have different truth values depending on which set-theoretical representation is chosen. Thus we must choose between them, but we have no way of making that choice. There is no reasoning we can use to argue that the number 3 is identical to the set $[[[\emptyset]]]$, or to the set $[\emptyset, [\emptyset], [\emptyset, [\emptyset]]]$; nor do we have even a strong case for preferring one set-theoretical conception over the other. It makes no sense to choose a particular set arbitrarily, since mathematical knowledge is supposed to be objective and independent of us. This is a serious problem for set-theoretic structuralism, since for any mathematical structure there exists a plurality of set-theoretical interpretations. Necessarily, if we identify the structure with one interpretation, we are making an arbitrary decision that is doomed to be inadequate since it leaves out all other set-theoretical interpretations.

¹⁸ Benacerraf [1965].

These ontological problems pose serious difficulties for set-theoretic structuralism, making it an untenable position despite its obvious attraction from a mathematical perspective.

1.3 *Sui Generis* Structuralism

Sui generis structuralism, which may also be called *ante rem* structuralism, is the name Hellman uses for Shapiro's version of mathematical structuralism. This view takes abstract structures to be things in themselves. This is a Platonic conception of mathematical structures, since as Hellman puts it "[t]he structure is ... a Platonic abstraction, ... a reified pattern."¹⁹ Shapiro claims that mathematical structures exist independently of the systems that instantiate them. Their concrete instantiations allow us epistemic access to the structures via abstraction and pattern recognition. This avoids the dilemma arising from Benacerraf's [1973]. However reifying structures in a Platonic way gives rise to some difficult metaphysical and ontological problems (which are discussed in greater detail in the following chapter on Platonism).

Shapiro uses an analogy with functional roles in human organisations, such as governments or sporting teams, to convey his ontology of mathematical structures. He distinguishes between two ways that we can refer to the places in mathematical structures: the *places-are-objects* perspective, and the *places-are-offices* perspective.²⁰ The purpose of using these two perspectives is to make clear the distinction between the positions in a structure and the objects that may occupy it. The *places-are-objects* perspective treats the positions in the structures as if they were objects, at least grammatically. A statement such as "the Vice-President is responsible for security" is an example of the *places-are-objects* perspective. In this case, "the Vice-President" is

¹⁹ Hellman [2001], p. 188.

²⁰ Shapiro [1997], pp. 82-84.

not a fully-fledged object, it is a position within an organisational structure that may be filled with a genuine object. The places-are-offices perspective is used when the statement refers to the item from the background ontology which fills that place in the structure. In the case of sporting teams and other human organisations, the background ontology is people. For mathematical structures, Shapiro claims that sets or other structures can fill the places. Presumably concrete objects can also fill some of the places in some mathematical structures. Examples of statements using the places-are-offices perspective include “the full-back is taller than the left winger” and “the von Neumann 2 has one more element than the Zermelo 2.”²¹

Statements of pure mathematics are in the places-are-objects perspective, since they treat mathematical objects as independent objects. They refer to mathematical entities as things in themselves, and in Shapiro’s ontology mathematical entities are genuine objects. From another perspective, namely in another theory, the places in a given structure may be considered offices that can be filled by other objects. However in the relevant mathematical theory the places of the structure are genuine objects in their own right. For example, in arithmetic the natural numbers are objects. Thus the number 2 is an object. From another perspective, the number two may be considered an office that can be filled either with $[\{\emptyset\}]$ or with $[\emptyset, \{\emptyset\}]$. However under Shapiro’s *ante rem* structuralism, places in structures are taken to be genuine objects and not just empty places in a structure that may be filled with genuine objects from some background ontology.

Hellman correctly observes that Shapiro’s analogy between mathematical entities and functional roles in human organisations is not a complete analogy. We do not consider the Wallabies’ fly half or the Vice-President of the United States to be “eternal, abstract, featureless positions.”²² For one thing, they are dependent on the

²¹ This second example is from Shapiro [1997], p. 82.

²² Hellman [2001], p. 192.

whole structure (the Australian national rugby team or the government of the United States) for their existence. If Australia had no national rugby team, the Wallabies' fly half would not exist as an abstract entity. Additionally, the places in these structures are ontologically dependent on the people who fill them, or at least on having been filled by people at some stage. If the United States had never had a Vice-President, then even if it had been conceived of as a potential position we could not refer to "the Vice-President" (except perhaps in the future tense). We cannot make any claims except hypothetical ones about an entity that has never existed. If a rugby team did not have the position of fly half, then the Wallabies' fly half would never have existed. It can only be considered a genuine object if we are talking in the places-are-objects perspective, referring to the person filling that position, or to all people who have filled it. We do not consider positions in these structures to be objects that exist independently, in their own right.

One of the metaphysical difficulties with *sui generis* structuralism is determining the answer to questions such as whether the natural number structure is identical to the substructure of the real numbers that corresponds to the natural numbers. Is the real number 5 identical to the natural number 5? Shapiro's response to this question is no, that they are not identical, however as Hellman notes this is only by convention. Shapiro chooses to decide that the natural numbers are not identical to the substructure of the reals that corresponds to the natural numbers. There is nothing about Shapiro's structuralism that enables us to determine whether real 5 is identical to the natural number 5. If places in structures are genuine objects, then one would expect definitive, objective answers to identity questions, not mere conventions. It seems contradictory to leave an identity question between objects up to an arbitrary decision.

Another problem with this variety of structuralism is known as the identity problem. It is a characteristic of *sui generis* structuralism that the places in a structure are individuated only by the relations they have to each other. This leads to a problem

with structures that are invariant under non-trivial automorphisms. For example, in the complex numbers i and $-i$ can be permuted. The identity problem is that the structure obtained from the complex numbers by mapping i and $-i$ onto each other is structurally indistinguishable from the original complex numbers. If i and $-i$ can be individuated only by their structural relations, namely the relations that they have to other elements in the structure, then they have the same properties as each other, since their permutation is invariant under automorphism. This means that *sui generis* structuralism is committed to identifying i and $-i$, or at least that it is unable to distinguish between them.

Similarly in the mathematical structure of the integers, 1 and -1 can be mapped onto each other, and the internal structure of the integers remains unchanged. Keränen argues that the structuralist (in the sense of Shapiro's *sui generis* structuralism) "must maintain that $1 = -1$."²³ There is no way to distinguish 1 from -1 within the structure, and according to this kind of structuralism a mathematical entity can only be individuated by its structural properties. This problem arises from the invariance under automorphism of certain mathematical structures, however Keränen's argument only holds if it can be shown that the only individuating properties are those that are invariant under automorphism.

The structuralist claims that mathematical structures can only be identified up to isomorphism, since they have no internal defining features. This entails that the places of structures can only be individuated by properties that are invariant under the automorphisms of the structure. This view is consistent with Benacerraf's claim in his [1965], that the only essential properties numbers possess are their structural properties. Earlier in this chapter I argued against this claim, maintaining that non-structural (and non-essential) properties can be used successfully to individuate a

²³ Keränen [2001], p. 317.

mathematical structure. However the *sui generis* structuralist is not committed to the possibility of specifying numbers using their non-essential properties, since structures, being abstract entities, can exist independently of any instantiations. Thus the *sui generis* structuralist would accept the claim that the only individuating properties mathematical entities possess are those that are invariant under automorphism. There is an automorphism that maps each positive integer onto its negative integer, and this automorphism preserves the relational properties of the natural number structure. Hence the (*sui generis*) structuralist has no way of distinguishing between 1 and -1 within the structure.

The identity problem is a serious one for *sui generis* structuralism. As is evident from Keränen's explication of this problem, it relies on the structuralist claim that only the essential or structural properties of mathematical entities can individuate them. To take an example, the number 9 has many properties that can be divided into essential and non-essential, and structural and non-structural. For mathematical entities (places in mathematical structures) their essential and structural properties coincide. The number 9 has the structural property of being the square of 3, and this is also an essential property. Structural properties cannot change since mathematical structures are not dynamic, they are recurrent patterns. They may be instantiated in dynamic systems, but the structures themselves are unchanging. Hence a property or relation that holds between two places in the structure is an essential defining property of the places in question, since it cannot change. If the number 9 were no longer the square of 3, then it would no longer be the same number. Similarly 9 is and will always be equal to the sum of 4 and 5. This is an essential defining property of 9 and it will never change. The number 9 also has plenty of non-essential properties that have extra-structural references. For example it has the property of numbering the planets in our solar system. This is not an essential property, for were Mercury to burn up and disperse, or if another planet were discovered, then 9 would no longer number

the planets in our solar system. It would still be the same number, but its non-structural and non-essential properties would have changed.

The identity problem rests on the structuralist claim that it is only the essential or structural properties that can individuate mathematical entities. It is claimed by structuralists such as Benacerraf and Shapiro that the places in mathematical structures can only be successfully individuated by their essential properties, those that hold between the place in question and other places in the same structure. In the section on epistemic structuralism later in this chapter it is argued that this claim is unfounded, and that a mathematical entity can also be individuated successfully by non-essential or non-structural properties.

1.4 Eliminative Structuralism

Eliminative structuralism characterises mathematical entities as *in re* structures. Rather than Shapiro's structures, which are Platonic forms, eliminative structuralism eliminates structures as things in themselves. According to this theory, structures are Aristotelian universals. They do not exist independently of the systems that exemplify them. Mathematical entities are ontologically dependent on their instantiations in the physical world. According to Aristotle, if there were no red things, then the universal property 'redness' would not exist. Similarly, under eliminative structuralism if there were no square things, then the mathematical structure of a square would not exist.

In eliminative structuralism the places-are-objects perspective is explained away in terms of the places-are-offices perspective. When we say something about a mathematical entity using the places-are-objects perspective, we are really just generalising over all the instances of the structure. We are saying something that is true of any concrete system that displays the mathematical structure in question. For instance when we say "a square has four corners," this is just a shorthand way of saying "every object that is square in shape has four corners." Under this theory every

mathematical structure must be instantiated in some system. Any mathematical claim is a generalisation over all instances of the structure. It is a statement about the properties of the structure that all instances have in common, but this common structure is ontologically dependent on its instances, as are all Aristotelian universals.

The main problem with eliminative structuralism is that the background ontology must be large enough to account for all of mathematics. Many mathematical structures are not instantiated in concrete objects, yet if they are to be explained under eliminative structuralism there must be some mechanism to account for their truth and objectivity. If all mathematical structures are considered to be *in re* structures, then those that are not exemplified do not exist. While we could find some way of accounting for references to such structures, there remains the problem of explaining their objectivity. We consider mathematical truths to be objective rather than arbitrary, and their applicability to, and explanatory power about, the physical world provides strong evidence to suggest that they are not arbitrary claims. However if they consist of references to non-existent structures, then their objectivity and usefulness becomes difficult to account for.

1.5 Modal Structuralism

Hellman's version of structuralism takes mathematical entities to be 'possible structures' quantified over by modal operators. Mathematical claims are understood to state "what *would* hold in any structure of the appropriate type that there *might* be."²⁴ This avoids ontological commitment to any structures or mathematical entities at all, and hence this approach dodges many of the difficulties arising from the ontology of

²⁴ Hellman [1990], p. 314.

structures. The basis for mathematical truth is logical possibility, and even Hellman admits that “modal primitives for mathematics are problematic.”²⁵

In modal structuralism mathematical statements are analysed using a second-order modal logic. Second-order logic is required because first-order logic is insufficient to express mathematical claims. The modal aspect is used to express the counterfactual nature of mathematical claims that is the hallmark of modal structuralism. Statements of mathematics are not taken to refer to objects that actually exist; instead they quantify over all objects that could possibly exist. Mathematics consists of statements that express what would be true of particular structures, if it were the case that they exist. There are no claims made about what actually does exist, thus on Hellman’s account the whole body of mathematical knowledge may be considered counterfactual.

There are many advantages that modal structuralism has over other kinds of structuralism, and *sui generis* structuralism in particular. Modal structuralism seems to have incorporated the beneficial aspects of each of these theories, while avoiding many of their problems. Like set-theoretic structuralism and eliminative structuralism, modal structuralism does not attribute an independent existence to abstract structures. This avoids the problems that Shapiro’s theory encounters as a result of considering empty, featureless positions in structures to be independently-existing objects. However modal structuralism shares many of the benefits of *sui generis* structuralism. Modal structuralism applies equally to any mathematical theory, and there is no particular theory that is taken to be primitive, as in the case of set-theoretic structuralism. These advantages make modal structuralism an attractive theory.

One difficulty that arises with modal structuralism concerns the link between the physical world and mathematical truth. In *Mathematics without Numbers*²⁶

²⁵ Hellman [2001], p. 205.

²⁶ Hellman [1989], Chapter 3.

Hellman extends his modal structural account of pure mathematics to include applied mathematics. Modal structuralism involves counterfactuals, so the antecedents of conditionals need not be true. In the case of applied mathematics this is no longer the case, since the antecedents *are* true. He provides a modal structural formulation of the statements of applied mathematics, which are taken to express what would be true in possible worlds containing models of certain structures. It seems that this method of taking a modal structural account of mathematics (which considers mathematical truths to be counterfactual), and adapting it to include applied mathematics, seems to downplay the applicability of mathematics to the world. The link between mathematical truth and the physical world seems almost to become a coincidence, rather than mathematics being a fundamental part of our interaction with the world. Hellman acknowledges that the implications of indispensability arguments for mathematical ontology are an important consideration.²⁷

I will not go into this issue in detail, since the purpose of this thesis is not to refute modal structuralism, but to provide an alternative in epistemic structuralism, and to situate it among the other available theories. Epistemic structuralism has common elements to modal structuralism, which remains a promising alternative. Further development of modal structuralism could result in an account of mathematical knowledge that explains the strong connection between mathematical truth and the physical world. To develop modal structuralism is not the purpose of this thesis; rather, my aim is to investigate the viability of taking our epistemic interaction with the physical world as a starting point, and trying to build an epistemology of mathematics that fits with a naturalistic epistemology and moderate realist ontology. I do not discount modal structuralism as a viable alternative, although I consider the problem of explaining the link between mathematical truth and physical reality to be a

²⁷ Hellman [1989], p, 95.

serious one. However I do not claim it to be an insurmountable problem, although it is not my aim to address this issue here.

1.6 Epistemic Structuralism

I shall use the term *epistemic structuralism* to describe the version of structuralism that I advocate. This is an epistemically-motivated theory that can be seen to be a hybrid between eliminative structuralism and Hellman's modal structuralism. There is a distinction between mathematical structures that have concrete exemplifications, and those that do not. Clearly some structures have concrete instances, for example an initial sub-sequence of the natural numbers. These are *basic structures*. Then there are some that, as far as we know, do not. A 28-dimensional sphere is an example of a structure that we can describe mathematically, but that we cannot observe. Mathematical structures that do not have concrete instances are *derivative structures*. As explained in Chapter 4 the dividing line between basic and derivative structures is not clearly demarcated. There are different ways that we can draw the line between them, but exactly where the line lies is not of crucial significance. The important thing is that there are some mathematical structures that we can observe via abstraction from our experiences of the physical world, and some that we do not believe to have any concrete instantiations. The distinction is not a simple one to make, as it involves drawing a line in a continuum, hence there are different criteria that we can use to draw the line. However there is a clear and important difference between those mathematical structures that we can observe, and those that are not exemplified in the physical world. The fact that it may be unclear which category a particular structure falls into is not a serious problem. The Chapter 4 gives much more detail about the distinction between basic and derivative structures.

It is possible to explain epistemic structuralism in terms of other varieties of structuralism already encountered. With respect to basic structures (those that are

instantiated and probably observable in the physical world), epistemic structuralism is essentially indistinguishable from eliminative structuralism. Structures are taken to be Aristotelian *in re* universals. They are not independently-existing abstract entities, rather they are ontologically dependent on all their instantiations. Recall that eliminative structuralism explains away the places-are-objects perspective in terms of the places-are-offices perspective. When we use the places-are-objects perspective, for example in the mathematical statement $7 > 3$, we appear to be making a claim about two abstract entities; however this perspective is a shorthand way of generalising over all instances of the numbers 7 and 3. Essentially the statement says that any group of 7 items has more items in it than any group of 3 items. This characterisation of mathematical structures has a significant advantage in that it does not have to deal with the ontology of independently existing, eternal abstract objects. Additionally, it can account for both our knowledge of mathematics, and for the applicability of mathematics to the physical world, because mathematical structures are a feature of the physical world. This means that when we express mathematical truths we are describing an aspect of the world, so the task of accounting for the applicability of mathematics to the physical world becomes redundant. Similarly the question of how we can have knowledge of mathematics is no longer mysterious, since we can abstract mathematical truths from our experience of the world.²⁸

The greatest problem for eliminative structuralism is that it is not equipped to account for all of mathematics. It is a good theory for those portions of mathematics that are physically instantiated, namely those portions whose subject matter is what I refer to as basic structures. However it fails to account for derivative structures, which form the subject matter of much of mathematics. By contrast, epistemic structuralism uses something closer to Hellman's modal structuralism to account for mathematics

²⁸ Chapter 3 gives more detail about how we acquire mathematical knowledge.

that refers to derivative structures, those that have no physical instantiations. Chapter 3 explains how we acquire knowledge of derivative structures, namely by recombining the mathematical concepts that we have already acquired from our knowledge of basic structures, with which we become acquainted by abstraction from their physical instantiations. This means derivative structures are grounded in mathematical concepts that we have abstracted from the physical world, even though we are not acquainted with the instantiations of derivative structures. In some cases it turns out that we later discover derivative structures in the physical world. Again, the conditions required to say that we have found an instance of a mathematical structures may be debatable, but it is not a crucial question. The important thing is that we may find some derivative structures in the physical world (in which case we may want to reclassify them as basic structures), but there are some derivative structures that have no physical instantiations, or at least none that are accessible to us. These may be considered to be hypothetical or fictional structures.

The ontology and semantics of derivative structures is discussed in greater detail in Chapter 4. The present concern is to specify epistemic structuralism as compared to other versions of mathematical structuralism. For mathematical statements that refer to derivative structures, epistemic structuralism takes a stance that is very similar to modal structuralism. Statements about derivative structures are claims that *would be* true of such a structure *if* it were to exist. There is no claim made about the existence of the derivative structure. It may turn out that we find evidence in the physical world for its existence, but this evidence is not required in order to make the modal claim about the structure.

A non-mathematical analogy may be used to clarify these modal claims. Consider a woman who was adopted at birth and knows nothing about her natural parents except that her birth mother was red/green colour blind. She does not know if she has any other siblings. Since colour blindness is a recessive trait that is passed down on the X chromosome, a woman must have two X chromosomes that carry the

trait in order to be red/green colour blind. Thus we can infer that if the adopted woman has any full brothers, they will be colour blind. (Since males have only one X chromosome and if it carries the trait, then they will be colour blind. For the mother in question, both her X chromosomes carry the trait, so all her sons will be colour blind.) This is a claim about a person or people that may or may not exist, and we can make the claim without knowing whether he exists. If it were discovered that the adopted woman has a full brother, then the claim that he is red/green colour blind is a true claim. If it is somehow determined that she has no brothers, then the claim about her brother remains a counterfactual modal claim. In the absence of any evidence either way, the claim can be made and is a perfectly respectable and useful modal claim.

Mathematical statements about derivative structures are analogous to the claim about the sibling about whose existence we are unsure. They are modal claims that would be true if such a structure were to exist, but often we do not know if the structure exists. Indeed it is difficult even to formulate what it means for a mathematical structure to exist, and Chapter 4 discusses the difficulty of drawing the line between existent and non-existent mathematical structures. However the question of existence is not the present concern; rather it is the nature of mathematical entities and the claims we can make about them. We can make mathematical claims that are hypothetical modal assertions without any commitment to the existence or non-existence of the mathematical structures in question.²⁹ A mathematician's primary concern is what she can deduce from certain axioms, rather than the origin of the axioms. Chapter 3 explains that mathematics is grounded in our experience of the world, which explains its objectivity, however it is not necessary to commit to the existence of every mathematical structure. Chapter 4 also explains how bridge laws can be used to determine the objective truth or falsity of statements that refer to

²⁹ The status of these claims is explained in detail in Chapter 4.

structures that may or may not exist in the non-mental physical world. These statements, which refer to derivative mathematical structures, are modal structural claims since they refer to possibly uninstantiated mathematical structures. We can make such claims without being concerned about the existence of the derivative structures in question.

Although epistemic structuralist claims about derivative structures are very similar to modal structural claims, there is a significant distinction. Epistemic structuralist claims – even about hypothetical mathematical entities – are grounded in our experience of the physical world. Hypothetical mathematical entities were derived by a process of abstraction from structures that we perceived, so we may say objectively true things about them (using bridge laws, see Chapter 4). So although these types of mathematical statements have a lot in common with modal structuralist statements, the epistemic grounding in our experience of the physical world is built in.

Epistemic structuralism, as a hybrid between eliminative structuralism and modal structuralism, has many benefits over other varieties of mathematical structuralism. One of the strongest of these is that it provides an account of mathematical knowledge that is pragmatic and reflects actual mathematical practice without assuming a mysterious faculty of mathematical intuition. This advantage may be attributed to the method of approach characteristic of epistemic structuralism, which is an epistemically motivated theory. The primary concern of most mathematical structuralists is to explain the nature of mathematical entities in such a way as to avoid the dilemma identified in Benacerraf's [1973]. Mathematical structuralism is a theory that takes mathematical entities to be structures, it is a theory of mathematical ontology and hence it is unsurprising that most varieties of mathematical structuralism come from an ontological approach. Hellman's concern, for example, is primarily ontological. He seeks to give the best possible account of what mathematical entities are. Although this view provides solutions to many ontological problems, it does not easily explain our knowledge of modal structures

nor the mechanism by which we refer to them. By contrast, epistemic structuralism arose as part of an explanation of the nature and justification of mathematical knowledge. It is in large part an epistemically-derived view rather than a purely ontologically motivated one. This accounts for its lack of ‘neatness’: i.e. for the distinction between basic structures as *in re* universals, and derivative structures as modal structures. This distinction arises in the course of explaining how we acquire mathematical knowledge. It is not as neat as *sui generis* structuralism or modal structuralism which have a uniform characterisation of all mathematical structures. However the epistemic derivation of epistemic structuralism makes it the theory that is best able to explain our knowledge of mathematical truths and the nature of these truths. Hence the ontology that arises has a strong advantage, and there are other benefits that epistemic structuralism has over other varieties of mathematical structuralism.

Recall that Shapiro’s *sui generis* structuralism reifies places in structures, considering them to be independent objects, and that this move undermines the accuracy of his analogy with places in organisational structures. This objection to Shapiro’s version of structuralism highlights one of the strengths of epistemic structuralism. Epistemic structuralism takes the physical instantiations of structures to be ontologically primary. Hence we can refer to the Wallabies’ fly half without considering it to be an eternal abstract entity that exists over and above all the people who have held that position. We have an abstract idea of (or a theory about) the Wallabies’ fly half that is both ontologically and epistemically dependent on the people who have occupied the position. Similarly our idea of the natural numbers is entirely dependent, both epistemically and ontologically, on our experiences with physical systems that instantiate sub-segments of the natural number structure.

Another problem facing *sui generis* structuralism was the identity problem, in which structuralists were compelled to identify the places in structures that were invariant under non-trivial automorphisms. For example, in the case of the integers,

sui generis structuralism led to the claim that $1 = -1$, which is clearly an unacceptable conclusion. Recall that this problem arose from the belief that mathematical entities, places in mathematical structures, can be individuated only by their essential properties, which are those that they hold in relation to other places in the same structure. For example the number 9 can be individuated by its property of being the square of 3, but not by its property of numbering the planets of our solar system. The reason for this assumption is that *sui generis* structuralism is a form of Platonism, which takes mathematical structures to be eternal abstract entities that are instantiated in physical systems but whose abstract conceptions are ontologically prior. Hence a mathematical entity must be able to be individuated by its purely abstract features, or its structural properties.

For epistemic structuralism, it is the physical instantiations of mathematical structures that are primary, both ontologically and epistemically. This means the epistemic structuralist is not compelled to reject the possibility of individuating mathematical entities by their non-structural properties. The number 9 can be individuated successfully by the property of numbering the planets in our solar system. This individuation is contingent, for if another planet were discovered then this property would specify the number 10 rather than 9. However the possibility of this occurrence does not prevent us from using this non-structural property successfully to individuate 9.

Thus epistemic structuralism does not fall victim to the identity problem laid out by Keränen. Although the integer structure is invariant under automorphism, the epistemic structuralist is not committed to identifying 1 and -1 since she has other properties than structural properties at her disposal for the purposes of individuating numbers. While the *sui generis* structuralist can only appeal to properties such as being half the magnitude of its successor, the epistemic structuralist can also use properties such as numbering the chairs in this room at this time. Since the non-structural properties were how we first came to acquire knowledge of numbers, these

properties are not considered non-mathematical as in Platonic accounts of mathematics. For the epistemic structuralist, mathematical structures are a feature of the physical world, they are built into our experience of the physical world. Our theories (or conceptual schemes) about mathematical structures are just tools to help us make sense of the world, they are not ontologically privileged, ideal versions of mathematical structures. Hence non-structural properties may be used to individuate mathematical entities just as structural ones can.

Individuation is likely to be contingent on the world being a certain way, but that does not make it an unsuccessful or useless individuation. While the number of chairs in this room may change, I can still successfully specify the number 1 using its property of numbering the chairs in this room at this time. A Platonically-based account of mathematics would take affront at the contingency of such an individuation for mathematical entities, but this is a bias grounded in the dogma of mathematics as a set of abstract eternal truths that are immune to the contingencies of the physical world. This dogma has led to Benacerraf's dilemma concerning mathematical truth and knowledge. An empirical grounding of our mathematical knowledge along with a modal structural interpretation of derivative mathematical truths provides both an explanation of mathematical knowledge, which is its primary motivation, as well as an account of the nature of mathematical entities; however it does so without sacrificing mathematical truth and reference and its applicability to the physical world.

CHAPTER 2: PLATONISM AND STRUCTURALIST ONTOLOGY

Mathematical entities are the things that constitute the subject matter of mathematics, they are what mathematics is *about*. Such entities include numbers (of all varieties), lines, points, planes, curves, knots, triangles, sets, groups, manifolds. For the purposes of simplicity and clarity, my discussion will involve primarily the natural numbers, although other mathematical entities are used to illustrate particular points. The reason for discussing primarily the natural numbers is that they are among the most basic of all mathematical entities and most others can be constructed in terms of them. Furthermore, the simplest of all mathematical statements, namely those that have these basic entities as their subject matter, are the ones that are known to almost all people and thus may be uncontroversially claimed to constitute knowledge. This chapter begins by evaluating a traditional theory of mathematical entities, namely Platonism. Since historically this has been the dominant theory of mathematical ontology, this chapter provides a detailed account of the epistemic structuralist view of mathematical entities and contrasts it to the Platonic approach.

2.1 Platonism and Realism

Perhaps the most traditional treatment of questions about the nature of mathematical objects is Platonism. This is the theory that mathematical entities exist independently as Platonic Forms, and are located distinct from the spatio-temporal world we inhabit. According to the Platonist, when we do mathematics we refer to these Forms, and things in the world may approximate the Forms. For example, there is a Form that is a perfect sphere, which may be defined in mathematical terms. There exist many physical copies of this perfect sphere, all of which are imperfect approximations to its perfection, which can only be captured in mathematics. Similarly, the number 3 exists in the Platonic realm as a Form, of which all instances of 3 things, for example 3 chairs, are imperfect copies (imperfect in the sense that they

do not consist of the number 3 in its purest form, since the number cannot be separated from the ‘chairness’ of the objects). In this sense, the natural numbers are attributed some independent existence over and above their exemplifications in the physical world.

Platonism has long been the dominant variety of realism with respect to mathematical objects. Often it is assumed that if mathematical entities are real then they are Platonically abstract, existing outside of space and time. Structuralism, however, provides a way of grounding them in the physical world.

One of the strongest arguments for realism with respect to mathematical entities is the Quine-Putnam indispensability argument. This argument is summarised by Colyvan as follows:

- (1) We ought to have ontological commitment to all and only the entities that are indispensable to our best scientific theories;
- (2) Mathematical entities are indispensable to our best scientific theories;

Therefore:

- (3) We ought to have ontological commitment to mathematical entities.¹

This is a valid logical argument that provides a strong case for the existence of mathematical entities independent of human thought. Naturalism and holism together provide support for the first premise, and I accept this since I am working with a broadly Quinean naturalism as outlined earlier and hold a broadly reliabilist epistemology that is compatible with holism.²

Field³ argues against the second premise of the indispensability argument, as he explains why mathematics is useful to science but not indispensable. He attempts

¹ Colyvan [2001], p. 11.

² See Colyvan [2001] for a more detailed evaluation of the Quine-Putnam indispensability argument.

³ Field [1980].

to demonstrate that we can do away with numbers altogether in our scientific theories, and that such a nominalised theory would still be reasonably attractive. This type of nominalistic view sees mathematical entities as instrumental constructions in order to explain their applicability in the physical world, but pragmatic because we do not *have* to use numbers to describe the world. However Field's project and others like it have been criticised for yielding unintuitive consequences (which are parallel to problems with nominalistic theories of universals) and also for technical inadequacies.⁴

While antirealist accounts of mathematical objects may have an advantage with respect to epistemic issues, this advantage comes at the great cost of making the truths of mathematics dependent on us and thus subjective,⁵ and no longer a part of objective reality. If mathematical objects are not real then a standard semantics cannot account for the objectivity of mathematical truths. Since the subject matter of mathematics has no independent existence, mathematical truth becomes subjective, since their truth value will depend on us or on the process by which we came to know them, rather than on independently existing entities. This leads to difficulties when antirealists have to explain why mathematical formulae can correctly predict the orbits of planets, the trajectories of bullets and so on, because in order to account for the applicability of mathematical truths they must resort to explanations that yield unintuitive results.

These concerns are much less problematic for realists, since they consider the statements of mathematics to be true in virtue of properties of independent mathematical entities. If the truths of mathematics are independent of us and objective, then the problem of applicability is less serious. However Platonism now

⁴ See introduction to Irvine [1990] for further discussion of the challenges facing Field's project.

⁵ As explained in the Introduction, I am following Putnam in associating objectivity of mathematical truths with mathematical realism.

faces a serious epistemological problem, which Brown referred to as the Problem of Access. As we saw in the previous chapter he and others have been unable satisfactorily to solve this problem, and it remains a fatal flaw of Platonism. In this section we shall see that an epistemic version of structuralism solves the problem of access by locating mathematical entities primarily in the physical world, rather than a Platonic realm. This approach maintains a natural explanation of why mathematics is so readily applicable to the physical world, since it arises from a physicalist account of how we acquire knowledge of mathematical entities.

Although Steiner distinguishes between ontological Platonism and epistemological Platonism, the two can be captured by the following six points from Irvine, which together describe a Platonic position:

- (i) Mathematical entities exist independently of human thought and of our ability or inability to obtain knowledge of them;
- (ii) Such entities are non-physical, existing outside space and time;
- (iii) Statements of mathematics possess truth-values, again independent of human thought and of our ability or inability to obtain knowledge of them;
- (iv) Such statements obtain their truth-values as a result of properties of mathematical entities (and not as a result of, *e.g.* properties of natural or formal languages, *etc.*)
- (v) It is possible to refer unequivocally to such entities; and
- (vi) to obtain knowledge of them.⁶

Ontological Platonism comprises points (i)-(iv), while epistemic Platonism involves the faculty of mathematical intuition that makes (v) and (vi) possible. Epistemic Platonism will be considered in the following chapter, and now ontological Platonism will be evaluated. Point (ii) is the most problematic claim, and this is the point that distinguishes Platonism from the form of structuralism which I am advocating. The

⁶ Irvine [1990], pp. xix-xx.

reason it is such a problem, as we shall see in the section on epistemology, is that it makes mathematical entities causally inert and so it is very difficult to explain how we can have knowledge of them, or even refer to them. While Steiner maintains that it is possible to accept the ontological aspects of Platonism and reject its epistemic implications, we shall see that in order to explain mathematical knowledge as we possess and acquire it, it makes much more sense to locate mathematical entities within the natural world.

Brown is a committed modern day Platonist, adhering firmly to his belief in the existence of Platonic objects. He claims that “Mathematical entities can be ‘seen’ or ‘grasped’ with ‘the mind’s eye,’”⁷ and while admitting that this is not really an adequate explanation, concedes that it is all he has. Most of his more successful arguments for Platonism can be adapted to be arguments for structuralism, even epistemic structuralism, because they are really arguments for mathematical realism. Most of the theories he argues against are species of formalism, conventionalism or instrumentalism, all of which reject the notion of an objective mathematical reality and seek to ground mathematical truth in mathematical language. He uses pictures and alternative representations for the same mathematical structures to argue that they share a common, objective mathematical reality.⁸ For example, knots can be described in a variety of notations that bring out different aspects of the knot. A more familiar example is a parabola, which can be represented by the formula $y = x^2$ as well as by a picture. While these different representations of mathematical entities do suggest that there is something ‘behind’ the representation that we are getting at (or trying to specify) when we use mathematical language, this works equally well as an argument for epistemic structuralism as it does as an argument for Platonism. There is nothing

⁷ Brown [1999], p. 13.

⁸ Chapter 7 contains a section entitled ‘Brown’s Argument from Notation,’ which provides a more detailed analysis of this argument.

in this style of argument to support point (ii) above. His arguments support only the existence of an objective mathematical reality. Structuralism is even more strongly supported by this type of argument, since it seems more plausible that when we make mathematical references we are specifying mathematical structures of which we have had some experience and with which we can interact, rather than Platonically abstract Forms that exist outside of space and time. This is supported by one of the main problems for Platonism, the problem of causal interaction, which is discussed in Chapter 3.

2.2 Structures are in the World

In giving an account of how truths about Platonic Forms can be used to describe the physical world, Brown explains that when mathematics is applied, we use a mathematical structure to describe an aspect of the world to which it is similar:

... a mathematical representation of a non-mathematical realm occurs when there is a homomorphism between a relational system **P** and a mathematical system **M**. **P** will consist of a domain D and relations R_1, R_2, \dots defined on that domain; **M** similarly consists of a domain D^* and relations R^*_1, R^*_2, \dots on its domain. A homomorphism is a mapping from D to D^* that preserves the structure in the appropriate way.⁹

This account of how we use mathematical structures to describe a “non-mathematical realm” requires that the relational system **P** possesses structural relations that are homomorphic to the mathematical system **M**. If the structure in question can be preserved via a homomorphism between the physical system and the mathematical system, then both systems must possess that structure. If the mathematical system consists of a domain and relations defined on the domain, as does the physical system, then the mapping from the physical domain to the mathematical domain must preserve the structure that the two systems have in common.

⁹ Brown [1999], p. 47.

Since the physical system and the abstract mathematical system have the same structure, then what they have in common, namely this common mathematical structure, must be the essential mathematical features of both. The relational system \mathbf{P} displays certain mathematical attributes, characterised by a domain D and relations R_1, R_2, \dots defined on D . The mathematical features of \mathbf{P} work in the same way, or are analogous to, the mathematical features of the system \mathbf{M} . Since the mathematical structure of \mathbf{M} as characterised by the domain D^* and relations R^*_1, R^*_2, \dots is homomorphic to that of \mathbf{P} , this means that both the physical system \mathbf{P} and the mathematical system \mathbf{M} have the same mathematical structure. The structure is captured both by D and R_1, R_2, \dots , as well as by D^* and R^*_1, R^*_2, \dots . Both \mathbf{M} and \mathbf{P} are instantiations of a common mathematical structure.¹⁰ We shall see that \mathbf{M} is actually nothing more than certain mathematical aspects of \mathbf{P} , considered independently of \mathbf{P} 's other (mathematical or non-mathematical) aspects.

In fact what Brown has given us here, instead of a description of how abstract mathematical truths describe a non-mathematical world, is an account of how we isolate a mathematical structure from a physical system. The process he has described as mathematical representation is really just the observation of a mathematical structure within a physical system. By specifying a homomorphism from \mathbf{P} to \mathbf{M} we have determined which are the mathematically significant aspects of the system. For example, if \mathbf{P} is a row of telegraph poles and \mathbf{M} is the natural number structure, then by finding a homomorphism from \mathbf{P} to \mathbf{M} we decide that the relations we want to retain in the move from the physical system to an uninstantiated one are just the sequential ordering and number of the physical objects, and we can leave out height, shape, woodenness, colour and so on. However if \mathbf{M} is a long, thin cylinder then we

¹⁰ A discussion of the formal sciences in a later chapter will provide further evidence that the essential mathematical features are preserved in the shift from concrete instantiations of mathematical structures (\mathbf{P}) to their abstract characterisations.

can leave out the sequential ordering of the telegraph poles as well as their number and also their woodenness and colour and instead, the homomorphism we seek will preserve the shape, height and circumference of the objects in question. This is how we abstract, in the sense of simplifying or leaving out, in order to get from a physical system which exemplifies a mathematical structure, to the mathematical structure itself, considered in isolation.

If the physical relational system **P** did not possess mathematical attributes, or did not exhibit a mathematical structure, then we would not be able to find a homomorphism between it and a mathematical structure. The fact that **P** consists of a domain with relations defined on it which is homomorphic to a mathematical domain tells us that the physical system itself instantiates this mathematical structure. The process of moving to the mathematical domain consists just of leaving out the non-mathematical aspects of **P**, so what we are left with (namely **M**) must still be physical since it is a part of the physical relational system. **M** is in fact nothing more than the purely mathematical aspects of **P**, it is not an independently-existing additional entity over and above the physical system. This is an important point, since it renders epistemic structuralism immune to the Third Man Argument often levelled against Platonism. Brown claims that “structuralism sees structures right in the non-mathematical world itself,”¹¹ but he has shown us himself that there are mathematical structures in the world, and thus the world cannot be non-mathematical.

2.3 Epistemic Structuralist Ontology

Some consider structuralism to be a variation on Platonism. This is misleading, for although *sui generis* structuralism may be considered a Platonistic variety of structuralism, there are several other varieties that are not Platonistic. Under

¹¹ Brown [1999], p. 58.

epistemic structuralism, mathematical objects have no existence independent of the physical world corresponding to the existence of the Forms in the Platonic realm. Epistemic structuralism is the view that most mathematical entities exist in the physical world, and are presented to us empirically as structures embedded in the physical world.

Through abstraction from our perceptions we gain knowledge of the structures themselves, even though these structures do not exist independently of their physical instantiations. We can think of and refer to them in the abstract (i.e. in isolation), but their abstract existence is only in our minds and their physical instantiations are ontologically primary.¹² To take the example of the natural numbers, we may perceive 3 dogs, 3 telegraph poles, 3 chairs, 3 houses and so on, and from this we abstract the common mathematical structure present in all these groups of objects, namely the number 3. We can also perceive some groups of 5 objects and realise that the numbers 5 and 3 are a part of the same structure and, as such, stand in a certain structural relation to each other. For example, we can add 2 more chairs to the original 3 and get 5 chairs. Then we can consider the numbers themselves independently of the objects, as we realise that adding 2 ‘somethings’ to 3 ‘somethings’ will *always* give 5 ‘somethings’. In other words, we have discovered the mathematical truth that $2 + 3 = 5$. This claim is a generalisation from many observations, and as such it can be considered an empirical claim. However it is stated in the abstract, leaving out the non-mathematical features of the objects it describes, so it is more useful and descriptive, and less cumbersome than many empirical claims. In this way, abstraction is the process of discovering mathematical structures in the physical world and expressing them as statements about abstract structures so that we may reason about

¹² Some mathematical structures do not, as far as we know, have physical instantiations. Chapter 4 deals with issues of ontology and epistemology pertaining to such structures.

them without having to consider all the extra-structural features of any particular instantiation.

This way of expressing mathematical statements has led many people to the misconception that the subject matter of mathematics is in fact Platonically abstract. There is an important difference between mathematical entities being abstract in the sense of existing independently of the physical world, outside of space and time, and *talking about* mathematical entities in the abstract. I wish to argue against the former, but maintain that we do the latter all the time. The subject matter of mathematics, or what we refer to when we do mathematics, exists in reality;¹³ by this I am advocating a form of realism. Mathematical entities are structures which are a part of the physical world, either explicitly, or because they are grounded in the physical world. When we recognise these structures as recurrent patterns, we find it useful to make statements about them in the abstract and reason about them as if they were abstract. However this way of talking about mathematical structures does not entail that they are abstract entities. We can speak of electrons in a conceptual or theoretical way, but they are still physical entities.

A good way of explaining how we might appear to be speaking abstractly about mathematical structures is by analogy to our concept of any other physical thing. My cat is asleep on my desk, she is a physical substance that exists independently of my observing her. I can look away or close my eyes and have a concept of her, an image of her in my mind. Using this idea that I have of her, I can refer to her and say true things about her, but this does not mean that the cat is in my head or in any way dependent on my mentality for her existence.

¹³ This may not be strictly true of all derivative mathematical structures, since some may be hypothetical entities. However most mathematical entities are real, and even the seemingly uninstantiated ones are grounded in the physical world, so we can make objective claims about them.

To develop the analogy further, I can observe the cat on my desk and also think of my other cat who is not in the room. By abstracting away the features that they do not have in common (size, colour, sex, and so on), I am left with an abstract concept of a cat. This concept is instantiated in both my cats, and if I see another cat outside I can recognise that it instantiates the same concept. My concept of a cat comes from my observations of cats and I would not have this concept had I never perceived a cat.¹⁴ The cats themselves may be thought of as thick particulars: particular substances that are ‘thick’ with properties, one of which is the property of being a cat. This property may be thought of as ‘thin’ in the sense that nothing can be a cat and have no other properties: there is no cat without size, shape, length of fur, colour, and so on. We can have a concept of the thin property, which is our concept of a cat. This does not mean we have any actual cats in our heads, nor that “cat” is a non-physical independently existing abstract entity. It is just an abstraction of the property from all the thick particulars that possess it. We do not have the thin property in our heads either, we just have a mental image of it.

Mathematical properties too may be thought of as thin particulars. The ‘four-ness’ of the pens on my desk and the ‘four-ness’ of the chairs around the kitchen table are a structure that these two physical systems have in common, and I can think of it in isolation from these instantiations. This abstract conception of a mathematical structure enables me to recognise other instances of it, if I see four trees I can tell that their number is the same as the number of my pens and chairs. Similarly, once I have obtained an abstract conception of a cat from my two cats, then if I see another cat I can tell that it too possesses the property of being a cat. The next section of this chapter makes a comparison between mathematical structures and other properties.

¹⁴ If I had never perceived a cat, someone could have explained to me what a cat is like. In this case I would have a concept of a cat, but it would be a very hazy and incomplete concept.

By referring to an abstract structure I am *not* postulating an abstract realm in which there exist pure, uninstantiated versions of all mathematical structures, nor do I wish to suggest that mathematical structures have an independent existence over and above their physical instantiations. The physical manifestations of mathematical structures are primary, and we have the ability to consider these structures independently of the non-structural properties of any given instantiation, however this does not magically infuse mathematical structures with an independent existence.

For the structuralist, individual mathematical objects do not have an independent existence, there is no singly-occurring entity which corresponds to the number 3. Rather, the number 3 exists only in virtue of its position in the natural number structure, and it is present in each group of 3 things. When we have 3 chairs, we may isolate the properties of this group of objects which belong to their numerical structure, as well as those properties that do not. For example the property of being half the magnitude of 6 is a property of 3 since all groups of 3 objects will share this property, however the property of being made of wood is not. Thus we may identify the numbers (and other mathematical entities) by their structural properties, and this allows us to reason about them in the abstract. On a metaphysical level, however, numbers only exist in virtue of being a part of a structure (the natural number structure) which occurs in a vast multitude of instances in the physical world.

It is worthwhile to clarify the relation between the structure as a whole and elements of the structure. The wider structure is distinct, but not separable, from the individual elements; for example the individual numbers make no sense outside the context of the number structure as a whole, similarly we cannot have the whole number structure without its constituent numbers. Mathematical items only exist in virtue of being a part of a structure, and are ontologically dependent on their structural context.

It is useful to recall the notions of essential properties and structural properties. An essential property is one that (in part) defines the object in question, and which is a

necessary condition of being that thing; if this property is removed, then the nature of the object would not be the same. A non-essential property is a contingent property, one that can be removed without altering the nature of the object in question. Structural properties are relations that hold between objects in virtue of their relative positions in the same structure. A structural property is given in terms of the wider structure of which the object is a part. A non-structural property is one that makes no reference to the structural context; it belongs to an object as an individual, independently of the rest of the structure of which it is a part.

In the case of the numbers, their structural properties and essential properties coincide. For example the number 9 has the structural property of being the square of 3, structural in the sense that it is given in terms of another element of the natural number structure. This is an essential property, because if the number was not the square of 3 then it would not be the number 9. It is part of the definition of 9 that it is the square of 3, it is a property that is *essential* to being 9. No number can give up any of its structural properties (the properties it has in relation to other nodes in the structure) and still be the same number. Thus the structure and its elements or nodes are ontologically intertwined, they cannot be separated from each other.

2.4 Mathematical Objects as Universals

It is useful to note that an epistemic structuralist account of mathematical entities has many parallels with a natural-realist account of universals, based on an Aristotelian approach to the problem. We saw above that the number 3 is present in all collections of 3 distinct items,¹⁵ which is similar to the way the colour blue¹⁶ is

¹⁵ If we observe a collection of 3 items, 3 is not the only number that is present. The number of atoms in 3 rocks, the number of limbs of 3 people, there are many numbers exemplified in any collection of items. However, this does not change the fact that when we see 3 people, 3 rocks and 3 trees, we can identify a common pattern that is exemplified in each of these collections of items. Furthermore, we can observe that this pattern is a node in the same structure as the common pattern that we observe in 5 fingers, 5 pencils and 5 spoons.

present in all blue things. A natural realist account of universals takes them to be real things in the world, rather than Platonic forms or nominalistic entities. To follow the Aristotelian conception, universals exist only in their instantiations in physical objects. Physical objects are ontologically prior to the properties and relations they possess. There is no postulation of independently-existing abstract entities, a universal is nothing more than a recurrent property or relation and does not consist of anything over and above its physical instantiations. It is present in each of its instantiations, and we may speak or think of it as an abstract entity by separating it from its instantiations. However this does not endow it with independent existence. We can observe that a fire truck and a pen have something in common, namely that they are coloured in the same way. Having made this observation, we are able to consider the common colour (red) in isolation, independently of its physical exemplifications. In so doing we have grasped the universal property “red” which is a real thing that exists as a part of the physical world. We have not intuited an object that exists in an abstract realm outside of space and time; we have just gained a concept of a property that exists in the physical world and is present in many instances.

Treating the natural number system (and other mathematical entities) as structures (or patterns) that recur in the physical world puts them in the same ontological category as Aristotelian universals (i.e. properties and relations). Both are real, non-abstract things which exist in many instances in the physical world. Since both are present in many distinct instances, we can speak of individual numbers in the same way as we can refer to properties independently of a particular instantiation. For

¹⁶ For the purposes of this discussion I will not draw a significant distinction between primary and secondary qualities, for it will cause me to stray from my focus. A property such as ‘blue’ is generally considered not to be a true property of the object that we perceive as blue, since it requires some interaction from the perceiver in order to appear blue. However in this discussion I merely wish to draw a comparison between Aristotelian universals and mathematical structures, so a mention of a secondary quality as a universal should not be taken as an indication that I identify both primary and secondary qualities as properties of the object in question. Rather, if I refer to ‘blue’ as a universal it can be construed as ‘the physical property of the object which causes it to appear blue to human perceivers.’

example statements such as “65,215 is a large number” and “red is a bright colour” refer to numbers and properties as things in themselves. Both of these statements are general claims stating truths that hold for all instances of the number or property; namely that all groups of 65,215 objects are large groups, and that all red objects are brightly coloured.¹⁷ When we speak abstractly about mathematical entities we are speaking in the same mode as when we refer to properties and relations as things in themselves. Although neither category of entities has an independent existence, it is often useful and much less cumbersome to refer to them as things in themselves. In both cases this way of speaking is nothing more than a useful tool, it does not endow the objects in question, which exist in recurrent instances in the world, with an independent existence over and above their physical instantiations.

General claims about mathematical entities or universals are extremely useful because they allow us to generalise across all instantiations of the structure or universal in question. Indeed, without the ability to make general claims, mathematics would be extremely cumbersome and mathematical progress much more difficult, if not impossible. However such general claims are misleading because of their treatment of mathematical entities, a feature of the physical world, as abstract entities. The distinction between Plato’s and Aristotle’s conceptions on universals is crucial here: Plato considered that all physical exemplifications of universals were imperfect copies of Forms, which were independently existing abstract entities, ontologically prior to the physical world. By contrast Aristotle held that it was physical objects that were ontologically prior, and that properties and relations only existed in virtue of being instantiated in the physical world. Epistemic structuralism takes the same position as Aristotle with respect to mathematical objects: namely that mathematical

¹⁷ These are, of course, relative claims and thus depend on context. If we are speaking of atoms in a glass of water, or inhabitants in a country, then a group of 65,215 may be considered a small group. Similarly if we are comparing a red object with fluorescent coloured objects then the red object might appear dull rather than brightly coloured.

structures are a part of the physical world, they only exist in virtue of being instantiated in physical systems, and although we can consider structures independently of their physical exemplifications using a process of abstraction, they do not exist as independent abstract objects.

A natural question that arises from this view of mathematical entities is: what about mathematical entities that have no physical instantiation? How can we make claims about mathematical entities like 28-dimensional spheres? This is the subject of the next two chapters, which explain the process of abstraction and the relationship between basic and derivative mathematical structures.

CHAPTER 3: ABSTRACTION AND MATHEMATICAL KNOWLEDGE ACQUISITION

This chapter is concerned with mathematical knowledge acquisition. The structuralist account that I am advocating uses abstraction from perceived instances of mathematical structures to arrive at mathematical beliefs. In this chapter I will outline this structuralist account and argue against the dominant Platonistic account. Further chapters explicate the precise nature of mathematical knowledge acquired in this way.

3.1 Structuralist Epistemology

McGee points out that for realists about mathematical entities, it is “a particularly urgent problem”¹ to explain how we can have knowledge of the referents of mathematical terms, or even refer to them. This problem is very serious if mathematical entities are considered to be Platonic (or abstract in the Platonic sense), while for the epistemic structuralist a fairly natural explanation is available. The structuralist account of how we come to know the truths of mathematics is partly an empirical story. In his paper “Mathematical Knowledge and Pattern Cognition”² Michael Resnik provides an account of how we recognise patterns in various physical phenomena, which then allows us to abstract the relevant mathematical structure. This is how the structuralist claims we get knowledge of mathematical entities, which are patterns that are repeatedly exemplified in the physical world. Resnik gives the following example:

Consider this sequence of dots

.

¹ McGee [1997], p. 36.

² Resnik [1975].

and let us try to explore it from a mathematical perspective. If we were impressed by the immediate succession of one dot after another we might make statements such as

- 1) no dot has more than one immediate successor
- 2) if one dot succeeds another then the latter does not succeed the former
- 3) there is a dot which succeeds no dot and every dot but it succeeds a dot
- 4) there is no dot between a dot and its successor

Or if we were impressed by the ordering of the dots we might come up with these other statements:

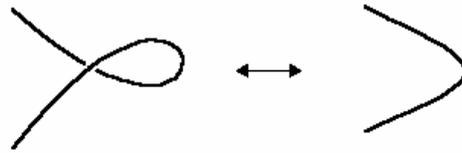
- 5) if one dot comes before a second and the second before a third then the first comes before the third
- 6) if one dot comes before another then the second does not come before the first
- 7) given two distinct dots one comes before the other
- 8) given any sub-sequence of dots there will be one in the sub-sequence which comes before the others.³

The axioms he gives can be directly observed by looking at the dots, however it is also possible to observe the same axioms in a row of houses, a stack of coins or any other sequential collection of items. From these observations we gain knowledge of the structure of a finite sequence. Resnik goes on to explain that we can then extend our knowledge of this abstract structure to arrive at the idea of an infinite sequence. Later in this chapter I will look more closely at how we do this, and we shall see that this extension is a species of abstraction.

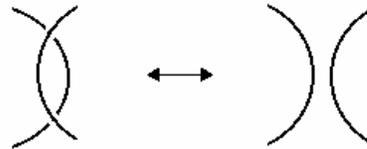
A second example of mathematical truths we can observe is taken from knot theory. If we manipulate a loop of string we soon discover that whether there is a knot in it or not, there are certain moves we can perform which do not alter the knottedness of the string. If there is no knot in the string then we can observe that the following manipulations and their inverses do not produce a knot (and if there is already a knot, they do not change it):

- (1) twisting a segment of string:

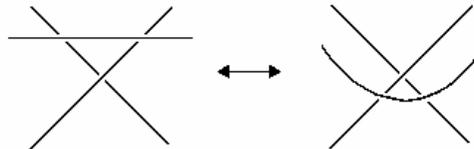
³ Resnik [1975], pp. 33-34.



(2) crossing one segment over another:



(3) moving a segment from one side to the other of an intersection:



We can observe directly from interacting with a piece of string that these manipulations do not knot an unknotted piece of string, nor do they untie or change any knots that are already present. These are the three Reidemeister moves under which knots are equivalent, and the fact that knots are invariant under these moves is observable by playing with a piece of string. If we use a different material, such as a heavy rope or an electrical cord, we will still be able to observe that they are subject to the same knot-preserving moves. This is another example of how our beliefs about mathematical structures originate in the physical world.

The crucial point in these examples is that they show how we have gained knowledge of a mathematical pattern instantiated in physical objects, even though the structure studied in mathematics does not consist of objects. A useful distinction between two ways of considering places (or nodes) in a structure is made by Shapiro,⁴

⁴ Shapiro [1989].

and is a good way to illustrate the structuralist conception of mathematical entities. This distinction was discussed in an earlier chapter. Recall that Shapiro calls these two ways “places-are-offices” and “places-are-objects.” His paradigm example concerns the natural numbers and when he refers to the relations between the places in a structure, it is clear that he is thinking of relations such as the less-than relation or succession. However, his two ways of considering mathematical objects can be applied to other structures as well.

Shapiro’s distinction arises from his conviction that mathematical objects are no more than places in structures (as in Benacerraf’s view), yet they are referred to in mathematics as if they are objects in their own right. The places-are-offices view treats mathematical entities as nothing more than positions in a structure, which are then filled by actual objects that are part of a system which exemplifies the structure. When these places are treated (grammatically) as objects in their own right, the places-are-objects perspective is being used. Statements of this sort are about the structures themselves and are independent of any instantiations of the structure, but will hold true in any system which exemplifies that structure. When we observe a specific row of dots on a page, we are using the places-are-objects perspective, since we observe and describe relations between *those* particular dots. The process of abstraction allows us to move to the places-are-offices perspective, as we realise that the same relations hold for any sequential collection of items.⁵

3.2 Platonist Account

Gödel is one of the most significant proponents of Platonism in mathematics. He advocates a kind of mathematical intuition by which we perceive the truths of set

⁵ Of course, some observed truths would vary between instantiations if they depended on the number of items. However successor relations and principles of ordering would always be the same.

theory (as distinct from geometry, which he considers to be physically-based). The reason he thinks we must have this faculty of mathematical intuition is that the objects of set theory “clearly do not belong to the physical world,”⁶ and yet “the axioms force themselves upon us as being true.”⁷ In other words, his view is that we seem to know the axioms of set theory with certainty, and yet the objects that these axioms describe are non-physical so they could not have been given to us through the five standard senses. He supports the idea that knowledge can come to us by means other than the senses by claiming that even our thoughts or ideas relating to physical objects contain more information than mere sense data. Gödel claims that the information given to us through this faculty of mathematical intuition is not in fact subjective, but is an aspect of objective reality. This is clearly a Platonistic view with respect to set theory, namely that the objects of set theory are non-physical yet objective⁸ (i.e. they could be described as Platonic Forms), and we have some faculty for perceiving truths about these objects with which we have no physical interaction.⁹ It will become evident that specifying the nature of this faculty of mathematical intuition is the greatest obstacle for Platonism with respect to mathematical knowledge.

One of the most serious problems with a Platonic account of mathematical entities is that we do not seem to interact causally with the Platonic realm, hence it is difficult to explain how we can have knowledge of the Forms. Benacerraf gives an account of this problem in his famous paper “Mathematical Truth,”¹⁰ which highlights the fundamental tension between semantics and epistemology in finding a satisfactory account of mathematical truth. Benacerraf argues that a causal theory of knowledge and reference taken with a Platonist account of mathematical entities is unable to

⁶ Gödel [1947], p. 483.

⁷ Gödel [1947], p. 484.

⁸ Refer to Chapter 2 for a thorough definition of Platonism with respect to mathematical objects.

⁹ This is clearly a Platonistic view because it encompasses points (i), (ii) and (vi) of the definition of Platonism given in Chapter 2, which deals with ontological issues.

¹⁰ Benacerraf [1973].

explain mathematical knowledge, since we can have no causal interaction with abstract entities that exist outside of space and time. However a standard theory of truth and reference forces a Platonist theory of mathematical entities, so we remain unable to explain how we can have knowledge of objects that are outside our causal realm.

This problem is also brought out by Steiner in his book *Mathematical Knowledge*,¹¹ in which he looks at various ways that we may come to know mathematical truths according to some of the dominant philosophies of mathematics. He distinguishes between ontological and epistemological aspects of Platonism, emphasising that it is possible to maintain consistency only while holding one and not the other of these two doctrines. Ontological Platonism asserts that the subject matter of mathematics consists of infinitely many objectively real entities, which exist in a Platonic realm distinct from the space-time manifold that we occupy. Although Steiner holds this doctrine to be true, he finds epistemological Platonism more problematic. This is the view outlined above as advocated by Gödel, namely that we have a faculty by which we intuit the truths of mathematics. This faculty is supposed to be analogous to sense perception, yet additional to the five senses. Steiner, an ontological Platonist, is unable to accept epistemological Platonism because of what may be a fatal flaw of the theory, namely that it fails to explain how we are able to interact with the ideal mathematical objects in order to gain knowledge of them. According to Platonism, mathematical objects are outside all causal chains and outside space and time. Thus there is for the epistemological Platonist the onus of accounting for how we may come to know truths about these entities when there is no causal interaction.

¹¹ Steiner [1975].

One way of avoiding the problem of causal interaction of epistemological Platonism is by denying the causal theory of knowledge. This is the tack taken by Maddy in her paper “Mathematical Epistemology: What is the Question?”¹² Her argument is that Benacerraf’s conclusion that Platonism is untenable, which she refers to as the Untenability Thesis (UT), rests on two premises: the first is that mathematical entities are causally inert; the second is the causal theory of knowledge, which she claims is now outmoded. She argues that reliabilism is a preferable alternative to a causal theory of knowledge, since there are other ways of justifying knowledge than through causal mechanisms. This is a disputable claim¹³ since many of the most prominent reliable knowledge-forming processes are causal.

However, even if we allow that non-causal, reliable mechanisms can produce knowledge, Maddy finds there is still an uneasiness about Platonism. She traces the basis for this uneasiness to two additional requirements that Benacerraf holds: first that there are some mathematical beliefs that are non-inferential; and, second, that the reliable mechanism which produces these beliefs must involve the objects of belief. Maddy identifies the source of these two requirements as the analogy between mathematics and science: the non-inferential mathematical beliefs are equivalent to perceptions in other sciences, and must be justified by a perception-like mechanism that involves the object of belief. Thus it is Maddy’s claim that “the underlying source of the persistent uneasiness about Platonistic epistemology comes ... from a strong assumption about the nature of mathematics: the science/mathematics analogy.”¹⁴

Maddy does not advocate rejecting the science/mathematics analogy, but rather claims that the task remains to describe a “perception-like base clause

¹² Maddy [1984].

¹³ See my discussion of Brown’s claim that there need be no causal connection in order to have knowledge, later in this chapter.

¹⁴ Maddy [1984], p. 53.

mechanism.”¹⁵ Presumably this must be developed in such a way as to reject the causal theory of perception, since the revised argument for UT (having dispensed with the causal theory of knowledge in favour of reliabilism) has as its premises (1) that mathematical entities are causally inert; (2) the causal theory of perception; and (3) that Platonism involves the science/mathematics analogy.¹⁶ I agree with Maddy that the analogy between mathematics and science is important to retain as it accounts for the indispensability of mathematics in science, and also for other reasons described in Chapter 6. However the causal theory of perception seems harder to refute, and the task of explaining our intuition of mathematical entities given their causal inertness has not been satisfactorily done. Even if we accept that the causal theory of knowledge has been displaced by reliabilism as our best epistemology to date, reliable knowledge-forming mechanisms are by and large causal. If there is no causal connection then we are suspicious of the knowledge as being occult and consider it to be dubious. This is reinforced when Brown’s attempt to show that there can be knowledge with no causal connection is examined, later in this chapter. Also, in Chapter 5, I will examine Maddy’s own account of mathematical intuition.

The notion of mathematical intuition is central to epistemological Platonism, and a satisfactory account of this faculty is required in order to combat the causal interaction objection. Steiner favours an empirical account of the faculty of mathematical intuition, while maintaining ontological Platonism as a theory of the nature of mathematical objects. He believes that the best approach is to argue for ontological Platonism independently, and then to look at how we come to know mathematical truths (which is our present concern) as a separate question. In trying to characterise the faculty of mathematical intuition according to causal theories of knowledge, Steiner runs into difficulties because of one of the most serious problems

¹⁵ Maddy [1984], p. 53.

¹⁶ Maddy [1984], p. 52.

for Platonism, what Brown calls “The Problem of Access,”¹⁷ which arises from the separateness of the physical spatio-temporal world which we inhabit and the Platonic realms where mathematical objects are to be found, if ontological Platonism is accepted. Steiner considers perception as a potential source (or foundation) for mathematical knowledge. This is similar to Maddy’s approach,¹⁸ which is an attempt to ground mathematical intuition in perception. Steiner finds perception to be an unacceptable way of accounting for mathematical intuition since there exist hallucinations or misleading perceptions. His argument is basically this: mathematical knowledge is supposed to be certain and indubitable, and since a subject is unable to distinguish between reliable and unreliable perceptual states the faculty of perception cannot be used as the basis for mathematical knowledge.

Eventually Steiner is drawn to Benacerraf’s structural description of mathematical entities,¹⁹ namely that the only useful properties mathematical objects possess are the structural relations they have to one another. He uses this notion to give the beginnings of an empirical account of mathematical intuition. He explains that we can abstract from physical things the structure of the mathematical entities that are present, so that having an intuition of a mathematical object means abstracting the object’s structure from physical objects. He gives the example of abstracting the structure of Zermelo-Fraenkel set theory (ZF) from dots arranged on a blackboard, explaining that the result of abstracting from the dots is not an intuition of a particular set, but of the structure of all the sets that make up the ZF ontology.

In support of this account of mathematical intuition, Steiner appeals to mathematicians who are convinced of the truth of various mathematical propositions without knowing their proof. This suggests that they possess a faculty which allows

¹⁷ Brown [1999].

¹⁸ Maddy [1980].

¹⁹ Benacerraf [1965].

them to apprehend these truths without going through a rigorous proof. The fact that mathematicians tend to agree with remarkable concurrence on the truths of mathematics suggests that these propositions are based on common elements of our experience of the world. Since many different cultures have come up with the same mathematical results, it seems unlikely that mathematics is socially constructed. Furthermore, if mathematical truths were a product of social uniformity or convention it would be difficult to explain their remarkable applicability to the physical world. Thus it is likely that the truths of mathematics somehow arise from our experience of the world, so they are either a product of how the world is (independently of us) or how we perceive the world (dependent on both our hard-wiring *and* how the world is). We can never interact with the world in a way that circumvents our perception of the world; however given the fact that there is a wide degree of consensus as to what the physical world is like, it seems reasonable to presume that even if our modes of perception influence the kind of impression we have of reality, there is still some objective reality that is independent of our experience of the world. Since mathematics is able to explain and predict what happens in the world, this objective reality must have mathematical aspects to it. These aspects are the structures that form the subject matter of mathematics, and they cannot be purely a product of our minds.

Brown admits that, for the Platonist, there is no adequate explanation for how we can acquire knowledge of entities that exist in an abstract realm, namely that the problem of causal interaction in mathematical knowledge acquisition, or the Problem of Access, is a serious difficulty for Platonists. His response to this problem consists of two main points: first he claims that perception of Platonic Forms is not much more mysterious than ordinary perception, since he believes the mind-body problem to be unsolved and there is no good account of how the physical process of seeing something results in a mental belief of that object's existence. This is not a very strong solution to the problem of access, because evidently we have much better access to cups and chairs than we do to the Platonic sphere. Even if a theory of

ordinary perception is lacking (and with a materialist account of the mind this need not be the case), we still have a much better chance of acquiring knowledge about the physical objects around us than of Platonic Forms.

Brown's second point consists of denying the claim that causal interaction is required in order to have knowledge. He does this by way of an example of knowledge in the physical realm which is not produced by a causal connection. The example he gives is from quantum mechanics, and involves an EPR-type set-up in which there are two detectors at either end of a room which are able, by way of polarisers, to detect whether photons emitted from a source in the centre of the room have the property *spin-up* or *spin-down*.²⁰ It is invariably the case that the outcomes at opposite ends of the room are correlated, so that one photon is up and the other is down. This is not due to any causal interference between the two detectors. Neither measurement can affect the other as they are made simultaneously and hence are outside each other's light cones. However, if we observe a measurement at one detector and find it to be spin-up, we can infer that the photon which arrived at the other end of the room is observed to be spin-down without having any causal interaction with the other detector.

Brown takes this to be an example of knowledge about the physical world without causal connection. The problem with the example is that in order to determine the spin property of the particle at the far detector, we need to have the prior knowledge that the two particles will always be correlated. This knowledge is derived from our scientific theory that describes the situation. It is essential that we have this information in order to infer the spin of the particle at the far detector, and acquiring this background piece of theoretical information provides the causal link. Thus this example fails to refute the causal theory of knowledge.

²⁰ See Brown [1999], pp. 16-17 for a more detailed description of the example and a diagram.

For Steiner, as an ontological Platonist, the fact that mathematical truths are uniformly agreed upon suggests that they are all about the same abstract objects. However, it makes even more sense that they are all about the same underlying structures in the physical world. Since all mathematicians' sensory perceptions arise from the same world, if there are certain mathematical structures that underlie the physical world, it would be expected that mathematicians will (largely) agree on the truths that they derive about these structures. Of course, in order for this to be the case, the process of abstraction, which begins with perception of the physical world, must preserve the essential elements of mathematical structures in the move from their physical instantiations to their abstract characterisations. Since it is the case that the truths of mathematics are more or less agreed upon, we can conclude that the process of abstraction which is in play does in fact preserve the integrity of mathematical structures. The rest of this chapter will examine the process of abstraction and seek to determine how it is able to preserve these essential structural relations.

3.3 Abstraction: How it Works

The distinction which Shapiro makes between the places-are-offices and places-are-objects ways of viewing structural entities brings out some crucial elements of a structuralist epistemology and ontology, and shows how central the notion of abstraction becomes in the process of knowledge acquisition in mathematics. In order for structuralism to provide an account of how we come to have knowledge of mathematical entities as seemingly abstract things, we must give an account of abstraction which demonstrates that the essential features of mathematical structures are preserved. From Resnik's example, we have seen how some axioms governing finite sequences can be observed directly from a sequential collection of physical items. Abstraction is the way we get from a row of houses, a row of dots, a stack of coins, etc. to the abstract structure of such a sequence. It is the process of

identifying and extracting the underlying essential structure in a physical system. The abstract structure²¹ that results from this process is one that captures the components of the physical system that are involved in a structural relation. This type of abstraction process is one of simplification.

For example if knowledge of the structure of a sequence is gained by abstraction from a row of telegraph poles, the abstract system has captured the structural essence of a finite sequence as exemplified by the telegraph poles, but nothing else about them. The process of abstraction has omitted to capture details such as their relative heights, the fact that some are made of wood and some are metal, and other properties which do not influence their relations as elements of a finite sequence. This is why we can extract *exactly the same structure* from a row of dots on a page, since abstraction allows us to capture the elements of the physical system which are mathematically significant, while neglecting facts such as the diameter of the dots, the fact that they are drawn in blue ink, and so on. This account of abstraction was advocated by Leibniz, who claimed that “[a]bstraction is not an error as long as one knows that what one is pretending not to notice is *there*.”²² This means that you can disregard all those properties of the physical system that are “irrelevant to some present concern”²³ without actually denying those elements, merely leaving them aside for the purposes of mathematical considerations.

A possible objection to this account of abstraction involves questioning why it is that we are able to draw out the mathematical structures instead of all the other common elements between groups of objects. For example, if we see five frogs sitting along a log, why should we be able to abstract the number 5 from what we see instead

²¹ Note that the structure of a sequence is not the only structure that we could extract from the row of dots. Different structures can be abstracted from the same physical system, depending on what the observer leaves out. For example, from the row of dots we could ignore the numerical aspect and acquire an abstract notion of a dot. This idea will be developed further later in this chapter.

²² Leibniz [1765], p. 57.

²³ Leibniz [1765], p. 57.

of the colour green? It is true that there are many other common properties that objects can possess (colour, size, texture, and so on), and that we can gain an idea of these properties as universals through a similar process by which we gain knowledge of mathematical entities. When we examine a group of objects we may choose to focus on various aspects of their existence. Looking at the five frogs, we notice that they are all the same colour, and thus have an impression of the property “greenness.” We might also notice various other features, such as the fact that the middle frog is the smallest, and the largest is on the far left. We may also notice that each frog has no more than one frog to its immediate left, and that the relation of “being to the left of” is transitive.

After looking at the frogs, we might see seven birds perched along an overhead electrical cable. From our perception of the birds we can abstract common properties such as “being feathered” and “having two legs,” but we also notice that the birds share some of the properties of the frogs, namely that each bird has no more than one successor, and once again the relation of “being to the left of” is transitive. We might notice that these properties are also exhibited by the telegraph poles on whose wires the birds are sitting. Thus we come to recognise certain common properties or patterns among all rows (or ordered sequences) of objects, despite the fact that they differ in many other properties (including number). These properties are the structural properties of the natural number sequence, and they interest us because they present themselves to us constantly in a multitude of different instantiations. This is not to say that “greenness” or “having feathers” are not recurrent properties of which we may acquire abstract knowledge, but since we are examining the process of acquisition of mathematical knowledge, it is our acquaintance with recurrent mathematical structures that is of interest. When we observe that certain structural properties are recurrent among all ordered groups of objects, we can begin to consider only the common mathematical features that all the groups of objects have. By ignoring factors like greenness, number of legs, woodenness, and all other non-structural properties,

we can discover the underlying mathematical structure shared by all these groups of objects. This process of simplification leaves us with the properties of the mathematical structure itself, and by considering the common elements between all observed instantiations we can deduce outcomes in the abstract structure which will then apply to all the physical instantiations of that structure.

Just as we can abstract both mathematical and non-mathematical features of any given system, so we can abstract different mathematical structures depending on which aspects of the observation we leave out. For example, if we see a loop of string and leave out its colour, thickness, the kind of fibres it is made of, and so on, we can focus on its circularity. In this way we have abstracted a circle from the loop of string. However if we have another loop of string that is knotted, this might prompt us to focus on the knot properties of the first loop of string. Now all the other irrelevant concerns (colour, thickness, etc.) are still irrelevant, but additionally the shape of the loop is irrelevant. All that we are concerned with now is the topology of the line formed by the loop, the fact that it does not intersect with itself and has no ends.

The reason mathematics is so successful in describing the physical world is that it is in large measure about structures that exist in the physical world, and because abstraction can capture the mathematical structures of a physical system. The structures described by mathematical theories are structures that exist in the physical world and underlie physical phenomena. Even though we can talk about them independently of their physical instantiation, this does not make them any less a part of reality. Mathematical structures exist *in* the physical world, and we can talk about them abstractly even though they do not exist abstractly, because they exist independently of any *particular* instantiation (although not independently of *every* instantiation). Speaking about structures independently of a specific instantiation can be seen as a type of generalisation or shorthand, as I can illustrate with Resnik's axiom about the row of dots. This axiom says that "every element in a finite sequence has no more than one successor." If we examine how we really came to have

knowledge of this axiom, namely by observing that it held for rows of dots, rows of houses and stacks of coins, then we could express it as “every dot has no more than one successor, AND every house has no more than one house situated directly to its right, AND every coin has no more than one coin directly on top of it” and so on. The axiom is a generalisation from all the perceived cases of a particular structure (in this case a finite sequence) and expressing it in such a way as would be applicable to any such structure we encounter in the physical world. It is also much easier to reason about finite sequences in order to determine more complex features of the structure when it is in this form, rather than having to deal with all the extra-structural features involved in physical instantiations. The reason this process works and still allows us to describe features of the world using the results discovered using abstract versions of the structures, is that the essential structural properties are maintained in the shift from the concrete instantiation to the abstract concept. An abstract concept of a structure possesses all the relevant structural properties of the concrete structure, but none of the other physical properties that do not affect the structural characteristics.

This characterisation of abstraction is subject to an objection that Frege made against Mill’s account of induction. In *System of Logic* Mill claims that all reasoning is based on induction, by which he means (non-deductive) inference from many specific observed instances to a general case.²⁴ He argues that the truths of mathematics are only certain insofar as they follow from their premises, or the first principles of geometry and arithmetic. These first principles are inevitably based upon inductive inferences from observation. He gives the example of the axiom of Euclidean geometry which states that two straight lines cannot enclose a space, arguing that this truth “is an induction from the evidence of our senses,”²⁵ since our degree of belief in this axiom is due to its constant confirmation in our experience of

²⁴ Mill [1872], Book II, Chapters V and VI.

²⁵ Mill [1872], p. 231.

the world. Similarly, truths of arithmetic find their basis in observation: $2 + 1 = 3$ “is a truth known to us by early and constant experience: an inductive truth.”²⁶ On this view, since any mathematical truth derives its certainty only from following deductively from the axioms, the certainty with which we hold any mathematical proposition cannot be greater than the certainty we attribute to inductive inferences based on experience.

Frege objects to this characterisation of mathematical knowledge on the grounds that mathematical truths are in fact logical rather than physical truths, thus claiming that Mill has confused the application of a mathematical law with the law itself. He argues that we cannot use induction in the process of discovering mathematical truths since this method fails to produce certain truths, as mathematical truths are generally held to be. If induction is the basis for mathematical truth, this means we can never be certain of mathematical knowledge. In short, induction “can never render a proposition more than probable.”²⁷ Furthermore, Mill’s account of mathematical knowledge becomes circular since we need the laws of arithmetic in order to develop probability theory.

This objection is mistaken in two ways: first, we do not need to be able to defend a rigorous theory of probability in order to hold a probable belief. I can be fairly certain that $2 + 2 = 4$ and realise that I am more certain of that than I am of the proposition that it will rain next Tuesday. It is not necessary that I grasp probability theory in order to hold probabilistically likely beliefs. Probability can be used as an external tool to explain belief formation, it need not be grasped by the subject who is forming the beliefs. Consider a small child who is asked, on an afternoon of a rainy day, if he thinks that it will rain tomorrow. He might look at the grey clouds outside and conclude that he believes it will rain tomorrow. He could also be asked if he

²⁶ Mill [1872], p. 256.

²⁷ Frege [1884], §9.

believes his mother will come home from work soon. Given that his mother comes home every afternoon at 4pm and it is now 3:30pm, he will say yes, he believes that also. With no training whatsoever in probability theory, or any kind of formal schooling, if asked which belief he is more certain of the child could answer that he is more certain of his mother's impending arrival than he is of rain tomorrow, as he has experienced consistent regularity in the time that his mother comes home, but knows that rainy days are not always followed by more rain.

Second, and more importantly, mathematical beliefs do not have to be entirely certain. In Chapter 8 we shall see that mathematical beliefs are fallible and that we do have probabilistic degrees of belief in mathematical propositions.

3.4 Types of Abstraction

There are two main types of abstraction: *abstraction by simplification* and *abstraction by extension or recombination*. When we abstract by simplification, we omit the features of a physical system that are not relevant to our present concern. For example we can observe a row of six fire trucks and gain by abstraction (in the sense of simplification) an abstract idea of the number six. In order to do this we disregard all their properties that make them fire trucks (except for their distinction from one another) and concentrate only on the number which they instantiate. By observing the same physical objects, we can also gain an abstract idea of the colour red: this time we leave out the numerical properties of the objects we see, as well as all their vehicular properties, and consider only their colour. Another example of abstraction by simplification is how we get the idea of a sphere from a cricket ball. The shape of the ball is close to that of a sphere, and we can ignore the stitches at the seam and some roughness in the leather to gain the idea of the shape of a sphere. We can also notice

that the abstracted shape of a cricket ball is the same as the abstracted shape of an orange, and this is how we come to have knowledge of spheres independently of their physical instantiations. Shapiro²⁸ gives another example of this sort of abstraction, which is analogous to pattern recognition: when we grasp the type of a letter, say “E,” we do so by observing several printed instances of the letter. We are able to focus on the essential pattern of the letter, that which all observed tokens of it have in common, while ignoring or leaving out the relative heights of the different instances and which colour ink they are printed in. This is how we get an abstract idea of the letter “E,” and it is in the same way that we gain an abstract idea of mathematical entities.

Islam²⁹ gives a good account of this type of abstraction, and uses it to show how mathematical theories can explain physical phenomena. Given a mathematical³⁰ theory T, there could be several models that fit T, but it is also possible to think of an unspecified model of T, and this is its abstraction.³¹ A model of T is any system (often a physical system, but not always) which exhibits the structure of T. For example if T is the theory of Euclidean geometry, then a model of T could be any collection of physical objects made up of lines, planes, etc. which would then exhibit relations in accordance with the axioms of T. However, a model of T does not have to be a physical system. Another model for the theory of Euclidean geometry is the set of all ordered triples. If each ordered triple is considered to represent a point in \mathbf{R}^3 , then the axioms of Euclidean geometry will describe relations between various subsets of the set of all ordered triples. Thus the set of all ordered triples constitutes another model for the theory T (and is isomorphic to three-dimensional Euclidean space).

²⁸ Shapiro [1997], p. 74.

²⁹ Islam [1996].

³⁰ T could be any kind of theory, but a mathematical theory is used here because these are most relevant to the present concern.

³¹ Note that ‘abstraction’ here is not used in the Platonic sense.

The other type of abstraction consists of abstracting from a system to get an extension of it. To do this, we recombine the concepts involved in the basic system in ways that do not occur in the original system. For example, the negative numbers are an abstraction from the natural numbers, since they are based on them. We are given the natural numbers in experience³² (using abstraction by simplification from various physical systems) but when we reason about the structure in an uninstantiated form we find that our theory about it is not rich enough, for example we cannot find answers to sums such as $7 - 12 = ?$. We can extend the natural number structure to postulate the negative numbers, and this abstract extension makes our concept of numbers (i.e. our theory about the number structure) richer and more powerful. It happens that once we have performed this abstraction, we find that the elements of the extended structure also exist in the physical world, for example we can perform some calculations and determine that a location is -3km west of some point of reference (which just means that it is 3km east). The structure of the integers is not as evident in our experience as the natural numbers, but when we have the concept of negative numbers we find that they are in the world also. However we can also use abstraction from the structures we observe to come up with structures that are not found in the physical world, such as n -dimensional geometries. We can observe one- and two-dimensional spaces and their relations to three-dimensional space, and having grasped these concepts we can abstract to higher dimensions. This is how we can reason about, for example, sixteen-dimensional geometry, even though we can never observe it.

An objection that Frege made to Mill's account of mathematical knowledge arising out of our experience of the world supports the notion of abstraction by

³² Although all we ever experience is finite sequences, we arrive at the idea of the natural numbers by abstracting from any given finite sequence. The finite sequence gives us the concept of iteration, which allows us to move beyond the number of elements given to us in experience, as illustrated by Resnik's example earlier in this chapter. Thus we acquire the concept of an indefinitely extensible sequence.

extension. On Mill's view our knowledge of the truths of mathematics is based in generalisations from observations. For example we know that 3 is the same as $2 + 1$ because we have put one object together with two objects and found that there are three, and our observations of the world constantly confirm this truth, which accounts for the high degree of certainty we attribute to mathematical propositions. Frege objects that if this is how we acquire our knowledge of the basic truths of arithmetic and geometry, then we cannot know the number zero because we can never observe it.³³ Using abstraction we acquire concepts that allow us to consider various numbers independently of their observed instantiations, and this is how we become acquainted with the natural number structure. Once we have an abstract idea of the natural numbers then we can manipulate them and come up with extensions to the structure we observed (abstraction by extension). This is how we come up with negative numbers, fractions, complex numbers and so on, but it is also how we first became acquainted with zero. The fact that zero was not discovered until long after people started to use basic arithmetic supports this claim,³⁴ and provides good evidence for the above account of abstraction.

The distinction between these two types of abstraction is important, because abstraction by simplification is how we extract mathematical structures from experience, while abstraction by extension leads to derivative structures. Some of these are found in the world, some are not, but the ability to theorise about them makes mathematics much more powerful. Chapter 4 deals with basic and derivative structures, and gives more detail on the abstraction used to get from basic to derivative structures. It should be noted that a theory of or about a mathematical structure will never be complete (because of Gödel's incompleteness results),

³³ Frege [1884].

³⁴ Arithmetic using zero first appeared in the work of the Indian mathematician Brahmagupta in the seventh century AD. See Boyer [1989], Chapter 12 for further details.

however this does not preclude our theories about mathematical structures from being extremely useful.

If I observe a finite segment of the natural number structure in a sequence of twelve cats and ignore all the features of the cats except for the number they instantiate, I have successfully acquired a concept of the number twelve. The theory I have about this number and others is incomplete, but my reference to the number does not leave anything out. To clarify this point, say that I take one of the cats in the sequence and consider it in the abstract (by simplification, by ignoring everything that I am observing except for that cat). This gives me an abstract idea or a concept of the cat and when I refer to it, I refer to it in its entirety as a rich and complex entity. Nothing is left out when I specify that cat, my reference is not limited to the aspects of it which I have fully grasped. The theory that I have about the cat is far from complete, I do not know the colour of the pads of its paws, its sex, anything about its chemical makeup, how its internal organs function or any details about its genetics. Still my reference to it specifies it exactly and completely, even though my theory cannot describe it completely. While many aspects of the cat are left out of my theory about it, my abstract concept of it is nevertheless able to capture it and refer to it in a robust way (in its entirety). Similarly, although arithmetic is incomplete, we can still refer successfully to numbers without omitting aspects of their existence.

Abstraction is the process of identifying and extracting the underlying mathematical structure in a physical system, or applying further simplification or extension to another mathematical structure. The abstract structure that results is one that captures all the components of the physical system that were involved in a structural relation. Although our mathematical theory is incomplete, we can still capture all the mathematical aspects of a physical system and our abstract conception specifies the mathematical structure entirely. It is because abstraction captures the relations of a physical system that mathematics is so applicable in the physical world. The structures with which it deals are structures that underlie physical phenomena,

and they only seem abstract when we talk about them independently of their physical instantiation. However, this does not make them any less a part of physical reality. Since mathematical structures come from the world, truths about our concepts of them can describe the world accurately when they are applied back to physical situations.

The reason that abstraction by simplification from experience is a legitimate process for acquiring concepts which apply directly to the physical world is that it does not involve any distortion.³⁵ Abstraction takes a structure in the world and ignores the physical attributes that are irrelevant to its mathematical properties. This is how we can get an idea of the same mathematical structure, such as the natural numbers, from so many different physical systems. We can see it exemplified in rows of houses, collections of marbles, boxes of cigars and indescribably many other groups of objects. This is because, when considering arithmetical properties, the process of abstraction looks only at the numerical properties of these groups, and ignores irrelevancies such as the shapes of the objects involved, their colour, texture, and so on (although at another time we may focus on some of these properties in order to find different mathematical structures, such as geometry). So abstraction by simplification is a process for extracting the mathematically significant properties of structures that are found in the world, while leaving out those elements which do not impact the workings of the mathematical system. Just because some of the finer details (such as unevenness on the surface of a wheel) are omitted in order to get at the underlying mathematical structure, this does not mean that our abstract concept is not representative of reality.

In these situations the notion of verisimilitude is important, because there is a margin for approximation. For example a soccer ball is an instantiation of a sphere

³⁵ We shall see in the chapter on basic and derivative structures that abstraction by extension can involve some degree of fictionalisation, and this can result in competing mathematical theories and undecidable mathematical claims. However when we abstract structures from perceptions of the physical world, there is no distortion involved.

despite some irregularity in its shape, whereas a rugby ball, while still round, is too far removed to give us the idea of a sphere. Rosenkrantz³⁶ provides an account of verisimilitude in which the truthlikeness of a theory depends both on its accuracy and its content. This lets us measure the relative closeness to the truth of various false theories. We need to take into account not only how accurately a theory describes reality, but also how rich that theory is in terms of what it can tell us about reality, namely its content.

The motivation for Rosenkrantz's analysis is the fact that all past scientific theories have been found to be inadequate, and have been overthrown by newer theories that are more accurate. It is considered likely that the same thing will happen to many, if not all, of our best contemporary theories. This phenomenon is related to the inability of any formal system to completely capture some aspect of the world, which is discussed in later chapters. Since there is no perfectly true theory, principles of verisimilitude have been developed to show how one theory can be closer to the truth than another. Rosenkrantz's theory is derived from the consideration that if two competing theories are equally accurate, then the one with more content is closer to the truth. The more content a theory has, the more detail it provides, the more room it has for error. Thus a theory with more content (and hence greater scope for error) that has the same degree of accuracy as another theory with less content, will be more truthlike than the latter theory. Rosenkrantz proposes a mechanism using a weighting function to measure the relative content of theories.³⁷ This allows us to evaluate competing theories based not only on their accuracy, but also the richness of their content.

For present purposes we can apply these principles to a mathematical theory and a physical system that instantiates it. A sheet of paper is an instantiation of a

³⁶ Rosenkrantz [1980].

³⁷ See Rosenkrantz [1980], pp. 466-470 for further details.

Euclidean plane, but the theory is not entirely accurate in describing the physical system. The paper is unlikely to be perfectly flat, and lines drawn on it with a pencil and ruler are not perfectly straight. Similarly a cricket ball is not a completely accurate instantiation of a sphere, as the leather is not perfectly smooth and it may bulge at the seam. In these cases it is preferable to use a false theory, namely theories of Euclidean geometry that the physical systems do not strictly instantiate, rather than attempting to increase the accuracy of the theory. Two-dimensional Euclidean geometry is extremely powerful in describing the relations between lines, curves, angles, points and so on that are drawn on a sheet of paper. Thus it is a theory with a great deal of content and explanatory power, even though it is not perfectly accurate. We may measure an angle to be a degree larger than its value obtained by calculation, however such a result is still more powerful than any that we would obtain by trying to capture all the variations of the paper, the thickness of the pencil's point, and so on. As Rosenkrantz notes, "a nearly vacuous truth may be far less truthlike than a false but quite accurate and highly specific one."³⁸ In the case of mathematics, the vacuous truths are those acquired by trying to capture all the irregularities of physical objects. If we abstract away from small variations in their shape, a cricket ball and a beach ball have the same shape whose surface area and volume we can easily calculate given their diameter. However if we take irregularities into account the two physical objects have completely different shapes that can only be captured by enormously complex equations, if at all. Thus in abstracting the mathematical structure from a physical system it is preferable to sacrifice a small degree of accuracy in order to have a far more powerful theory.

³⁸ Rosenkrantz [1980], p. 478.

3.5 Abstraction and Epistemic Structuralism: Why it Works

The above exposition of how the process of abstraction lets us move from a physical system to an abstract mathematical structure is a fundamental part of the structuralist account of how we acquire mathematical knowledge. It is necessary to provide this explication of the mechanics of abstraction in order to show how we come to know mathematical truths under the structuralist picture. In order to demonstrate that structuralism is a useful theory, it is necessary to show that it yields a satisfactory account of how we come to have a body of mathematical knowledge that so successfully describes the world. Thus, it is also necessary to show that the process of abstraction is reliable. There are two crucial tasks involved in developing an account of abstraction in mathematical epistemology: we need to give an account of how abstraction works in order to show how we may come to have mathematical beliefs at all; but we also need to develop an explanation of why the process of abstraction is a reliable (or truth-preserving) mechanism, in order to explain why the beliefs that we gain via abstraction constitute mathematical knowledge. While it is possible that these two explanations may coincide, it is not the case that the account of abstraction provided above implies that it is a reliable process. Thus we look to factors other than the actual workings of the process itself, in order to show that it is a reliable process.

Since structuralism claims that we acquire mathematical knowledge by abstracting from concrete structures, reasoning in the abstract structure, and projecting the abstract results back on the physical instantiation of the abstract structure, it is crucial that the process by which we make the move from the concrete structure to the abstract preserve the essential mathematical relations of the system. While this is not implied by a description of how the process of abstraction works, the reliability of abstraction as a truth-preserving process is evident from the fact that our mathematical knowledge consistently shows itself to be reliable in describing and predicting events in the physical world. The reliability of abstraction as a process that preserves

mathematical relations is indicated by the fact that truths arrived at via abstraction and then projection from an abstract system back to the physical structure are accurate, and predictions cohere with observations. The reliability of results gained through abstraction are sufficient to show us that it is a reliable process, for were it an unreliable process, the results would not be so startlingly and consistently accurate.³⁹

An obvious response to this argument for the reliability of abstraction is to deny the claim that abstract mathematical results, when transferred into physical systems, are true. However, if we restrict our analysis to cases where the structures or systems are well-defined,⁴⁰ and also eliminate those cases in which errors have been made in the abstract reasoning, mathematical truths in their abstract and concrete instantiations concur extremely well. For example, there is no physical instantiation of the natural numbers in which $2 + 3 = 5$ is not true, as long as we disregard objects that are not well-defined or well-formed. This may seem to be question-begging, but it is a required clause.⁴¹ Truths about the natural numbers are abstracted from certain physical systems, and in order to apply the truths to the physical world we must restrict ourselves to those same sorts of physical systems. This is the same as restricting theories of fluid dynamics to physical systems composed (at least in part) of fluids: it is not question-begging, but restricts the theory to a domain of application that contains the subject matter of that theory.

³⁹ Much of mathematical practice is not concerned with application, and consists solely of reasoning about mathematical structures independently of their physical instantiations. Indeed, the greatest challenges of mathematics are to construct proofs of mathematical results, and proofs are purely about mathematical structures. Nevertheless, results that are proven by pure mathematicians can be translated into truths about the physical world (if suitable models of the mathematical structures can be found); the objectivity of mathematics is preserved by the link to the physical world.

⁴⁰ This is what Wilson [2000] calls the Stability Condition for mathematical application: the rules of arithmetic apply to objects that are stable or well-defined, excluding such things as raindrops that can coalesce, resulting in claims such as $1 + 1 = 1$, and rabbits that can reproduce and lead to claims that $1 + 1 = 16$.

⁴¹ See the “Shift in Meaning” argument in Chapter 6 for further consideration of this clause, and how to use it in a non-question-begging way.

Castañeda⁴² demonstrated that we cannot use another arithmetic than the standard one we use, and still describe the world adequately. His argument is analysed in Chapter 8, and he succeeds in proving that the arithmetic we use is a built-in feature of our physical world. The arithmetic we use is a feature of the physical world, and arithmetical results can accurately describe aspects of the physical world. Hence we must have some way of extracting the arithmetic from the world in the first place. This is reliable evidence that the arithmetical beliefs we have are true of the world and that abstraction is capable of providing us with knowledge about real mathematical structures. The fact that the process of translating truths discovered in the abstract structure back into some physical instantiation of that structure tends to result in concurrence, is a good indication that abstraction is a reliable process.

Further evidence for the success of abstraction in preserving truth is found in the formal sciences. These are studies in which the structures in question seem closer to their physical models than some mathematical structures. This ‘closeness’ just means that the abstract structure contains many features of the physical system. For example in queuing theory, the structure of a queue could be instantiated in a row of people. The abstract structure captures not only the linearity and cardinality of the row, it also contains certain manipulation rules that exist in the physical system, such as that a new member can only join at the back of the queue, and the only member who can leave the queue is the one at the head. So while a row of people instantiates both the natural number structure and the structure of a queue, *more features* of the physical system are captured in the latter abstraction; in other words we leave less detail out when we abstract the queue than the natural numbers. Because of this closeness, the formal sciences perform a useful psychological function. They demonstrate in an obvious and accessible way how truth in an abstract structure can

⁴² Castañeda [1959].

correspond to truth in a physical model. Many people would find it easier to accept that an abstract concept of a queue will yield a truth about a real queue than that an arithmetic truth corresponds to a truth in the physical world, however there is no significant difference between the two.

This evidence shows that the process of moving to abstract mathematical structures from their physical instantiations is a process that works; it results in an accurate capture of mathematical aspects of the physical model. The fact that the process provides us with a mathematical model whose truths hold in the physical system shows us that abstraction is a reliable process. An analogy might be useful to show why this is the case. In the days when the earth was thought to be at the centre of the world, people knew that the sun would rise every morning. They were not well enough acquainted with the workings of the solar system to know exactly *why* the sun came up every morning, but the fact that it had consistently risen every day with no known exceptions was enough for them consider it a reliable bet that it will rise tomorrow. Similarly with abstraction, it may be the case that we do not know enough about the process itself to show why it *must* be reliable, but the fact that it has always proven to be reliable gives us sufficient reason to presume that it will continue to do so. This, I believe, is a reason why we can use abstraction to make the shift from a physical structure or several physical instantiations of the same mathematical model, to an abstract structure, and know that truths discovered by reasoning in the abstract structure will also hold in all of its physical instantiations. We shall also see in a later chapter that we can never be completely certain of any beliefs, even mathematical ones, so based on its reliability and accuracy, we can take abstraction to be a reliable process for generating knowledge, because it results in a body knowledge of which we are more certain than most.

3.6 Iterativism

In his paper “A Number is the Exponent of an Operation”⁴³ Hand argues for iterativism as a preferable alternative to structuralism in providing an account of the natural numbers. One of the reasons he gives is that abstraction, which is a required component of epistemic structuralism, is unnecessary for our gaining an understanding of arithmetic. Hand’s discussion of iterativism is useful to show how we can get an idea of infinity from our experience of the world. This is an important question for any physicalist account of mathematics and Hand’s theory has some valuable insights regarding the mechanisms by which we can move from experiences of finite sequences to grasping an idea of infinity. However, we shall see that he fails to refute structuralism because it is not possible to do away with abstraction in the process of moving from physical mathematical systems to an abstract idea of an uninstantiated mathematical structure.

Iterativism is a theory which explains the acquisition of our knowledge of the natural numbers (also referred to as the omega sequence), claiming that it arises from the iterative nature of procedures such as counting. Hand explains that according to iterativism, the natural numbers are taken to be controls on increments or iterations. Rather than thinking of a number as an object, it can be considered to be a regulation on the counting process, limiting the iterations after a certain point. Roughly speaking, this means that any number can be defined as a certain degree of iteration, going so far and no further. This is a recursive definition, since each number is defined in terms of previous numbers. The number one, as a unit, is the basic starting point, then two is defined as the result of performing an iteration on a unit, three is defined as the result of performing another iteration on two, and so on. So, for example, the number seven

⁴³ Hand [1989].

is described as what you get if you perform as many iterations as you need to get to six, and then perform another.

The way that Hand arrives at iterativism is by examining the faculties required of a subject in order to gain knowledge of the natural number sequence, and deciding that abstraction is unnecessary to the process of acquisition of this knowledge. The only sequences that we have access to through experience are finite sequences, and thus they can only function as initial segments of the whole natural number structure. We do have some impression of infinite sequences because of the continuity of space-time; for example we know that there are infinitely many points in a finite spatial segment.⁴⁴ However, these points are so tiny that we cannot separate them, and so we cannot count them. The sequences with which Hand is concerned are those that can be regulated by procedures such as counting, for his claim is that “the source of arithmetical truth [is] in our counting procedures, or more generally, in the iterative character of procedures among which counting is the most obvious.”⁴⁵ Under this countability requirement, we can assume that we only have experience of finite segments of the natural number sequence. To gain knowledge, or an idea of the natural numbers, we must be able to make the move from these finite segments to the infinite sequence which is the natural numbers. This means that we must recognise these sequences as being sub-sequences of a similar, but infinite, sequence.

Hand explains that the conceptual step involved in acquiring this knowledge is recognition of the fact that any finite sequence can be extended indefinitely. He maintains that this realisation breaks down into two parts: first we must see that any finite sequence can be extended, and then we have to recognise that this extension is indefinitely iterable. He claims that abstraction on its own cannot provide a person with the conceptual tools necessary to conceive of the successor of a sequence of

⁴⁴ The claim that we get an impression of *actual* infinity in this way will be refuted in a later chapter.

⁴⁵ Hand [1989], p. 243.

length n without the existence of a concrete sequence of length $(n + 1)$. The mistake which he makes here is to presume that the idea of indefinite extensibility must come from abstraction by isolating the required components from the physical system. It is true that we can only abstract finite sequences by simplification from physical systems, however once we have the concept of a finite sequence then we can simplify further and leave out the upper limit of the sequence. Thus we come to have an abstract idea of an infinite sequence.

Once we have an abstract idea of a perceived mathematical structure, we can abstract from that idea by simplifying or recombining its concepts in order to come up with an extension. Hand argues that any sort of structuralism would have to incorporate the idea of iteration⁴⁶ in order to give us the conceptual tools necessary to arrive at the concepts we have. He gives this type of theory the name of “iterative structuralism.” He then goes on to argue that the structuralist aspect of this theory is superfluous, and all that is required is the concept of indefinite iterability. He claims that “these non-abstractionist capacities are all that are needed, and hence that structures may be disposed of. The concept of indefinite iterability underlies our ability to understand and decide basic arithmetical statements such as numerical statements, and also to understand, if not always to decide, quantified arithmetical statements.”⁴⁷

Thus Hand’s claim is that abstraction has no role in the acquisition of our knowledge of the natural numbers, and that the concept of indefinite iterability is all that we require. Granted, it may be possible to provide an account of our understanding of the natural numbers in terms of iteration, with no reference to structures. However the possibility of such a reduction does not mean that this is how we come to be acquainted with the natural numbers, nor does it constitute a

⁴⁶ This notion of iterability can be considered to be a species of abstraction, in the sense of extension.

⁴⁷ Hand [1989], p. 260.

convincing argument against the claim that numbers are places in structures. Intuition and personal experience will tell us that we *do* gain our ideas of the natural numbers through abstraction from our experiences of the physical world. When we see a row of eight telegraph poles, we notice that they exhibit a common pattern with the eight birds perched along the wires running between the poles. Then we notice that this common feature is similar in essence to other groups of objects, and that the numbers in the groups can vary. From this we can use abstraction (in the sense of extension) to increase the size of the group indefinitely, and so we have an understanding of the natural number structure. Once we have grasped the indefinite extensibility of the natural numbers we realise that they have no largest (or last) number, and this is all that is required for an idea of infinity.⁴⁸

In contrast to Hand's account, consider how children learn to count. The first thing they learn to do is to chant the numbers by rote, and they find great pride in learning to count to ten for the first time. This mechanical memorisation tends to come before the ability to count actual objects. Hand might argue that since children learn to recite the numbers before they learn to count things in the world, their knowledge of the numbers comes from understanding the principle of indefinite iterability. They learn a number, and then learn the next number, and the next, and so on. Only after they have some idea of the numbers can they start applying this knowledge to the world around them. However, I do not think a child can be said to have knowledge of the numbers until they know how to apply them. When a child is learning to count (before they begin to count *things*) she is learning nothing more than vocabulary, a string of words. A child who can recite the numbers one to ten but cannot yet count a row of blocks does not have *knowledge* of the numbers. Knowledge only comes when she learns to apply the vocabulary to things in the

⁴⁸ Our knowledge of infinity will be examined in greater detail in a later chapter.

world. When a child chants the numbers from one to ten they have no meaning for her, they are simply a string of words which go together, much like a nursery rhyme. Only when the child is shown how to apply those words to physical things does she begin to grasp their meaning and acquire the concept of the natural number sequence. Thus, it is by recognising the common structure in all the groups of objects to which the number words can be applied, rather than from the abstract counting procedure, that a child gains an understanding of the natural numbers. This comes by considering this common structure in isolation, ignoring all the non-numerical properties such as the colour and shape of the blocks, that the child has grasped the natural numbers, and this must involve abstraction.

There is another reason that structuralism is a stronger theory than iterativism: as well as providing a much more intuitive account of our knowledge of arithmetic, it extends to all other mathematical structures that we know of today. It is a significant advantage of structuralism that it provides an explanation for *all* our mathematical knowledge. If we were to adopt iterativism we would have to find a different theory to account for our knowledge of geometry, as well as all the other branches of mathematical knowledge.

Another problem for Hand's account is that iterativism must allow for abstraction at some stage, for we need it to acquire the idea of a finite initial segment of the natural numbers. We must somehow get from five trees to the number five in order that we can conceive of the numbers six, seven, and so on. We need the faculty of abstraction in order to make the move from our perceptions to our uninstantiated concepts of mathematical entities. While iterativism is more concerned with how we make the conceptual shift from finite sequences to infinite sequences, it still cannot dispense with abstraction altogether, as Hand would perhaps wish. Iterativism as it stands has no link to the world, and thus is at a loss to explain whether our knowledge of mathematics originates in the world and, if not, how it can describe the world as well as it does.

Hand claims that abstraction is unnecessary because it does not provide us with the notion of indefinite iterability, but we do still require it to get the idea of a finite sequence which can function as an initial segment of the natural numbers. Even if we do get our understanding of the natural numbers from the concept of indefinite iteration, we still need abstraction to get our original idea of a finite sequence and our notion of the indefinite. Indeed the concept of iteration can be incorporated into our understanding of abstraction, since once we have abstracted the basic structure of a finite numerical sequence, we can then extend it by further application of principles that we acquired when we grasped the original numerical sequence. Thus abstraction is an essential part of the process of acquiring our knowledge of mathematics, and since we can use it not only to extract structures from our experience of the world but to extend these structures once we have an abstract idea of them, we rely on abstraction for our understanding of all mathematical concepts.

Having provided an account of how we use abstraction to acquire knowledge of basic mathematical structures, the next chapter explains how we apply abstraction to mathematical concepts in order to come up with those mathematical structures that are not obviously found in the physical world. Once this work has been done, I turn to considerations arising from Skolem's paradox to investigate more closely the nature of mathematical knowledge.

CHAPTER 4: BASIC AND DERIVATIVE STRUCTURES

So far epistemic structuralism has been explained as a theory of mathematical knowledge acquisition, using abstraction by simplification and extension. It has been argued that there are many mathematical structures in the physical world that are available to us as part of our perceptions. However when we work with abstract concepts of mathematical structures we can abstract further from these and arrive at mathematical structures that are not in the physical world, yet we still can say objectively true things about them. This chapter investigates the notion that some mathematical structures are a part of the world, and derivative structures may not be. An explanation is given of how derivative structures are nevertheless objective, and we can make claims about them that are objectively true or false, using bridge laws.

4.1 Two Levels of Mathematical Structures

If we examine the epistemic structuralist account of mathematical knowledge acquisition, it becomes evident that we can divide the truths of mathematics into two categories: those that correspond directly to truths about the physical world, and those that do not. The latter are true of structures that are extensions of (or abstractions from) those that we can directly perceive in our experiences of the world. The first kind, which are about structures that exist in the physical world, are often immediately available as a part of our perceptual experience, combined with inference to the best explanation. Examples of such statements are “ $2 + 2 = 4$ ” or “the sum of the angles of a triangle on a flat surface is 180° .” These statements can be verified by empirical means, by going out and counting groups of objects or by taking a protractor and measuring angles of triangles (allowing, of course, for the slight error-margin of our

perception). These claims of mathematics are constantly verified and never refuted¹ by our perceptual experiences, which makes them somehow built into our experience of the world.² This is because these statements are about *basic mathematical structures*, the mathematical structures that exist in the physical world and are instantiated in physical systems with which we interact every day.

When we make a statement such as “ $2 + 2 = 4$,” which is an abstract³ way of talking about the natural number structure, it can be thought of as shorthand for “whenever you take two objects and add two more objects, you’ll end up with four objects.”⁴ Talking about mathematical structures in this way is significantly different from claiming that mathematical entities are purely abstract. The epistemic version of structuralism that I am advocating has as one of its most basic claims the notion that our knowledge of mathematics originates in our experience of the physical world. In an earlier chapter I explained how we move from our physical experiences to having ideas of mathematical structures, but even when we speak about mathematical structures independently of their instantiations it is not the same as claiming that they are independently existing abstract things in the Platonic sense. The mathematical concepts that we have are just mental copies of the actual structures that exist in the physical world, they exist as ideas in our minds and are given to us via our perceptual experience.

This grounding in the physical is essential because it shows why mathematical truths apply to the physical world: because they are really *about* the physical world itself. Recall that one of the primary benefits of epistemic structuralism is that it is a

¹ As long as all the requisite conditions are fulfilled. In a later chapter we shall see that the Stability and Correct Counting Conditions, as specified in Wilson [2000], must be satisfied in order for experience to contribute toward supporting or refuting a mathematical claim.

² Chapter 8 provides a more precise account of how mathematical truths are ‘built in’ to the world.

³ This is an abstract statement in a non-Platonic sense, recall that we can refer abstractly to mathematical structures without ascribing to them an independent existence.

⁴ Assuming, again, that the required conditions are met.

form of realism, thus retaining the objectivity of mathematical truths; another advantage is that by grounding mathematical knowledge in the physical world it is natural for the truths of mathematics to describe the physical world as well as they do. Concrete or physical exemplifications of mathematical structures are both ontologically and epistemically prior to their abstract counterparts, which exist nowhere but in our minds. Just because we often seem to speak abstractly in mathematics does not mean that mathematical entities are essentially abstract, it just means that this mode of reference can be useful because it allows us to ignore all the extra-structural features of the systems that exemplify a particular structure, and isolate their structural qualities. When we do this directly, without abstracting further or extending the structures, we are referring to basic structures that actually exist in the physical world.

When we take the structures that we perceive directly and extend them as described in the previous chapter, then we use abstraction (in the sense of extension) to acquire concepts of what can be considered to be *derivative structures*. Derivative structures may be defined as mathematical structures that are not found in the physical world, while basic mathematical structures do exist in the physical world. The distinction between basic and derivative structures is not a clear or simple one to make, for two reasons. First, we can make the distinction along either epistemic or ontological lines. Second, it is not clear what it means for a mathematical structure to be found in the physical world.

To make the epistemic distinction basic structures are defined as those that we know to be in the physical world, and derivative structures as those that we cannot find in the physical world. However there may still be derivative structures that are part of the physical world although we are not aware of this fact. Often it is the case that we come up with a derivative structure and then it turns out that it is a part of the physical world. Thus there is a difference between those structures that we *know* to be a part of the physical world, and those that actually *are* a part of the physical world,

regardless of whether we consider them to be basic or derivative. The ontological way of distinguishing between basic and derivative structures is to define basic structures as those that are in the physical world, whether or not we know that they are. It is not important which way we make the distinction, as long as we are aware that there is a difference. When I refer to derivative structures I will generally be using the epistemic distinction, for if we cannot find the structure in the world then for our purposes it is treated as a derivative structure.

This leads to the second problem with distinguishing between basic and derivative structures. There are various ways that we can consider a mathematical structure to be a part of the physical world. A structure such as Euclidean geometry may be considered to be in the world in a very obvious sense, since any flat surface constitutes a Euclidean plane and we can directly observe truths of Euclidean geometry in the physical world. However we have found that space is actually curved, so Euclidean geometry does not accurately reflect the physical world. For many numbers, the question of their existence is not clear cut: π is an irrational number so we cannot ever specify it completely in numerical terms, yet it describes the ratio between the circumference and the diameter of a circle. Complex or imaginary numbers such as $\sqrt{-1}$ do not seem to exist in the physical world, yet they are used in calculations to describe physical phenomena.

Where we draw the line between basic and derivative structures is not crucial. A broad conception of basic structures could be used, encompassing all structures that have any bearing on the physical world at all, including instrumental use which is discussed in a later section. Alternatively a very narrow understanding of 'basic' could be employed, which pushes more structures into the derivative category. Where we draw the line, or indeed whether we make a clear demarcation at all, is not so much a philosophical problem as a scientific one, so it is not crucial to my project. What is important is that some mathematical structures are basic and some are derivative, and we need to explain how we come up with derivative structures and the

mechanism by which we can state objective truths about them. It is important to have some link to the physical world in order to explain both their objectivity and the fact that we seem to find many derivative structures in the world after all.

4.2 From Basic to Derivative Structures

There are two main ways that we extend basic structures to get derivative structures.⁵ The first is by recombining the concepts that are involved in the basic structures. For example from the natural numbers we have the operations of addition and multiplication as well as their inverses, subtraction and division. When we apply these to the natural numbers in a way that does not maintain closure, for example by dividing 10 by 3, we acquire concepts of derivative numerical structures, such as the rational numbers. Another example of a derivative structure arrived at by recombining the concepts we have in the basic structures is a higher dimensional geometry. We experience three-dimensional space, as well as perceiving two-dimensional systems (for example, drawn on paper) and one-dimensional systems. Thus we have the concept of a dimension and understand what it is for one space or manifold to have greater or fewer dimensions than another. Using these concepts we can imagine a 15-dimensional space, for example. By studying the relation between a circle and a sphere, we can extend these basic structures and consider a 15-dimensional sphere. We can define these kinds of entities because they are made up of essentially the same concepts that are involved in basic structures, only they are used in a combination that is not (or at least is not believed to be) found in the physical world. This recombination of concepts is generally the process by which we acquire knowledge

⁵ There could be more ways that mathematicians move from basic to derivative structures, this account is not supposed to be exhaustive but indicative of how the process works.

about mathematical structures that are uncontroversially derivative, since it can yield structures that are in no sense instantiated in the physical world.

Another way that we can extend basic structures to obtain derivative ones is by applying further abstraction or simplification to basic structures. We leave out certain features of the system and come up with a more abstract one. An example of this type of extension is group theory. If we start with arithmetic or basic set theory, we have operations such as addition and multiplication on a given domain. To develop a group we leave out certain features of the operations, and say that we have just one general operation. It could be either addition or multiplication, but we study it in a more generalised, less specific way, so the particulars do not concern us. Continuing in this way we could leave out even more specific features of the system to come up with a more highly abstract one, such as category theory. We can see that set theory, group theory and category theory form a hierarchy of increasing generality. All three structures, namely a set, a group and a category, could be instantiated by the same concrete system. Each theory focuses on different features of the system, depending on which features we are interested in. The more general the theory we use, the more isomorphic systems we can find because a more general theory, such as category theory, has less specific criteria than set theory.

This kind of extension by simplification is analogous to the abstraction by which we first come to know the basic structures that are found in the physical world, as we saw in the previous chapter. When we abstract from arithmetic or set theory to develop the notion of group theory, we are leaving out some specific aspects of the system in order to consider a more general structure. This is the same process that we use to abstract a basic structure from our experience of the physical world. Just as a stack of coins possesses features that do not concern us when we are interested in the properties of a finite sequence that it possesses, so there are features of sets that do not concern us if we are interested in exploring group theory. Once we have group theory we can abstract by simplification even further to get to category theory. This is how a

single physical system does in fact exemplify various mathematical structures, depending on how far we take the abstraction. In other words, we can abstract from perceptions of a certain physical system and then apply the same kind of simplifying abstraction to the resulting mathematical system, and both of the resulting mathematical structures will be instantiated in the original physical system. This method of generating derivative structures tends to yield structures that are disputably derivative, since they are more removed from the physical system than the structure from which they were abstracted, however they can still be considered to be grounded in the physical world. Whether or not we call them derivative is not significant, the main point is that the way that we come up with them gives them a link to the world that has been found reliably to produce objective claims about them.

4.3 Ontological Status of Derivative Structures

There are many different derivative structures we could abstract from basic structures that we have abstracted from the physical world, however some of these extensions are better at describing the world than others. Thus there is a sense in which some derivative structures are better than others, because they are more useful in describing the world. A later section in this chapter deals with the instrumental value of derivative structures, and this is one way of finding or locating some derivative structures in the physical world. Since the derivative structures are grounded in basic structures they are in a sense contained in them, in the same way that the average commuter is contained in a list of all commuters. Derivative structures are real and objective because they originated from basic structures, but they are captured by concepts that increase our understanding of the world.

There are also derivative structures that do not appear to have instrumental use. These are less obviously connected to the physical world, but we still want to be able to say things about them that are objectively true or false. A later section in this chapter deals with the semantics of statements about such derivative structures, but

here my concern is what types of entities they are. They cannot be physical structures since we do not know whether they are instantiated in the physical world (or since, as far as we know, they are not in the physical world). To maintain consistency with epistemic structuralism derivative structures cannot be Platonically abstract entities. While they are derived from basic mathematical structures, if they do not exist in the world then derivative structures must be fictional entities. Lehman⁶ discusses various views along these lines. If-thenism is the claim that mathematical statements are conditionals, which allows them to be true even if mathematical entities do not exist. Russell's postulationism is a similar theory, namely that the axioms of mathematics define a structure, but carry no ontological or existential implications. The deductions from these axioms then constitute the truths of mathematics.

The view of derivative structures that I wish to put forward is related to these theories but with an important difference. If mathematical entities do not in fact exist then it is difficult for if-thenism to account for our conviction that there are both mathematical truths and mathematical falsehoods. If all mathematical claims are conditionals whose antecedents claim the existence of fictional items, then any such claim would be trivially true.⁷ Postulationism runs into difficulties with applied mathematics; like formalism, it can account for why we have the mathematical truths that we have (because they are deduced from the axioms), however if the axioms are arbitrary then why do mathematical truths describe the world? The solution to this difficulty is to ground the axioms in the physical world. When we postulate derivative structures like 28-dimensional cubes and prove a result about them, the resulting truth is a conditional whose antecedent is the claim that there is such a thing as a 28-dimensional cube. While as far as we know this is a false claim, the properties of a 28-

⁶ Lehman [1979].

⁷ It is acknowledged that such claims are trivially true in classical logic, but this is not necessarily the case in alternative logics such as relevance logics or paraconsistent logics.

dimensional cube are derived from what we know about geometry and dimensionality from our experience of the world, so it is not a purely random or arbitrary claim. Since the antecedent of the conditional is composed of concepts that directly capture aspects of the physical world, the derivative structure is not a pure fantasy. Their derivation from the physical world is what gives derivative structures their objectivity, their grounding in reality. In a later section of this chapter I will use Hinckfuss' notion of bridge laws to show how we can make true claims about derivative structures that may not exist.

Basic mathematical structures are features of the world, and derivative structures are abstracted from these. We often derive these extensions in order to gain added insight about the mathematical structures of the world. Truths about basic structures translate directly into truths about the world, while truths about derivative structures may or may not. It is often the case that derivative structures that we gain by abstraction are then found to be extremely useful in describing the world. An example of such a mathematical structure is the imaginary or complex numbers. The number -1 has no square root, but mathematicians have postulated the existence of the square root of -1 . By giving the square root of -1 a name (" i ") they are able to use it in calculations which can be useful not only in describing the complex numbers, but also for explaining certain problems that are unsolvable in the domain of the real numbers. These problems originate in the real numbers, but since they have no solutions in that domain, the complex numbers are a useful tool that give us a better understanding of certain problems than if we were restricted to the real numbers. An example of this sort of problem is quadratic equations whose coefficients are all real, but which have complex roots. A quadratic equation has the form $y = ax^2 + bx + c$, and to find the roots of this equation (values of x for which $y = 0$) we can apply the quadratic formula: $x = -b \pm \sqrt{(b^2 - 4ac)}/2a$. If $b^2 < 4ac$ we will have a negative number under the square root sign, and so the equation will have its solution in the complex numbers.

While we can never find i in the physical world as numbering some objects or as a measurement, it has instrumental value because we can use it to solve equations with which we describe the world. Since i is applicable to physical systems, it is a mathematically useful concept in the same way that averages are. We can make many claims about complex numbers, in the same way that we can make many claims about the average commuter. Say the average commuter drives for one hour and forty five minutes every day. This is a true statement, and useful in determining how much traffic is on the roads during peak hour and so on, however we cannot go out and interview the average commuter. Many derivative structures have this sort of status; we can say many things of them and they can increase our understanding of the world, although they do not exist as things in themselves in the way that basic structures do. However they are not purely fanciful entities, the claims we make about them are objectively true or false since they are based on physical facts, in the same way that claims about the average commuter are grounded in concrete facts.

How is it that we come to have the derivative structures we do? We explore extensions of (or abstractions from) the basic structures, whether because of a need to gain increased insight into the workings of a basic structure, or because of curiosity. As we explore various extensions we often find that some are superior to others, in that they tell us more about the world, or perhaps they possess features which are interesting in their own right. Utilising the same methods that we do when reasoning about basic structures, we can discover truths about derivative structures. These facts still retain their objectivity, because they are grounded in the physical world in the same way that knowledge by description is grounded in knowledge of things with which we are directly acquainted.⁸ The fact that derivative structures arise out of basic

⁸ Russell [1912] introduces the notions of knowledge by acquaintance and knowledge by description. This distinction is also discussed later in this chapter.

structures means that they can give us insight into physical phenomena, even though their truths might not directly tell us anything about the world.

As Kitcher points out, the success of a theory in helping us to understand the world is more important than how we came up with the theory in the first place.⁹ As we successfully apply a theory, the justification of the theory and the justification for its use are based increasingly on its successful use, rather than on the principles that generated it in the first place. However the process of abstraction that let us get to the derivative structure is still significant because it provides the link that explains the applicability of the theory which pertains to that structure. As for how we came to use one particular derivative structure instead of another, this is a question that is not the concern of my project. Kitcher correctly identifies such questions as being primarily the domain of psychologists and philosophers of mind, and as these are not the topics of this thesis I leave these questions to others. The only way to decide on a particular derivative structure over another, which is pertinent to current concerns, is that one extension or abstraction from a basic structure may be more powerful in accounting for phenomena that we desire to explain.

In one sense we have a great deal of freedom when deciding which derivative structures to use, however we strive to find the structures which have greatest explanatory power and are the most useful in describing our world. When faced with a limitation of our mathematical theory, for example the natural numbers' inability to solve $7 - 12$, we might investigate various different derivative structures as possible extensions which are available to us. Postulating the negative integers is one option, another might be to introduce a number called "*nil*" which is the answer to any calculation that has no solution in the natural numbers. So $7 - 12 = nil$. However this extension might not be the best we can find. For example, it gives us the result

⁹ Kitcher [1983], pp. 9-10.

$7 - 12 = 7 - 50$. This result is less powerful than the result that $7 - 12$ is 38 more than $7 - 50$, because the latter lets us distinguish between various numbers that all would have been called *nil* under the other extension, and this characteristic makes it a more effective theory (in this context). We are likely to find that the structure of the natural numbers plus *nil* is not much more useful in describing the world than the natural numbers alone. In this example the negative numbers are a more powerful extension, because they increase the explanatory power of our theory and this extension is confirmed by its applicability to our experience. In this way we are led to certain derivative structures, which can be considered a part of the world in the sense that they are abstracted from the basic structures. This shows why truths about derivative structures are objective, because if we could invent derivative structures with complete randomness, there would be no connection to the world and the objectivity of mathematical truth would be lost. Even those derivative structures that do not appear to have any connection to the physical world are grounded in basic structures, and this gives them their objectivity.

The notion that derivative structures are connected to the physical world and are objective is supported by the story of the Indian mathematician Ramanujan. He was exceptional in that he had no formal mathematical training, indeed Hardy who became his mentor and wrote about his life¹⁰ estimated that he had access to only one or two mathematical texts. Despite this lack of training, Ramanujan discovered many theorems in algebra, and was eventually brought to Cambridge to work with Hardy. It seemed that Ramanujan had no concept of proof, but was so mathematically gifted that he was able to intuit some results of higher mathematics. Hardy introduced him to the notion of proof and during his time at Cambridge he worked on proofs for his theorems, some of which were already known but many of which constituted new

¹⁰ Hardy [1978].

advances. Many of these are now proved but some of them remain unproved to this day, and some of his conjectures have now been proven false.¹¹

The fact that Ramanujan was able to work in isolation and come up with theorems that other mathematicians had already proved suggests that the structures he was working with are an objective feature of the world, and that we do not pick and choose randomly which derivative structures to work with. Certain derivative structures have greater instrumental value than others. When it becomes evident to us (or to a mathematician) that we require a certain structure to make sense of a particular phenomenon, we could test various abstractions from the ones we have encountered thus far. As in the example above, there is generally an extension which seems the best or the most natural, which has the greatest instrumental value in explaining the phenomenon (or phenomena) that led to the dilemma. This extension is the derivative structure that we will end up using; while it is not necessary that we choose a particular one, there is often one which seems the most natural. We can still consider such derivative structures to be objective, because they are in a sense embedded in the basic structures and can be described in terms of them.

The view that derivative structures are objective is supported by the fact that nature continues to display mathematical structures that we do not necessarily expect it to, for example fractals which can be thought up as abstract geometrical concepts and are useful in their own right, yet are also found in a great multitude of instances in the natural world (for example in certain types of vegetation). Another good example is $\sqrt{2}$: this is an irrational number that we can arrive at by reversing the operation of squaring. However, we can also get $\sqrt{2}$ by drawing a right angle triangle such that the two perpendicular sides are each one unit of length, and then the hypotenuse will be $\sqrt{2}$ units long. The fact that an irrational number, which is a part of derivative

¹¹ Kanigel [1991].

structure, arises in a physical form indicates that some derivative structures are built in to the physical world. Once we are aware of derivative structures we often start to find them in physical systems, for example we can be -5 kilometres west of a marker, or halfway to a destination. In these cases the integers and rational numbers give us a more concise, better way of capturing mathematical phenomena. While we could describe these situations in terms of basic structures, derivative structures give us a more sophisticated way of relating the same information.

This indicates that many derivative structures, as well as the basic ones, are a part of how the world is, although they are available to us in subtler ways than basic structures which are more obviously observable as elements of our perceptual experience. Since derivative structures are not purely created by us but are extensions of basic structures arising out of our manipulations of the structures we find in the physical world, statements about them can still be considered to be objective in the same way that statements about constructions such as averages are. Their objectivity arises from the fact that derivative structures are natural extensions of the basic structures, constructed from elements of our experience of the world. Indeed, this claim is supported by the remarkable concurrence among mathematicians (surely greater than among any other scientific communities) and by stories such as that of Ramanujan, in which mathematicians working independently in vastly diverse geographical locations are able to come up with the very same results as each other.

4.4 Instrumental Use of Derivative Mathematical Concepts

We saw earlier that the distinction between basic and derivative structures is not entirely clear-cut for two reasons. One is that we cannot always know whether a mathematical structure is in the world or not, and the other reason is that it is unclear what it means for a structure to be found in the physical world. Many structures are obviously in the physical world, for example various sub-segments of the natural numbers, but some structures are found in the world in an instrumental sense,

although they are not given to us directly in perception. However there are also purely derivative structures which do not appear to have any bearing on the physical world, and these exist only as concepts (as far as we know). They are abstract in the sense that they are mental entities, rather than in a Platonic sense, but there is always the possibility that we could one day find an application that they have in the world.

Some derivative structures (not all) can be viewed as instruments that tell us things about the world. Since they can be defined in terms of actual things in the world they are objective, even though we only have access to them as concepts. Instrumental use is one way of highlighting the objectivity of *some* derivative structures, however they are still derived in the same way as all others. Derivative structures that have instrumental value are not *more* objective than those that do not, however the fact that many are useful in this way shows that the methods we use are reliable.

One example of an entity that has instrumental value but cannot be found in the world is the average commuter. Information about the attributes and habits of the average commuter lets us make fairly accurate estimates about road usage, traffic flow, pollution and so on. However we can never observe or interview the average commuter, as we know this entity only as a concept. The concept we have of the average commuter is dependent on us for its existence, however once the concept is defined we can make statements about the average commuter which are true or false in virtue of independent facts in the world. The average commuter is a theoretical entity constructed out of real, objective facts about the world. Although the average commuter is a concept rather than a thing in the world, this does not preclude us from saying objectively true things about him or her, since the concept is constructed out of things that do exist in the physical world. It can be considered to be a kind of shorthand for conveying information about a large group of real, existent people. When we claim that the average commuter travels for one and a quarter hours each day, that is an abbreviated way of saying that Commuter 1 travels for one hour and ten

minutes, Commuter 2 travels for half an hour, Commuter 3 travels for two hours ... and so on. Using the concept of an average we can convey the same information in a summarised form that is much more useful than providing a list of the habits of thousands of motorists.

Russell's¹² distinction between knowledge by acquaintance and knowledge by description is helpful for understanding basic and derivative structures. Knowledge by acquaintance is the kind of knowledge we have of something we experience, hence it is knowledge of that of which we are directly aware. Knowledge by description is the knowledge we have of something with which we are not directly acquainted, but that we can describe in terms of things with which we are acquainted. Russell gives the example of being in the presence of a table: we have only knowledge by description of the table, since we do not know it directly. What we do know directly, what we are acquainted with, are the sense data associated with the table. We are acquainted with the colour brown, a flat surface and so on, and we have knowledge by acquaintance of all of these sense data. In contrast, our knowledge of the table itself is only via our knowledge of these sense data, so it is knowledge by description. We know the table in the sense that we can describe it as being brown and flat, but we do not have direct or immediate knowledge of it because we are not acquainted with the table itself.

An important corollary of this distinction is that acquaintance is a crucial factor in the formation of all our knowledge. Anything that we know by description can be traced back to things we know by acquaintance; the terms in which something is described must be reducible to things with which we are acquainted. While this is not the foundation for our knowledge, since we need concepts and reason in order to become acquainted with anything in the first place, acquaintance is an essential component of our knowledge of the world. There are varying degrees by which an

¹² Russell [1912].

item that we know by description can be removed from the items that we are acquainted with, the constituents of a description need not all be particulars that we are acquainted with, but they must all be ultimately reducible to things that we know by acquaintance. Russell gives the example of Julius Caesar¹³ as something that we know only by description, because evidently we could not be directly acquainted with Julius Caesar. We might, for example, know him as “the founder of the Roman Empire” (or some other description). The Roman Empire in turn is not something with which we are acquainted, but the description we associate with the Roman Empire could be “the ancient civilisation I read about in a history book,” and this description provides the link to something with which we are directly acquainted, namely a history book, or at least its representative sense data. We might also consider a description of the Roman Empire that arises from a visit to Rome. Different people will have different descriptions for the same thing, but all descriptions are ultimately reducible to things that we know by acquaintance.

This analysis is a useful analogy that helps to clarify the distinction between basic and derivative structures: we know basic structures by acquaintance because we experience them (or at least, aspects of them¹⁴), and derivative structures are known by descriptions which are ultimately reducible to basic structures. In other words, derivative structures are defined in terms of concepts that we have experienced, or been acquainted with, through our knowledge of basic structures. However Russell’s distinction is merely a useful analogy for the relation between basic and derivative structures, it is not a characterisation of our knowledge of mathematical structures. Indeed, it makes a good analogy for various stages of mathematical knowledge.

¹³ Russell [1912], pp. 58-59.

¹⁴ We do not experience basic structures completely, our sense data acquaint us with only an aspect of the structure from which we infer the whole, in the same way that we infer an object from our perception of some aspects of it (its colour, shape, etc.). The incompleteness of our acquaintance with basic mathematical structures is discussed in a later chapter.

Whenever we perform an abstraction we begin with either sense data or some mathematical concepts with which we are acquainted, and infer a mathematical structure that can be described in terms of what we were abstracting from. For example we are acquainted with notions of both spheres and circles, having abstracted from many perceptions of both these shapes, and from these concepts we can abstract to the concept of a four-dimensional sphere. This entity can be described in terms of spheres and circles, as it has many properties in common with both of these and stands in the same relationship to a sphere as the sphere does to the circle with respect to dimensionality. The example below demonstrates further the process of abstracting from one mathematical structure to get to another.

4.5 Example: Number Systems

A good example of basic and derivative structures is the various number structures. The natural numbers are an example of a basic structure, because they are given to us in our experience of any group of discrete items. However even this fairly obvious and simple example is not completely straightforward. It is only finite sub-sequences of the natural numbers that can be considered truly basic structures, because we can never apprehend an infinite number of objects. In fact it is unclear whether there exist an infinite number of things, so it may even be the case that the entire natural number structure does not exist. However, we saw in the section on abstraction how we can start with the finite initial sub-sequences of the natural numbers that we get from experience, and extend them to get an idea of the entire natural number structure. Although the whole of the natural number structure is never given to us in experience, we grasp the indefinite extensibility of finite sub-sequences of the natural numbers by interacting with the sub-sequences and observing that we can always increase a collection by one. This gives us an idea of the entire natural number structure. The natural number structure can still be considered a basic structure since segments of it are instantiated by physical objects and we can observe

and interact with its concrete exemplifications. Also we can abstract away the concept of a last element to consider the structure in its entirety. Infinity certainly is a part of the natural number structure and is something that we cannot observe (and which may not exist),¹⁵ however we experience many instantiations of initial sub-sequences of the natural numbers and can grasp their indefinite extensibility by leaving out (abstracting away) the notion of a finite endpoint. This allows us to get an idea of the natural number structure which forms the basis for all other numerical structures.

When we manipulate physical systems that instantiate the natural number structure, there are certain operations which are quite natural to perform on numbers. Addition and subtraction are probably the most obvious, because to any collection of items we can add more items, and the resulting collection of items exemplifies a new number which is the sum of the original two groups. Just as natural as this operation is its reverse, and thus we also have the operation of subtraction. When we subtract a greater number from a smaller number the result is not within the natural numbers, and this is how we are led to extend the natural number structure and get the integers. Negative integers are not obviously instantiated in physical systems, in the sense that we can never count -5 objects. Nevertheless, having come up with this derivative structure we find that we do in fact use the integers in many real-life applications. For instance, they may be used when we ascribe a direction to a measurement, for example if I wanted to say how many metres north of the lamppost I am standing. If I start off 3 metres north of the lamppost and walk south for 5 metres, then I will be -2 metres north of the lamppost (or 2 metres south of it). The negative numbers are also used to express debts, for example the fact that I owe someone 5 dollars leaves me with -5 dollars. While we are never faced with -2 chairs, the negative numbers have

¹⁵ A later section on infinity in this chapter goes into this issue in more detail.

other kinds of physical applications and can in this way be considered a part of the physical world.

The same sort of story is the case for the rational numbers: multiplication comes from manipulation of physical systems (for example we might wonder how many pebbles it takes to make a rectangle with sides of 3 and 7 pebbles, which gives us " $3 \times 7 = 21$ "). Once we have grasped the concept in its physical form we can consider the notion of multiplication independently of grids of pebbles, and this gives us an abstract idea of the operation itself. When we have that, we also have its reverse, namely division, and by applying that to numbers that do not divide into each other with no remainder we get the rational numbers. As with the negative numbers, we come up with this structure in order to explain applications in which we already use the rationals. This derivative structure has enormous instrumental value, because in our applications of mathematical theory to the physical world we often use ratios, and in the cases where numbers do not divide evenly into each other it is useful to have a coherent way of completing the calculations. The rational numbers are also used in more direct physical applications, for example we could observe that a container is $\frac{3}{4}$ full. If the container can hold 12 units of fluid, we could also express this truth by saying that it now has 9 units in it. Once we have the rational numbers we have an alternate way of expressing this physical fact which might be more informative. Issues regarding various notations will be discussed in greater depth in Chapter 7.

A similar story about the operations of squaring and square rooting gives us some of the irrational numbers such as $\sqrt{2}$. Other irrational numbers are given to us by other means, for example π , which comes from dividing the diameter of a circle into its circumference. This last example helps to bring out an important aspect of the process of extension from basic to derivative structures. Working with a basic structure and coming up with a derivative structure is not a purely creative process; we have arrived at the derivative structures that we have because of a need to cope with something that arises out of our perceptual experiences. For example we are

given addition and thus subtraction from our interactions with the world, and in certain circumstances they do not yield an answer within the basic structure of the natural numbers. However, because there are situations that we need to describe such as being in debt or walking backwards, we are led to the concept of negative integers. While it is not a matter of necessity that we have the derivative structures that we do, there does seem to be a fairly natural way of extending basic structures in order to explain physical phenomena. We could experiment with various extensions and find that a particular one is found in the actual world (such as the negative numbers), while others might have significant instrumental value (such as statistics). Whether we include instrumentally useful structures in the category of basic or derivative is unimportant for present considerations. The crucial point is that *all* mathematical structures are abstracted, whether directly or by a chain of abstraction, from our experience of the physical world.

4.6 Infinity

While the natural number structure can be considered a basic structure since we observe segments of it all the time, it could be argued that infinity is a part of the natural number structure and since we cannot observe infinity as a completed totality, we cannot even consider the natural numbers to be a part of the physical world. There is some contention as to whether or not we can observe infinity, since we sometimes have the impression of experiencing infinity. Usually this is when we experience something very large, seemingly without end, or else the other aspect of infinity, namely the infinitesimal. A finite interval can be divided up into smaller and smaller intervals, with no apparent limit. It is argued that by considering all the tiny segments that make up a given interval, we can perceive infinity as a completed totality.

Despite these impressions that resemble infinity, I agree with Hilbert that we cannot really experience the infinite as a completed totality, and that the infinite may not exist as a part of the physical world. Consider our purported experiences of

infinity, which consist either of the infinitely large or the infinitely small. When we experience the infinitely large, the experience is usually of an unbounded and vast collection or distance, such as the stars, the grains of sand on the earth, or the distance one could possibly travel in a straight line up into the sky and away from the earth. None of these could be considered experiences of infinity as a completed totality, because the limits of our perception prevent us from experiencing an infinitely large collection or distance. We are given the impression of the infinitely large by the unboundedness of the collection, we can imagine that (technology and lifespan permitting) we could travel away from the earth without ever reaching a boundary, and it seems that we could never finish counting the stars or the grains of sand that make up the earth. This means that we cannot *actually* perceive infinity as a completed totality, but we can experience *potential* infinity in the indefinitely large.

The notion of potential infinity originated with Aristotle in the *Physics*. It was Cantor who realised that Aristotle denied the existence of actual infinity, and that he claimed the needs of mathematics were served by potential infinity. There was little study of infinity from the time of the ancient Greeks until 1784, when the Berlin Academy prize proposal called for a precise characterisation of infinity, in recognition of the use of infinitely large and infinitesimally small quantities in mathematics and the lack of an adequate account of these principles. Lagrange and other judges of the Berlin Academy prize credited the ancient Greeks with meticulous avoidance of actual infinity, because of their use of the exhaustion method. This method historically justified Aristotle's claim that Greek geometry does not require actual infinity, however this claim has been strongly attacked.¹⁶ Late in the nineteenth century infinity became a crucial concept, since the characterisation of real numbers

¹⁶ For example, Cantor argued against the sufficiency of potential infinity for mathematical purposes. Aristotle's notion of potential infinity was also attacked by Knorr and Hintikka, among others. See Kouremenos [1995] for further details on Aristotle's conception of infinity and surrounding debate.

that emerged introduced completed, or actual, infinities into mathematics. Cantor argued that use of potential infinity required the existence of actual infinity, which would then be ontologically prior. Gauss, however, strongly objected to the existence of actual infinity and its use in mathematics. Aristotle's characterisation of infinity as potential, rather than actual, is an enduring notion that has proven difficult to refute.¹⁷

Some might agree that we cannot perceive the infinitely large as an actual completed totality, but claim that we can experience a completed actual infinity in the infinitely small. By considering all the points in a given finite interval, we can experience the infinitesimal as a completed totality. However, I would argue that even in the case of the infinitely small, we do not genuinely experience infinity. We can recognise the *potential* for dividing the interval into smaller and smaller segments without reaching a limiting size, but we do not perceive each of the infinite number of points within the interval. We are limited by our perceptual apparatus, and can only divide the segments so far. Again, all that we experience is the potential for infinity, because we realise that in theory there is no limit to how many times we can subdivide the interval, however we cannot experience all the points in the segment. Thus we are once more given an idea of the infinite by realising the potential for infinitely small segments, but we do not experience actual infinity in the number of subintervals.

Both these ideas of infinity, the infinitely large and the infinitely small, are the result of our recognising the potential for the infinite or the infinitesimal when we experience something very large or very small. In both cases we do not experience actual infinity in its entirety, but we cannot perceive any bounds to how large or how small something can be. It should also be noted that there is some evidence from modern physics to suggest that neither the infinite nor the infinitesimal exist in the

¹⁷ Kouremenos [1995] provides a detailed account of Aristotle's characterisation of infinity.

physical world. Investigations into the size of the universe and the curvature of space have led physicists to postulate a finite but unbounded model of the universe, in which space curves back onto itself. To illustrate and clarify this model of the universe, a two-dimensional analogy is made with the surface of a sphere, which is an example of a space that is finite in magnitude but has no boundary. If the universe is finite in a similar way, that would imply that there is a limit to how long a distance can be, or how many particles exist. The distances and numbers involved would be so enormous that we could not easily imagine them, which is why we identify very large things with the infinite. Nevertheless, if this model of the universe were true then there would be some limit and thus actual infinity would not exist.

Another argument against the existence of infinity as a part of the natural number structure comes from van Dantzig. In his paper “Is $10^{10^{10}}$ a Finite Number?”¹⁸ he argues that infinity is not a natural number, and that the largest natural number is actually smaller than $10^{10^{10}}$. He explains that if the natural numbers are those that we can come up with by incrementing numbers that we are acquainted with, modern physics dictates that this process is limited by time and so there is an upper limit to the natural numbers. We can get to numbers larger than those we can construct in this way if we add together previously constructed numbers, however van Dantzig claims that “[t]hese sums do not exist as natural numbers in the first sense, but only in a new sense.”¹⁹ He believes that the properties of natural numbers cannot be considered to be proven for these larger numbers, we can only postulate that the large sums have the same properties as the directly constructible natural numbers. Using multiplication, taking powers and so on, we can get to higher-order classes of numbers, but van Dantzig claims that none of these can be considered to be finite numbers. We can describe and refer to infinite or transfinite numbers, but they cannot be considered to

¹⁸ van Dantzig [1955].

¹⁹ van Dantzig [1955], p. 274.

be finite or natural numbers in the original sense, so we cannot assume that they have all the familiar properties of natural numbers.

There is also some physical evidence to suggest that the infinitesimal may not exist in the physical world, that there is some limit to how far space can be broken down. When a space gets smaller and smaller, there is a point at which the standard laws of physics break down and no longer apply, space and time separate and the rules of quantum gravity take over. This happens at the horizon of a singularity, and is associated with the Planck-Wheeler area, $2.61 \times 10^{-66} \text{cm}^2$.²⁰ This area is so much smaller than anything that we can observe that we could never come close (in our ordinary interactions with the world) to experiencing a limit to how small space can be, which is why we have the impression that space is indefinitely divisible. However this finding suggests that the infinitely small may not exist, that there is a limit to how far space can be divided up. While I do not claim that it is conclusive, there is some physical evidence to suggest that the infinitely large and the infinitely small may not exist in the physical world, that there could be a limit on how large or how small something can be. The reason we intuitively think otherwise is that we experience *potential* infinity when we encounter something very large or very small, since these limits are so far beyond the reach of our perceptual faculties. However we do not ever experience actual infinity as a completed totality, and there is some physical evidence to suggest that actual infinity may not exist.

Given that we have no impression of actual infinity, we still have the idea of infinity as a completed totality which functions as a useful entity in our mathematics. The natural numbers as an indefinitely extensible sequence are a basic structure, given to us as a part of our experience of the world. As Hilbert²¹ explains, we introduce the notion of actual infinity in order to make our system work better. We do not perceive

²⁰ Thorne [1994].

²¹ Hilbert [1925].

infinity, but in our experience (although perhaps not in reality) there is no largest natural number. We imagine that there is one, and we call it infinity, and now we can use it in our mathematical system. Hilbert explains that the role that the infinite plays is that of an idea, “a concept of reason which transcends all experience and which completes the concrete as a totality.”²² It is useful in the same way that i , the square root of -1 , is useful. Both i and ∞ are extensions that we have added to more basic numerical structures, and although neither may be found in the world, both have useful functions that enable us to discover more about the world. As Hilbert points out, we introduced i in order to preserve the laws of algebra, and for analogous reasons we introduce the notion of an unending straight line in geometry. These are both ideal elements, concepts in our thinking that we use in mathematics to make the system work better. Similarly, the notion of a completed infinity is an ideal element that we have introduced to preserve our logic, even though it is not something we can experience and may not correspond to an actual entity in the world.

If indeed infinity is not to be found in the physical world, then it is a purely fictional mathematical element that we have derived from the mathematical concepts acquired from our experience of the world. In the next section of this chapter I will explain how Hinckfuss’ mechanism of bridge laws can be used to make true claims about infinity even if it does not exist.

In the same way as with i or a 15-dimensional cube, we are not acquainted with infinity but know it only by description. We can never experience infinity, but based on our experiences of initial sub-sequences of the natural numbers, we can introduce the concept of infinity as a completed totality because our mathematical system benefits from being able to refer to “the number that is larger than all other numbers.” Since we can never experience infinity and cannot find it in the physical

²² Hilbert [1925], p. 201.

world (if it is there at all), the only way we can refer to it is via the numbers that we know. In our experiences of the world we encounter many countable collections of things, and by augmenting these collections it soon becomes evident to us that there are indefinitely many numbers, that we will never run out of them. However we can refer to infinity by making use of the procedure of counting, with which we are familiar. We can imagine that someone (or something) counted with increasing speed, counting “1” half a second after starting, “2” a quarter of a second later, “3” an eighth of a second later, and so on. Once a whole second had passed, they would have counted to infinity. Of course this counting procedure is impossible to perform, even for a machine, but by specifying the procedure we can indicate what we mean by ‘infinity’.

4.7 Truth in Derivative Structures

While the grounding of mathematical structures in the physical world preserves the objectivity of mathematical truth, it is less intuitive when faced with derivative structures that are obviously not found in reality. It is not clear what a statement about a 15-dimensional sphere means, and we certainly cannot find evidence of any such entities; still we have some sense that statements about it are objective. When we are faced with a truth in a derivative structure, such as a 15-dimensional geometry, we can say that it is true not of our physical world, but that it is hypothetically true. Derivative structures that have no instantiations in the world are known to us purely as mental entities; they do not exist in an abstract realm, only as ideas or concepts in our minds. We get them by abstracting from structures that we get directly from the world, and we can derive statements about the relevant derivative structures. These kinds of statements express relations in structures that are not found in the world, but are still grounded in those that are. For example, a statement about the properties of a 15-dimensional sphere does not tell us about anything that exists. This sort of statement can be thought of as an hypothetical

universally quantified statement: $\forall x[S_{15}(x) \supset P(x)]$.²³ However, we can know truths about this structure with a high degree of certainty because we know many things about 15-dimensional geometry, and the reason for this is that we live in a 3-dimensional geometry. We can observe 1- and 2-dimensional geometries by focusing on certain types of spaces, and by comparing these, we can speculate on what sorts of truths would hold in a 15-dimensional geometry. This system is fictional and hypothetical, but it still has its roots in the physical world.

Hilbert argued that the world is finite, and that infinity is a concept that we introduce in order to make our mathematical system work better. In order to carry the certainty attributed to mathematical claims from the finite, which was a real part of the world, to the infinite which did not really exist, Hilbert attempted to provide a way of getting from the finite to the infinite that preserved truth. He aimed to demonstrate that infinitary mathematics did not reflect reality, but was an instrument that we use. In order to make this claim but retain the certainty of all mathematics, he had to show that infinitary mathematics was a reliable instrument. Ultimately his program failed, because of the incompleteness and undecidability results of the first part of the twentieth century. This meant that we could not get from finitary to infinitary mathematics and retain absolute certainty.

The failure of Hilbert's program need not rule out the possibility that infinitary mathematics is a reliable instrument. He was unable to demonstrate that we can use the mathematics of the infinite with absolute certainty, however if we remove the requirement for absolute certainty then infinitary mathematics can still be considered a reliable instrument. Despite the fact that we cannot *prove* its *absolute* reliability, we have a lot of evidence that the infinite can be used in mathematics to provide reliable results. The fact that we do use infinitary mathematics and it is reliable in telling us

²³ For all x , if x is a fifteen-dimensional sphere, then x has the property P .

true things about the world constitutes evidence of its reliability. It may not be absolutely conclusive evidence, but it is good evidence nonetheless.

Hinckfuss²⁴ provides a way of making assertions using an extravagant (and false) ontology, but using a mechanism of bridge laws we can do so without committing to the extravagant ontology nor any other false claims. This is how we can make claims about derivative structures that may not exist, however the claims are still meaningful and do not need to be reduced to claims about basic structures in order to explain their truth or falsity. We need not commit to the existence of derivative structures in order to make true claims about them. Hinckfuss distinguishes between presuppositions and assertion commitments: a presupposition is a common supposition to most dialogues within a community, where “[s]uppositions are to the locutions they govern as the antecedent of a conditional locution is to its consequent.”²⁵ He defines an assertion commitment that one has with respect to a proposition p as when “other participants could reasonably expect one to assert p , given the history of the dialogue, were one asked whether or not p were true.”²⁶ Suppositions do not have to be true in order to operate within a dialogue that makes true claims, however in order to make true or useful claims one must seek to avoid false assertion commitments. In the case of statements about derivative structures the presuppositions in question are the existence of the derivative structure as an objective aspect of the physical world. We can make claims governed by these presuppositions without being committed to the assertion commitment of their existence.

Hinckfuss’ notion of presuppositions governing statements lets us make claims that refer to a more extravagant ontology than one to which we wish to commit

²⁴ Hinckfuss [1993].

²⁵ Hinckfuss [1993], p. 600.

²⁶ Hinckfuss [1993], p. 598.

ourselves. The following extract explains the mechanism of bridge laws which facilitate this:

We have some economical propositions, $E_1, \dots, E_{n(e)}$. We have some bridge laws, $B_1, \dots, B_{n(b)}$, which equate [ontologically] economical propositions to propositions concerning the fictional items within the prodigal ontology ... and which, together with $E_1, \dots, E_{n(e)}$ entail the [ontologically] prodigal propositions, $P_1, \dots, P_{n(p)}$. We have also some propositions, $A_1, \dots, A_{n(a)}$, which derive from a presupposed auxiliary theory A concerning the fictional items. From these and $P_1, \dots, P_{n(p)}$, we deduce our [ontologically] prodigal conclusion, P_c . Finally, we have a bridge law, B_r , which allows us to deduce an [ontologically] economical conclusion E_c .²⁷

In the case of mathematical claims about possibly uninstantiated mathematical structures, derivative structures are the fictional items that constitute the prodigal ontology, the basic structures from which they were derived from the economical ontology, and the bridge laws are the operations that we used in order to come up with the derivative structures, namely abstraction as explicated in Chapter 3. This ensures the objectivity of mathematical claims about derivative structures, and allows us to make such claims without making ontological commitments to structures that we do not know to be a part of the physical world.

An important feature of Hinckfuss' mechanism for making claims that refer to a prodigal ontology is that we need to understand the presuppositions that are in place when we make the claims, but the claims need not be translatable into claims that refer solely to our economical ontology. He gives an example²⁸ of Mary who makes the claim "I shall go for a walk." She makes this statement in response to the question: "Suppose it is fine on Saturday. What will you do?" As it turns out it is not fine on Saturday so Mary does not go for a walk, however her claim "I shall go for a walk" does not commit her to a false assertion commitment. Her utterance is governed by

²⁷ Hinckfuss [1993], p. 608.

²⁸ Hinckfuss [1993], p. 600.

the presupposition that it will be fine on Saturday, and if this presupposition turns out not to reflect how the world is, still it does not falsify Mary's statement if properly understood. When a claim is made under a supposition then it must be considered a conditional: in this case "If it is fine on Saturday I shall go for a walk." We do not have to analyse her claim "I shall go for a walk" as a false statement given that it turns out that it rains on Saturday, we just need to understand the context of her locution.

This is also how we can make statements about derivative structures. When I say "all sides of a 28-dimensional cube have the same length" (whatever it may mean to be a "side" in a 28-dimensional geometry) my statement must be understood as the consequent of a conditional claim, whose antecedent is an existence claim about 28-dimensional cubes. My claim is governed by the presupposition of the existence of such items, but in making the claim I do not commit myself to their existence, I do not have to believe in them as a part of my ontology. It may turn out that they exist²⁹ but I can make my claim irrespective of whether or not they actually do. Furthermore, in the case of claims about derivative structures, it is their grounding in concepts that we acquired from our experience of basic mathematical structures that allows our conditional claims about them to be objective. Our idea of a 28-dimensional cube is derived by extending our notions, acquired through experience, of 2- and 3-dimensional geometries and of shapes in them, and given this grounding in the physical world we can make objective conditional claims about their properties. The consequents of these claims will turn out to be true in the event that we find out that, in some way, 28-dimensional cubes do exist.

Having said that derivative mathematical structures that do not exist in the physical world are fictions, how can we say true things about fictitious items if in fact

²⁹ It does not serve the present purpose to debate what it means for a mathematical structure to exist: it may be either in an obvious way like the number 2, or in a more subtle way like the average plumber. Whether or not we count instrumental use as qualifying a structure for inclusion in our ontology is not a crucial concern.

they turn out not to exist without all such claims being vacuously true? Take the example of Sherlock Holmes: he is a fictitious character who is composed out of concepts that we have from our experiences of reality, such as a human being, male, a detective, preferring a certain type of hat, and so on. Given that he is composed out of concepts that we are already acquainted with, we can infer certain truths about him. Since he is a human being, he has (or at one time had) a mother and a father, he does not have three heads and he is a carbon-based organism. These are all truths that we can infer from the supposition that Sherlock Holmes exists. There are also truths that we cannot determine from this supposition, such as whether he had a second cousin living in Bulgaria. This is due to the incompleteness of our reference to Sherlock Holmes, which is the same as the incompleteness of our reference to mathematical structures.³⁰

We need to be able to distinguish between the two statements “Sherlock Holmes had a mother” and “Sherlock Holmes’ mother was a flamingo,” because although Sherlock Holmes did not exist we want to be able to say that, in some sense, the first claim is true and the second is false. Hinckfuss’ mechanism allows us to do this, since if we understand that they are being uttered under the presupposition that Sherlock Holmes existed, we can regard them as hypothetical claims with the existence of Sherlock Holmes as the antecedent and various claims about his mother as the consequent. The semantics of these claims are the same (for our purposes) as those of claims made about derivative structures that seem not to exist, such as infinitely large collections or 28-dimensional cubes. However it should be noted that Sherlock Holmes is far more fictitious than most mathematical structures. Generally we are led to derivative structures through a desire to understand some phenomenon, and it is often the case that we find them in the physical world in some sense.

³⁰ Later chapters discuss incompleteness of reference in more detail.

While most derivative structures seem to arise fairly naturally because of their connection with the physical world, this is not the case with all of them, especially once we get up into the more abstract derivative structures. Sometimes, in order to keep progressing with a problem, we have to make choices that determine what kinds of derivative structures we end up with, and we are not always well-informed about whether or not we are making the right decision. It is important to note here that we do tend to think that there *is* some ‘right’ decision, or possibly several right decisions, even if we do not know what they are. These situations in which we are forced to make choices that determine our derivative structures account for the existence of contradictory claims within mathematics and logic. When we make these choices, the truths about the resulting mathematical structure are contingent on our having made the correct choice (if in fact there is a correct choice), and we are less certain of these truths the less certain we are of our decision.³¹

A good example of this type of phenomenon comes from whether or not to accept the Axiom of Choice.³² This axiom and its negation are both consistent with the other axioms of Zermelo-Fraenkel set theory (ZF), so two branches of set theory have arisen. One includes the axiom of choice (AC) among its axioms, and the other includes its negation. It seems reasonable to use either of these systems, since both AC and \neg AC are consistent with ZF, however they are obviously incompatible with each other and thus the results derived using one contradict some results derived using the other. The problem is that both of these systems are potentially correct, and we have no indication of why we should use one rather than the other,³³ however we

³¹ The high degree of certainty with which we hold most mathematical beliefs shows that at the lower levels we have not made difficult choices of this type; the derivative structures which we have are quite obvious and natural.

³² The Axiom of Choice states that for any collection of nonempty sets, we can choose one element from each set.

³³ These days the majority of mathematicians tend to agree that AC should be accepted, rather than its negation, but there is no conclusive consensus.

obviously cannot use both since they yield contradictory results. This example raises some interesting issues to do with mathematical truth, certainty and contingency.

When we make a decision such as to accept AC, all the mathematical truths that we derive from this point forward will be contingent on AC being true. A later chapter examines certainty in mathematics in greater detail, but it is a mistake to assume that we hold all mathematical beliefs with absolute certainty. Truths about basic structures are held with a very high degree of certainty, because these are usually verified and continually confirmed by experience (which is why they often seem self-evident), and can be thought of as generalisations about a range of physical systems. When we move to more highly abstract derivative structures we lose this direct link with the world, however as many derivative structures are natural extensions of the basic ones we observe in the world, truths about these structures retain a high degree of objectivity. The only choices we have made are fairly uncontroversial and intuitive, for example to extend the natural numbers to include all integers, which is hardly a controversial choice. This means that we maintain a high degree of certainty in truths about this derivative structure. The less controversial the decisions we make in coming up with a derivative structure, the more certain we are of truths about it. Mathematical truths are contingent on the assumptions we have made in arriving at the structures that are their subject matter. For the majority of our mathematical knowledge, we have not made any controversial decisions in coming up with the derivative structures involved. Most of them are natural extensions of basic structures, which is why mathematics is considered so objective. AC is a striking example, because the dilemma it poses shows that usually we don't have such a dilemma. It demonstrates that the objectivity and comparative certainty³⁴ with which we hold mathematical knowledge comes from the fact that even the derivative

³⁴ Compared to other (more fallible) kinds of knowledge.

structures which are not found in the world *per se*, are somehow embedded in it in the form of natural extensions to the basic structures to which we have direct perceptual access.

We are given basic structures as a part of our experience of the world, and many of the concepts that are associated with these structures lead us to consider derivative structures that are not immediately available as a part of our experience of the world, but are still defined in terms of concepts that come from the physical world. Sometimes we are faced with decisions that lead us to different derivative structures, but it is not clear which is the correct decision; in other words, it is unclear which derivative structure is the correct one, or the one that will be the most useful in describing the physical world, and in fact there may not be just one. In this situation, such as with AC, we must make a decision one way or another in order to make some sort of progress, but the resulting derivative structure is not of the same kind as the ones which are forced upon us. The truths of statements about such a structure are contingent on the decisions or choices we made in arriving at the structure in the first place. For example a result which has been proved in ZFC (Zermelo-Fraenkel set theory plus AC) depends upon the truth of AC, which is undecidable. In Chapter 8 we shall see that certainty, even about mathematical propositions, comes in degrees. If we prove a proposition using AC, then we are only as certain of it as we are of AC (unless we have some additional justification for it which could increase our certainty).

In the following chapter I return to the question of the nature of mathematical knowledge. It is argued that Skolem's paradox supports the view that formal systems are unable to capture the intricacies of natural phenomena, and mathematical phenomena are no exception.

CHAPTER 5: SKOLEM'S PARADOX

Cantor's diagonal argument proves that the set of all sets of integers is uncountable. The Skolem-Löwenheim theorem proves that for any first order theory there will always exist an enumerable model. How can these two results be reconciled? Putnam¹ has argued that the tension between these two results refutes what he calls 'moderate realism'. Either we are forced to accept Platonism in order to explain our knowledge of non-denumerable sets, or we must abandon classical truth theory. This is a version of Benacerraf's famous dilemma,² which requires philosophers of mathematics to choose between Platonism and a standard theory of truth and reference. Epistemic structuralism is one form of moderate realism, so this chapter argues against Putnam's interpretation of the implications of Skolem's paradox. I will defend a version of moderate realism, one in which a variety of mathematical models may all be said to correspond to mathematical reality. This is suggestive of a wider phenomenon, namely the inability of formal systems to capture aspects of reality in their entirety.

5.1 Benacerraf's Dilemma

Benacerraf, in his famous paper "Mathematical Truth,"³ brings out one of the most important tensions in the philosophy of mathematics. He argues that any account of mathematical truth explains *either* the epistemology *or* the semantics of mathematics, each at the expense of the other, and that "we lack any account that satisfactorily brings the two together."⁴ If a standard theory of truth and reference is adopted for mathematical truths, this traditionally forces a Platonic theory of

¹ Putnam [1983].

² Benacerraf [1973].

³ Benacerraf [1973].

⁴ Benacerraf [1973], p. 663.

mathematical entities. Standard semantics entail that mathematical terms have referents, and these are most often supposed to be Platonically abstract mathematical objects. Benacerraf favours a causal theory of knowledge, and this seems to make mathematical knowledge impossible if mathematical entities are Platonic, and hence outside our causal realm. This dilemma is a significant one for philosophers of mathematics, and can be resolved in one of two ways. The first method is to stick with Platonism, and to give an account of mathematical intuition that explains the existence and legitimacy of mathematical knowledge. This is the route taken by Gödel⁵ and followed by Maddy,⁶ who attempted to explain mathematical intuition in terms analogous to sense perception.

Maddy points out that we only ever perceive an aspect of any perceptual object, we never experience any object in its entirety. From this she argues that we can perceive sets, although the only aspect of them we experience is their members, and she goes on to explain how a perceptual faculty of mathematical intuition could work. Her argument relies on the analogy between our experiences of physical objects and our experiences of sets, arising from the fact that in each case we only experience aspects of the entity in question. However there is a significant difference between our perceptions of aspects of a physical object giving us an impression of the object itself, and our perceptions of members of a set giving us an impression of the set itself. The difference is that in the case of perceiving sets the aspects that we perceive, namely the members of the set, are things in themselves that are independent of the set. The red book, as a member of the set of things in this room, constitutes a distinct object independently of the other things in the room; by contrast, the redness of the book is not independent of the shape of the book, and neither constitutes a complete object when considered independently of the book itself. This suggests that the book is an

⁵ Gödel [1947].

⁶ Maddy [1980].

actual thing in the world, aspects of which hang together in our perception, since they are all aspects of the same thing. However there is no similar connection between the members of a set, they only hang together insofar as we consider them to do so. Sets are derivative structures that we are acquainted with via a process of abstraction as described in earlier chapters, but we cannot perceive mathematical entities if they are supposed to be Platonic.

The other way to resolve Benacerraf's dilemma is to start with an epistemology of mathematics and use it to determine the nature of mathematical entities. This is the route that I have taken in formulating epistemic structuralism, and detailed accounts of mathematical entities and our knowledge of them have been given in other chapters. This is similar to the tactic chosen by Putnam,⁷ however rather than locating mathematical entities in the material world, as I have done, he favours a more nominalistic account, based on the notion of proof. I will evaluate Putnam's argument against moderate realism in light of Skolem's paradox. He considers that there exist three main positions on reference and truth: extreme Platonism on the one hand, that has the problem of reference; verificationism on the other hand, that replaces the classical notion of truth with an account in terms of verification or proof; and moderate realism, which comes somewhere in between, as it attempts to preserve classical notions of truth and reference without positing unnatural mental powers, as the Platonist is forced to do. The position of epistemic structuralism falls into the moderate realist camp, but Putnam argues that moderate realism is threatened by Skolem's paradox. Thus I will first outline the Skolem-Löwenheim theorem and the resulting paradox, and then consider Putnam's argument.

⁷ Putnam [1983].

5.2 The Skolem-Löwenheim Theorem and Skolem's Paradox

The Skolem-Löwenheim theorem suggests that there is a relativity of set theoretic results, since it follows from the theorem that if a theory of first-order logic has an intended non-denumerable model, it also has a model with a denumerable domain. The theorem states that if a set of sentences (a theory) has a model, then it has an enumerable model. From this theorem arises Skolem's paradox, which is that a sentence that says there exist non-denumerably many sets of natural numbers can be true, even though the domain of the interpretation contains only denumerably many sets of natural numbers. In other words, we can prove a sentence that asserts the existence of uncountably many objects in a model with only a countable domain.

The reason this seemingly contradictory situation can come about is that there is a function called an enumerator which enumerates the domain, but this function is not in the model. The intended non-denumerable model contains a true proposition that states "there exist uncountably many objects," because in that model there is no enumerator function that can count the elements of the domain. In order for the domain of a model to be enumerable there must exist within the model a function that establishes a bijection between the elements of the domain and the natural numbers, called an enumerator function. In the intended non-denumerable model there is no such function. However an enumerator function for the domain of this model *does* exist. It does not exist in that model, but in a different model. Thus from the point of view of another model, one which contains an enumerator function for the domain of the intended non-denumerable model, the domain is denumerable.

Therefore there cannot be a first order theory using a classical finitary language that has a genuinely uncountable model, since by the Skolem-Löwenheim theorem any theory with a non-denumerable model will also have a model that is denumerable. If a theory has an intended non-denumerable model, this model cannot be captured by set theory, since there can be no 'genuinely' uncountable models because every model is countable relative to some other model. This result suggests

two very significant corollaries for set theory: first, that set theoretical results are relative rather than absolute and second, that no axiomatic system can capture our intuitive concept of set theory. One way to defuse the severity of this result is to reject the notion that a classical finitary language is adequate to provide an axiomatisation of set theory. This was the method used by Zermelo, who insisted on the infinitary nature of mathematics, in which case a finitary language will necessarily be inadequate to capture our intuitive understanding of set theory. His objection is discussed later in this chapter.

Another potential way to evade the problem of the inability of axiomatic systems to capture our intuitive notion of set theory is to deny the existence of uncountable sets. Indeed, Jané⁸ claims that Skolem's argument was that there is no good evidence for the existence of uncountable sets. He correctly notes that Cantor's diagonal argument proves only that the set of all sets of integers is uncountable, it does not provide any evidence for the existence of this set. However even if this set does not exist, the Skolem-Löwenheim theorem still implies the relativity of set-theoretical notions. If the intended model of ZFC has a countable domain, there is still a theorem of ZFC which states that there are uncountable sets, since we can take the power set of the set of all integers. So even if our intuitive belief in the existence of uncountable sets is wrong, the relativity of set-theoretical notions remains a corollary of the Skolem-Löwenheim theorem. Thus the denial of the existence of uncountable sets fails to evade Skolem's paradox.

⁸ Jané [2001].

5.3 Putnam's Argument Against Moderate Realism

In the first chapter of his book *Realism and Reason*,⁹ Putnam argues that moderate realism – namely the view that tries to maintain classical notions of truth and reference without postulating unnatural mental powers – is the position most seriously threatened by Skolem's paradox. Epistemic structuralism assumes a form of moderate realism, so his concerns must be addressed. His claim is that Skolem's paradox forces a trade-off: we must either postulate an unexplained mysterious faculty of mathematical intuition (as Platonists must do) or we are forced to abandon classical truth theory.

Putnam argues for the relativity of the truth value of certain mathematical statements, and claims that "Skolem's argument ... casts doubt on the view that these statements have a truth value independent of the theory in which they are embedded."¹⁰ This means that our intuitive notion of set cannot be captured by axiomatic set theory, and Putnam shows that since *something* must capture our intuitive notion of set (or else what is axiomatic set theory failing to capture?) we are led once more into Benacerraf's dilemma. The Skolem result can be used to argue for Platonism, since we must acquire our intuitive understanding of set theory somehow, and if it cannot be formalised by axiomatic set theory then we must have some mysterious faculty by which we acquire this notion. As Putnam points out, this option will be rejected by "the naturalistically minded philosopher," who would never accept the notion that our mathematical knowledge originates from such an occult faculty. In earlier chapters we have seen some of the difficulties that arise from holding Platonism with respect to mathematical entities.

⁹ Putnam [1983].

¹⁰ Putnam [1983], p. 10.

Putnam's preferred solution to the problem is instead to reject classical causal theories of truth and reference in favour of verificationism. This view analyses truth and reference in terms of verification and proof, rather than by truth conditions or correspondence with reality. On this view, Putnam explains, our understanding of the statement that a certain set (such as the set of the real numbers) is non-denumerable is no longer attributed to a mysterious faculty, but consists of our knowing what it is to prove that the set is non-denumerable. The relativity of set theoretic results is no longer a problem under this analysis, because our understanding of notions like non-denumerability of sets is based on "an evolving network of verification procedures."¹¹ The claim that the real numbers are non-denumerable is made true by our understanding how to prove that it is true. Putnam analyses our understanding of language in terms of how we use it, so reference is linked to use rather than to the world. This approach does dispel the problem of Skolem's paradox, however it opens up the difficulty of how such a non-realist semantics can explain the independence of mathematical claims from their knowers. Thus a preferable approach is to attempt to refute Putnam's claim that Skolem's paradox is fatal for moderate realism, thereby retaining a standard semantics for mathematical knowledge while avoiding the pitfalls of Platonism.

Putnam presents an argument for the relativity of " $V = L$," the claim that all sets are constructible.¹² Gödel's intuition, which Putnam notes is shared by many other set theorists, was that if set theory is consistent then $V = L$ is false, even though it is consistent with set theory. In trying to determine whether this intuition makes sense, Putnam proves the following theorem:

Theorem: ZF plus $V = L$ has an ω -model which contains any given countable set of real numbers.

¹¹ Putnam [1983], p. 22.

¹² L is defined as the class of all constructible sets, and V is the universe of all sets.

An ω -model for a set theory is defined as “a model in which the *natural numbers* are ordered as they are ‘supposed to be’; i.e., the sequence of ‘natural numbers’ of the model is an ω -sequence.”¹³ Using the above theorem Putnam explains the construction of a model for the entire language of science in which $V=L$ is true. He describes how to construct a formalisation of the entire language of science that satisfies $V=L$. Any model that contains a set isomorphic to OP (which is the correct assignment of values, with operational constraints coded in) can be extended to a model for this formalised language that is standard with respect to the assignment of values to physical magnitudes. If Gödel’s intuition is correct, the model described must not be the intended model, even though it satisfies all required theoretical and operational constraints. The only way that $V=L$ could be false is if we add $V \neq L$ to the axioms of ZF as a theoretical constraint. This means that the truth value of $V=L$ depends on which theoretical constraints we adopt for the model in question, in other words the truth value of $V=L$ will vary between different intended models. There is no objective way of deciding whether $V=L$ is true or not, it depends only on whether or not we decide to adopt it as one of the theoretical constraints for a model.

There are similar arguments for the relativity of both the axiom of choice (AC) and the continuum hypothesis (CH). We can find models for the entire language of science that satisfy AC or CH. Whether or not these are the intended model depends on whether or not the falsity of AC or CH is coded in to ZF as a theoretical constraint.

Putnam takes these arguments to show that (given our classical analysis of truth) realism must be false, since “the realist standpoint is that there is *a fact of the matter* ... as to whether $V=L$ or not.”¹⁴ He does not consider another plausible possibility: that on a realist conception these statements ($V=L$, AC and CH) are neither absolutely true nor absolutely false, and that their truth value depends in part

¹³ Putnam [1983], p. 6.

¹⁴ Putnam [1983], p. 7.

on the theory in which they are embedded. They could be true in some models and false in others, without contradiction. There may be no objective fact of the matter as to whether they are true or false, if the reality of the world is not complex enough to make such claims decidable. Putnam points out that it makes no sense to decide arbitrarily that AC is true, or even to decide that we will take it to be true, for example because it has so many useful implications. Its truth would be only by convention, and we could not say that a theory which includes its negation is false, merely that it contradicts our convention. However for some mathematical claims, we do not know whether they are true or false in an absolute or fixed sense, but we can use models in which they are defined as true or false. This certainly happens in mathematical practice, for instance theorems have been proved that are contingent on the truth of the Riemann hypothesis, even though that result has not been proved. We can work within models in which an undecided mathematical proposition is assumed to be true or false, and until the assumed proposition is proved or disproved we do not know which of these models reflect mathematical reality.

5.4 Truth Value of Mathematical Claims

According to epistemic structuralism, we have access to mathematical reality in the same way that we have access to other aspects of the physical world. We observe basic mathematical structures, abstract from their physical instantiations, and using the concepts that this process provides us we can derive further mathematical truths. Some of these turn out to be true of the physical world, some have value as instruments to gain greater insight into mathematical reality, and others are hypothetical or fictional truths that may or may not turn out to refer to the world. These claims are about derivative structures, and because of our limited access to mathematical reality we sometimes do not know whether the assumptions we made in coming up with them correspond to mathematical reality or not. We can postulate various mathematical systems that all describe mathematical reality as far as we

know, but in some cases it is difficult to tell which is the right one, even if they contradict each other. Thus we can come up with various mathematical models that are all candidates for being the privileged model, the one that reflects mathematical reality.

This is how there can be mathematical statements that lack any absolute or fixed truth value, at least from our perspective. Statements whose truth value is unknown (possibly unknowable) have no truth value for our purposes. In order to work with these statements, we need to make assumptions as to whether or not they are true, bearing in mind that we do not know which is the case. Since we must work within the limits of our theories, we can treat such statements as lacking truth values. The axiom of choice (AC) is an example of such a statement. We can postulate a model in which AC is true, and another in which AC is false, and if each of these could be the privileged model, AC has no absolute truth value. It is a claim about a derivative structure, one that we have not found in our experience of the world, although it is possible that one day we may. Since there are models in which AC is true and models in which AC is false, none of which have been ruled out as candidates for being the model that accurately reflects mathematical reality, we cannot say whether *in reality* AC is true or false, or even if it has a truth value. Different models that make differing claims about its truth value can be useful for different purposes, so in some circumstances it may make sense to claim that AC is true or false, in order to derive truths in a particular model. However we cannot claim that AC has a fixed truth value across all models, or know whether it has a fixed truth value in reality.

Mathematical statements that lack truth values in this way are generally about derivative structures. Statements with no truth value are to be expected when we deal with derivative structures, since they no longer correspond directly to actual structures, but rather are composed of concepts that we have acquired in the course of our experience with basic mathematical structures. Claims about derivative structures

may not correspond to facts in the world, so it is not surprising that they lack truth value. As long as there exist mathematical statements that are undecidable in this way, any formal system that we use to capture mathematical reality will fall short in some way. Gödel's incompleteness results suggest that we will never be able to find *the* model that captures mathematical reality, however we can develop extremely rich models that are good candidates for being the privileged model.

5.5 Zermelo's Refutation of Skolem's Paradox

Zermelo was one member of the mathematical community who believed he had a solution to Skolem's paradox. Van Dalen and Ebbinghaus¹⁵ provide an analysis of Zermelo's refutation of Skolem's paradox, claiming that it fails to refute the paradox but that it provides a revealing insight into Zermelo's epistemological convictions. Their discussion is useful in considering the implications of Skolem's paradox.

In essence Zermelo's refutation of the paradox centres around his infinitary convictions about the nature of mathematics. He concludes that Skolem's paradox is based on a "finitistic prejudice," namely the assumption that "every mathematically definable notion should be expressible by a finite combination of signs."¹⁶ Van Dalen and Ebbinghaus reveal Zermelo's views on the nature of mathematics, namely that it is infinitary in nature and can only be apprehended a priori, in a Platonic sense. He considered mathematics to be "the logic of the infinite,"¹⁷ and thus believed that a first-order approach would fail to capture the richness of mathematics. For him, using finitary combinations of symbols is just the way that our inadequate intelligence tries to approach what he considers to be 'true' mathematics, which is "the conceptual and

¹⁵ van Dalen and Ebbinghaus [2000].

¹⁶ van Dalen and Ebbinghaus [2000], p. 145.

¹⁷ van Dalen and Ebbinghaus [2000], p. 150.

ideal relations between the elements of infinite varieties.”¹⁸ Thus Zermelo sets himself the goal of constructing an infinitary system of logic. In order to elude the finitary prejudice he drops any restriction to the axiom of separation, since he believes mathematics to be about the infinite, so there should be no limitations on how subsets can be formed.

Zermelo's refutation of Skolem's paradox follows from his attempt at constructing his infinitary logic. He sets out to construct a system of logic in which Skolem's paradox cannot arise. Ultimately he failed to achieve this aim. In 1937 he wrote what he considered to be a refutation of Skolem's theorem, but as van Dalen and Ebbinghaus point out, his argument “amounts to a proof of the set-theoretical statement saying that, given a denumerable set M , any subset of the powerset K of M that is closed under arbitrary unions and intersections and whose union is M , either is finite or of the same cardinality as K .”¹⁹ Thus while it may seem that Zermelo missed the point of Skolem's theorem, van Dalen and Ebbinghaus suggest a more likely theory: Zermelo considered that a first-order approach to set theory was doomed to failure because of the infinitary nature of mathematics. He was absolutely convinced of this, and as a firm Platonist he strongly believed that there exists a unique set-theoretical universe. He considered that the only way to resolve Skolem's paradox was to abandon its finitary assumption and move to an infinitary logic.

For Zermelo, Skolem's paradox reinforced his infinitary convictions about the true nature of mathematics. He believed that the paradox rests on the assumption that all of mathematics is expressible using a finite combination of symbols, an assumption that he refers to as the “finitistic prejudice.” The fact that this assumption leads to a paradox confirms Zermelo's infinitary convictions, since from a falsehood any absurdity can be derived. He considers that the contradiction apparent in

¹⁸ van Dalen and Ebbinghaus [2000], p. 153.

¹⁹ van Dalen and Ebbinghaus [2000], p. 156.

Skolem's paradox confirms the erroneous nature of the finitistic prejudice, and he attempts to refute it by coming up with an infinitary logic. As van Dalen and Ebbinghaus explain, Zermelo's failure to refute Skolem is due to his firm belief that true mathematics is infinitary in nature. This belief is based on intuition more than reason. Modern physics provides evidence to suggest that the physical world is not infinite, nor that the infinitesimal has any physical manifestations. The previous chapter argued that we have evidence only for potential, rather than actual infinity. Thus Zermelo's refutation of Skolem, which rests on his infinitary conviction, is not sound.

5.6 Implications of Skolem's Paradox: Inadequacy of Formal Systems

Zermelo's refutation of Skolem's paradox contrasts with Skolem's own resolution of the situation. The paradox faced is that any theory can be represented in a denumerable model, so we can have a theory that contains the true proposition P "there exist uncountably many objects," yet this theory has a countable model. The way this situation comes about is that the non-denumerable domain has an enumerator function, a function that maps the domain onto the natural numbers (effectively, a function which 'counts' the domain, or establishes a bijection between the elements of the domain and the natural numbers). The way that such a function can exist without refuting P is that the enumerating function is not in the model in question. The model that contains P is an uncountable model in the sense that there is no function in the model that can count the elements in its domain. However there is a way of enumerating its domain even though the function required to do so is not a part of the model. This is how P can be provable, even though there exists a different model containing the enumerator function, namely a countable model for the theory.

This resolution to the paradox has some interesting philosophical implications, since it suggests both the impossibility of a genuinely uncountable theory as well as the relativity of set theoretical results. We have seen that these results lead to the idea

that we can postulate various mathematical systems that all describe mathematical reality as far as we know, but we have no way of knowing which is the right one, even if they contradict each other. This is how there can be undecidable mathematical statements, and any formal system that we use to capture mathematical reality will fall short in some way.

On this view, our attempts at capturing mathematical reality are similar to our attempts at capturing any other aspect of the world in formal systems, none of which are completely adequate. An example of this situation is the failure of physicists (so far) to find a unified 'Theory of Everything' which encompasses both the large-scale phenomena explained by relativity theory and astronomy as well as the smaller-scale events described by quantum mechanics. These shortcomings parallel the relativity of set theoretical results, as revealed by Skolem's paradox. The paradox tells us that set theoretical results do not all have an absolute truth value, some will have different truth values that vary relative to the interpretation under consideration. This is analogous to the lack of a unified 'Theory of Everything' in modern physics: Einstein's laws of general relativity hold in most circumstances, but when the dimensions involved are extremely small these laws break down. There is the chance that a unifying physical theory may one day be found, but currently it is so far off that a comparison may be made with the Skolem-Löwenheim result. Just as no formal axiomatisation of set theory can ever capture our intuitive notion of a non-denumerable set, there is no theory of physics that can explain all physical systems (the difference being that the latter has not been proven to be unattainable).

Another example of the failure of formal systems to capture aspects of reality concerns our use of natural languages. It has been argued that formal languages are not powerful enough to capture our ordinary use of language. Haack,²⁰ in a discussion

²⁰ Haack [1978], p. 73.

of singular terms and the denotation of names, points out that it has been argued (for example by Schiller and Strawson) that there are subtleties in natural language that are beyond the scope of formal languages. She explains that often the pragmatic aspects of discourse simply cannot be captured by formal logic. It seems reasonable to expect that there would be a similar difficulty with mathematics and set theory, namely that there are some things that an axiomatic system cannot tell us about a mathematical structure, and this is what Skolem's paradox shows. The fact that a formalised set theory cannot capture our intuitive understanding of set is a perfect example of the inability of formal systems to capture all aspects of any given conceptual entity.

Schiller and Strawson's claim that formal methods are inadequate for the subtleties of natural language is consistent with the incompleteness of arithmetic. If a formal system such as arithmetic falls short of being able to capture everything about mathematical reality, then it makes sense that it will not capture every true statement of mathematics. This is analogous to the inability of any one scientific theory to explain the world completely. Another manifestation of this principle is in first-order logic, which is undecidable. Reasoning is something that we do as a part of our interaction with the world, and first-order logic is an attempt to capture how we reason. That this formal system is unable to classify every statement as either true or false (undecidability) is to be expected if we accept that formal systems will always fall short of capturing some aspect of the world completely. In science, mathematics, logic and linguistics we see this principle confirmed, that formal systems are insufficient to capture every aspect of natural phenomena.

Putnam gives the following analysis of the philosophical problems arising from Skolem's paradox: we have explained our understanding of language in terms of how to use it, and then we have tried to find out what models 'out there' we can find for the language. He points out that something strange is happening here, because our understanding of the language is supposed to determine reference, and yet the language lacks interpretation. Perhaps this is the wrong way of approaching the

matter. Rather than taking language as our starting point and analysing how we understand language, we need to consider what is 'out there' in the world as *at least as fundamental* as our concepts. When we observe the world (which is where most of the models are, at least for basic mathematical structures) we need concepts, expressed in language, in order to make sense of the structures that we observe, as well as logic in order to understand our observations. Just as we cannot have an isolated observation which is independent of concepts, it is impossible to understand a language meaningfully if it is considered in isolation from our experience of the world. Language is a tool for us to try to explain and pin down all the features of the world, and the fact that it comes up short and cannot completely capture various systems in the world is just how it is. Language is not as powerful as we might wish, but this is not such a serious problem.

As well as his infinitary conception of mathematical reality, Zermelo holds a view that hints at this feature of formal systems. For Zermelo, mathematics is our inadequate way of trying to capture mathematical reality. Since the language of mathematics is finite and he believes that mathematical reality is actually infinitary, there are bound to be inadequacies. He was convinced that the nature of mathematics was infinitary, and that our formalisations have not yet come that close to capturing the reality of mathematics, since we are limited by the symbols we use. His views have brought out the point that I wish to emphasise about formal systems: that they almost always fall short when attempting to capture some aspect of reality. He believed that the reason a formal axiomatisation of mathematics fails to capture mathematical reality is that the axiomatisation uses only a finite number of symbols, and is therefore unable to reflect the infinitary nature of mathematical reality. Skolem's paradox also tells us that formal systems cannot capture mathematical reality, but for a different reason: that set theoretical results are relative rather than absolute. Both these cases suggest that our attempts at explaining mathematical reality using formal systems has fallen short, however that does not imply that we have not

come close. Using analogies with language, we can see that we can still get close to explaining mathematical reality, even if we cannot capture it in its entirety.

Language may be unable to capture completely things that we perceive about the world, however this is not fatal for its usefulness. As we saw previously, reference is rarely if ever complete. I can successfully pick out people by using their names or by pointing to them or by describing them for example as “that tall blonde girl wearing a blue T-shirt.” My reference can succeed without capturing everything about the person. I can use language successfully to refer to this person, but my reference will not tell me whether or not the person has a sister, and indeed my concept of that person might not include this information. The incompleteness of the reference does not prevent it from picking out the right person, nor from making true statements about her. Similarly, in set theory the language is incomplete, it does not capture everything about the structure in question. However this is not a fatal problem, we can still express true statements. Skolem's paradox and Gödel's incompleteness results do not require us to abandon classical notions of truth and reference. Although we cannot capture set theory with a formal system, we can still refer successfully to set theoretical entities and say *some* (indeed, many) true things about sets. The true claims that we can make are true in virtue of reflecting (true) facts about the set theoretical structures which we have come up with by extension from perceived structures in the world. This is why they can be considered to be objective truths.

Once we start doing things like accepting AC as true, we become less certain of the set theoretical results derived. Results proven using AC are contingent on its being true, and we can only be as certain of their truth as we are of AC. Our mathematical methods now become even more like those of empirical science, since we make assumptions (for example that AC is true) and continue working, deriving results that are contingent on the assumptions we make. The more evidence we have for AC the more certain we are of results that follow from it. In addition, the more intuitive (or ‘correct-seeming’) results we derive using AC the more convinced we are

of its truth, even though we cannot be completely certain. It may actually be the case that AC has no objective truth value, that there is nothing about the world that makes it true or false.

This outcome is not so surprising, since the notion of a set of sets is part of a derivative structure that is not directly observable in the world. It is debatable whether sets themselves are basic structures or derivative structures. All that we perceive directly are the members of the sets, not the sets themselves. It was argued earlier, based on Maddy's account of mathematical intuition,²¹ that we cannot perceive sets directly, which implies that they must be derivative structures (under an epistemic interpretation of what it is to be a derivative structure). Given that sets themselves are derivative structures, then it is evident that sets of sets are also derivative since the notion of a set of sets is derived from a derivative structure. It is not surprising that there are statements about such a structure that have no objective truth value, since the structure is not (as far as we know) instantiated in objectively existing systems in the world, but is a concept that we have come up with to help us understand the world and the structures in it. Nor is the conclusion that AC (as well as some other set theoretical claims) has no objective truth value devastating to mathematics, since we can continue to do mathematics (albeit producing fallible theorems) without knowing whether AC is true or false or has no truth value.

This view fits with the incompleteness of reference and language, and this is simply a fact about our interaction with the world. Indeed it is a desirable outcome, since if a formal system is factually complete it is untestable. If a formal theory is complete then it will be a maximally consistent set. In order to test a theory we have to be able to add premises to it and see whether the truths derived reflect reality. If the

²¹ Maddy [1980], see discussion earlier in this chapter.

theory is already complete then there are no true premises left that we could add, so we would have no way of testing the theory.

This conclusion also implies that mathematical knowledge, rather than being held as a privileged realm of certainty, can be regarded as a species of scientific knowledge and understood in the same way as the rest of our knowledge about the world. It is a body of knowledge that attempts to capture an aspect of physical reality. Mathematics is concerned with studying the mathematical structures that occur in our world. In the following chapter I argue that mathematical methods are not significantly different from scientific methods.

CHAPTER 6: MATHEMATICS AS SCIENCE

My aim in this chapter is to show that there are significant similarities between mathematics and other physical sciences, in terms of both content and methodology. Both are studies of physical systems and consist of attempts to understand the workings of some aspect of the world. The most significant difference between mathematics and the physical sciences is that the systems or structures studied in mathematics are much more general (or metaphysically simpler) than those studied in other sciences, and thus they apply to a much wider range of situations in the physical world. This high degree of generality has led to the misconception that mathematics is about abstract objects, rather than about underlying structures in the world. By showing that the way we acquire mathematical knowledge is not very far removed from how we acquire scientific knowledge, I hope to lend further support to my epistemic structuralist claim that the subject matter of mathematics consists of structures, many of which are instantiated in physical systems.

In order to determine the similarities and differences between mathematical and scientific knowledge I will examine the methods used by mathematicians in generating mathematical knowledge. I hope to dispel the popular belief that mathematical knowledge is produced by deductive proofs from axioms that are known with certainty. I favour a view that mathematics is quasi-empirical (to use the term of Lakatos and Putnam, who favoured this view as do others such as Tymoczko). This makes mathematics much like the rest of science in many respects. By examining the role of proof I intend to show that its perceived epistemic function as the basis for generating mathematical knowledge is not in fact the case, but rather that it is only one element in a more scientific, quasi-empirical process. Then I shall make use of examples from the formal sciences as given by Franklin, which serve to bridge the gap between mathematical and scientific knowledge and methods, and discuss considerations arising from the use of computer proofs in mathematics.

6.1 Mathematics as Quasi-Empirical

The axiomatic-deductive method is traditionally used to explain mathematical statements in such a way as to make their truth apparent. This practice has led to the mistaken conception that this is also the method by which mathematical truths are discovered. As Lakatos explains, “[c]lassical epistemology has for two thousand years modelled its ideal of a theory ... on its conception of Euclidean geometry.”¹ The model of Euclidean geometry, or the axiomatic-deductive method, is the process of using rules of deduction to prove consequences of certain initial axioms. The Euclidean method does not specify where the axioms and rules of deduction should come from, although the positivists took the axioms to be necessary truths. This method preserves certainty: the operations of deductive logic are truth-preserving, certainty ‘flows downward’ through this model from the axioms to the mathematical results which are their consequences. Since axioms are often taken to be necessary truths or self-evident truths, they are considered to be certain, and the axiomatic-deductive method preserves this certainty.

In contrast to this method is what Lakatos² and Putnam³ have termed ‘quasi-empiricism’, which is a theory that argues against the axiomatic-deductive method as reflective of actual mathematical practice (at least not of *all* mathematical practice), and that this is not the way that mathematical truths are justified. This is the idea that mathematical truths are often acquired by quasi-empirical methods that rely on a process of observation and confirmation, and not solely by the axiomatic-deductive method. Quasi-empirical systems are distinguished from Euclidean or axiomatic-deductive systems. Lakatos⁴ emphasises that a deductive system that is quasi-

¹ Lakatos [1978], p. 33.

² Lakatos [1978].

³ Putnam [1975].

⁴ Lakatos [1978], p. 34.

empirical may be either empirical or non-empirical in the “usual sense,” depending on whether or not the basic axioms are derived from experience. However what sets a quasi-empirical system apart from an axiomatic-deductive one is that propositions may be confirmed or falsified by experience (whether basic propositions or deductions from non-empirical premises). Following Tymoczko⁵ I shall use the term ‘quasi-empiricism’ to refer broadly to the general approach pioneered by Lakatos and Putnam, rather than to any one individual’s particular theory.

Examination of mathematical practice and the role of proof will show that the notion of mathematical knowledge acquisition as purely axiomatic-deductive is a myth, and that there is a significant empirical component. This will lend support to the theory that mathematics is the study of primarily physical structures rather than Platonically abstract entities, by way of an argument that may be known as the Argument from Quasi-Empiricism:

- P1: If quasi-empirical methods are used (at least in part) to generate mathematical knowledge, then the subject matter of mathematics must be (a) objective, and (b) a part of the physical world.
- P2: Quasi-empirical methods are used (at least in part) to generate mathematical knowledge.
- C: The subject matter of mathematics must be objective and a part of the physical world.

This section of this chapter is concerned with establishing the second premise. The first premise is not controversial under the assumption of natural realism: quasi-empirical investigations involve observation and confirmation, and under natural realism these come from interactions with the physical world. In the naturalist metaphysics the physical world is objective, so if we can acquire knowledge of a subject through interaction with the physical world then the subject matter must be objective.

⁵ Tymoczko (ed.) [1998], p. *xvi*.

The axiomatic-deductive method has long been associated with foundationalist philosophies of mathematics, which have dominated until the twentieth century. The traditional view of mathematics was that it consisted of theorems rigorously derived from self-evident axioms. When the antinomies threw mathematical certainty into doubt, this led philosophers of mathematics to re-examine the axioms, as they were considered to be the foundations of mathematical certainty. Attempts were made to ground the axioms of mathematics in either logic, formal systems or the human mind. Quasi-empiricism challenges this foundationalist approach and looks instead to the actual practice of mathematics in its search for an account of mathematical knowledge acquisition. In the introduction to his book *New Directions in the Philosophy of Mathematics*, Tymoczko aims to “challenge the dogma of foundations.”⁶ He explains how the assumption that all mathematical results were derived using the axiomatic-deductive method led philosophers of mathematics to seek to find the foundations of mathematical knowledge. He claims that quasi-empiricism leads to a view of mathematical practice as an interaction between mathematicians and independently existing mathematical structures.

Quasi-empiricism is consistent with a coherentist, rather than a foundationalist story about the acquisition of mathematical knowledge. Instead of building from an axiomatic foundation using rules of deductive inference, mathematical knowledge is acquired by taking various kinds of evidence into account and reasoning about which beliefs to accept or reject, with the aim of building the most coherent web of knowledge. This view is called *quasi-empiricism*, because it does not claim that mathematical knowledge is acquired by purely empirical means. Empirical evidence has an important role in providing us with our mathematical knowledge, however there is also room for other kinds of evidence. When faced with a new theorem, we

⁶ Tymoczko (ed.) [1998], p. xv.

can believe it with a high degree of certainty because it is confirmed by our perceptual data; however this degree of certainty might increase if we are shown a deductive proof for the theorem. So while our mathematical knowledge is grounded in our experience of the world, there is room for other types of verification than empirical evidence to increase our body of mathematical knowledge.

Tymoczko distinguishes between realism and constructivism in the philosophy of mathematics as the respective views that mathematical truths are discovered, or that they are invented.⁷ Realism bears great similarity to Platonism, claiming that mathematical entities are real and abstract, existing in an independent mathematical universe. Constructivism is the view that mathematical truths are inventions of the people who do mathematics. Tymoczko claims that physicalism, the view that all objects exist as a part of the spatio-temporal universe, implies constructivism. This claim belies an unstated conviction that there is no room in the physical world for mathematical entities. The kind of structuralism for which I am arguing would allow for a physicalist realism about basic mathematical entities and hence also about the truths of mathematics: they are about relations between spatio-temporal objects,⁸ which are mathematical structures. I hope to show that this view will sit well with a quasi-empirical account of how mathematicians go about their business.

In his paper “A Renaissance of Empiricism in the Recent Philosophy of Mathematics?”⁹ Lakatos contrasts the Euclidean ideal of mathematical practice with a quasi-empirical account, which is most often associated with scientific knowledge, and claims that mathematics too is quasi-empirical. The Euclidean model is a system of deductive reasoning, in which truth flows downward from certain axioms. In this

⁷ Tymoczko (ed.) [1998], p. *xiv*.

⁸ When we say that the subject of mathematics is structures which are located in physical reality, sceptics may be tempted to ask “where and when is the number 2?.” The answer is that the number two is in every pair of socks, binary stars, pairs of molecules, etc, and exists at every time when there are at least two things in the universe.

⁹ Lakatos [1978].

model mathematical results are discovered by truth-preserving inferences, so there is never an uncertain proposition in mathematics. Lakatos introduces the slogan “foundations and accumulation of eternal truths”¹⁰ for this ideal. By contrast, the slogan for the quasi-empirical method is “growth and permanent revolution.”¹¹ This is the more scientific method of hypothesis testing through observation or experimentation.

Lakatos explains that the way we acquire mathematics is quasi-empirical rather than Euclidean. This view of mathematical knowledge parallels the scientific method of accumulating a body of knowledge. Although we tend to be very certain of most of our mathematical knowledge, mathematical beliefs can be fallible and revisable. Brown explains that “[t]he fallibilism involved is that of having the wrong concept, not the fallibilism of having the wrong beliefs about the right concept.”¹² A mathematical belief can be overthrown, but this involves a conceptual change rather than a straight out negation. When a mathematical belief is disconfirmed by experience this means that we are dealing with a different kind of system than the one to which the belief refers. This shift in meaning is discussed in greater detail later in this chapter.

Although it is really only the notation which changes, this can constitute a very significant shift in our beliefs. This is also what happens when there is a paradigm shift in science: the phenomena we are describing do not change, but we use different concepts to describe what we observe. Newtonian mechanics gave us a theory about our observations which let us describe them and predict future observations. When relativity became the dominant paradigm this was just a different, slightly sharper and more accurate way of describing *the same facts about the world*,

¹⁰ Lakatos [1978], p. 35.

¹¹ Lakatos [1978], p. 35.

¹² Brown [1999], p.23.

so in revising our physical theory we devised a new way of describing the world that gave us better insight into its workings.

This is the same as what happens in mathematics, except that in mathematics we are often working with less specified and more accessible systems than in the physical sciences, so we need less sophisticated structures to describe the facts that we observe. This simplicity means that the resulting theory is less likely to require a conceptual shift, which is why we have not seen as many mathematical revolutions to parallel the scientific revolutions. From this also comes our misguided belief that the truths of mathematics are entirely certain in a pure and privileged sense. The quasi-empirical nature of mathematical knowledge supports the claim that the way we acquire and verify our mathematical beliefs is the same method by which we acquire scientific knowledge. The primary difference between mathematics and other physical sciences is that since the structures we study in science are so much more complicated, the theories are more susceptible to being overthrown than mathematical theories, which deal with much simpler subject matter.

6.2 A Posteriori Confirmation of Mathematical Truths

A common objection to an empirical component in an account of mathematics is that the propositions of mathematics cannot be confirmed or disconfirmed by experience. Casullo advocates an empirical account of mathematical knowledge, arguing that empiricist epistemologies work well because they give a unified account of all our knowledge based on fairly well understood physiological processes.¹³ Apriorists unnecessarily complicate epistemological accounts by postulating a non-experiential source of knowledge which is vague and mysterious; we have very little understanding of how ‘mathematical intuition’ or ‘the mind’s eye’ works in terms of

¹³ Casullo [1988].

our cognitive functioning. Casullo divides empirical accounts of mathematical knowledge into two varieties: inductivist and holistic. The inductivist empiricist maintains that there are at least some propositions of mathematics that can be confirmed or disconfirmed individually, and epistemically basic mathematical propositions can be confirmed by inductive generalisation from experience.¹⁴ This type of view is contrasted with an holistic empiricist position, which focuses on the wider theoretical context of mathematical propositions, maintaining that only an entire theory, rather than individual mathematical propositions themselves, can be confirmed or disconfirmed by experience. Casullo explains that the latter type of theory is a descendant of the former, motivated in part by the desire to avoid objections by apriorists. Since inductivist empiricist theories are more susceptible to the apriorist objections and holistic accounts are descended from them, Casullo focuses on the inductivist empiricist position in his defence against the three main objections to empirical accounts of mathematical knowledge. He claims that if he succeeds in proving that inductivist accounts are not susceptible to these objections, then holistic accounts will also be immune.

The three objections from apriorists against which Casullo defends empirical accounts of mathematical knowledge are: the Irrefutability Argument, the Argument from Certainty and the Argument from Necessity. The Argument from Certainty takes as premises the claims that mathematical propositions are known with certainty, and that inductive generalisations cannot yield certain propositions, and draws as its conclusion that mathematical propositions are not known on the basis of inductive generalisations.¹⁵ Casullo defends inductivist empiricism against this argument by taking issue with the use of ‘certainty’, however I shall not focus on this argument as I

¹⁴ Casullo [1988], p. 43.

¹⁵ Casullo [1988], p. 53.

have my own preferred way of accounting for the certainty of mathematical propositions, which is explained in Chapter 8.

The Argument from Necessity has similar form to the Argument from Certainty: its premises are that mathematical propositions are necessary propositions, and that necessary propositions cannot be known on the basis of experience; hence mathematical propositions are not known on the basis of experience.¹⁶ Casullo acknowledges that the first premise is controversial after Quine rejected necessity and the a priori, however in his defence of inductivist empiricism against this argument he disputes the Kantian assumption that if the general modal status of a proposition can only be known a priori, then the truth value of the proposition can only be known a priori.¹⁷ However as with the Argument from Certainty, I present a different way of accounting for the purported necessity of mathematical propositions in Chapter 8, so I shall focus here on Casullo's defence against the Argument from Irrefutability.

The Irrefutability Argument, originally made against Mill, is an important one for any account of mathematical knowledge that has a physicalist or empiricist component. The Irrefutability Argument as stated by Casullo is:

- (1) No experiential evidence can disconfirm mathematical propositions.
- (2) If experiential evidence cannot disconfirm mathematical propositions, then it cannot confirm such propositions.
- (3) Therefore, experiential evidence cannot confirm mathematical propositions.¹⁸

As Casullo points out, the argument is valid and the second premise is uncontroversial, so we must focus on the first premise. Casullo comes up with a disconfirming scenario to demonstrate that premise (1) is not true. This scenario is

¹⁶ Casullo [1988], pp. 61-62.

¹⁷ Casullo [1988], p. 63.

¹⁸ Casullo [1988], pp. 47-48.

essentially an extension of Mill's, but it incorporates the prerequisite that the participants in the scenario believe that two conditions are met. Wilson refers to these conditions as the Stability condition and the Correct Counting condition and characterises them as follows:

Stability condition: Neither the operations of counting and combining, nor the interactions of the objects to be counted, produce any changes in the relevant features of the objects being counted.

Correct Counting condition: No miscounting (say, by repeating a number or missing an object) takes place.¹⁹

Casullo's scenario involves a participant who is justified in believing that these conditions are met having an experience which provides disconfirming evidence for the proposition $2 + 2 = 4$. Wilson explains that in evaluating this scenario it is useful to consider two aspects of number: the cardinal aspect of number has to do with its being the cardinality of a set, and the arithmetical aspect of numbers has to do with the relations of numbers to each other. In trying to separate these two aspects of number she considers what our counting procedures would be like if they were not integrated with arithmetic. Tallying is a more basic procedure than counting, namely the act of establishing a one-to-one correspondence between objects and tally marks. Wilson calls the set of tally marks "a 'pure' (that is, an 'arithmetic-free') representation of the cardinality of the collection."²⁰ However it turns out that it is impossible to distinguish between tallying two distinct collections and tallying the union of the two collections. The first argument for this impossibility is that if the Stability and Correct counting conditions are combined with the fact that tallying cannot distinguish between accumulations and new starts, then the result of tallying two collections will be the same as the result of tallying their union. The other argument assumes that Casullo's scenario actually holds, and the results of the two

¹⁹ Wilson [2000], p. 58.

²⁰ Wilson [2000], p. 75.

tallies are different. This is quickly found to be epistemically impossible.²¹ The equivalency of the results tallying two distinct collections and tallying their union shows that even Wilson concludes the cardinal and arithmetical aspects of number cannot be separated as Casullo's scenario requires.

Wilson's argument against Casullo's disconfirming scenario is convincing, and effectively demonstrates that the arithmetical aspect of number cannot be separated out from the cardinal aspect. This result supports Castañeda's argument against Gasking, in which he shows that the mathematics we have is the mathematics we *have* to have, given the world that we live in.²² Wilson has convincingly argued that Casullo's scenario fails to demonstrate that the propositions of mathematics can be disconfirmed by experience. However Casullo's disconfirming scenario is fatally flawed from the beginning because of the way he sets up his task: he looks to find a scenario in which a *true* proposition of mathematics is *falsified* by experience. If the Stability and Correct Counting conditions are satisfied, then this scenario should be impossible or, as Wilson has shown, incoherent.

The problem with these accounts of the Irrefutability Argument is that they look to the wrong mathematical propositions for evidence that experience can disconfirm mathematical propositions. The proposition $2 + 2 = 4$ should not be able to be disconfirmed by experience (nor by any other means) since it is true. Mathematical truths should only be able to be confirmed by experience. In order to determine whether premise (1) of the Irrefutability Argument is true we should investigate whether a *false* mathematical proposition can be disconfirmed by experience. Such an investigation would quickly show that mathematical propositions *can* be disconfirmed by empirical evidence. We could choose any false mathematical proposition, say $2 + 2 = 15$, and go out to test it in different physical systems in which the Stability and

²¹ See Wilson [2000], pp. 79-80 for a full account of the epistemic impossibility.

²² See Chapter 8 for a more detailed discussion of the views of Gasking and Castañeda.

Correct Counting conditions hold. This test would certainly provide us with strong disconfirming evidence for the proposition $2 + 2 = 15$, and this evidence is a posteriori. It is misleading to claim that mathematical propositions cannot be disconfirmed by experience on the basis of the lack of disconfirming scenarios for true mathematical propositions. There are countless disconfirming scenarios for false mathematical propositions, just as there are confirming scenarios for true mathematical propositions. Thus the first premise of the Irrefutability Argument is false, and hence mathematical propositions can be confirmed by experiential evidence.

6.3 The Shift in Meaning Argument

It has been argued above that only true mathematical propositions can be confirmed by experience, and only false ones can be disconfirmed. In some situations we are confronted with an experience that appears to violate these principles, but in such situations we can evoke an escape clause, which may be called the Shift in Meaning argument. It is explicated as follows:

Shift in Meaning Argument: if an experience disconfirms a true mathematical proposition or confirms a false mathematical proposition, then the object of the experience is not a system that instantiates the mathematical structure to which the proposition refers.

A typical application of this argument is as follows: take the proposition of Euclidean geometry which states that the internal angles of any triangle add up to 180° , or two right angles. If we measure the angles of a triangle and find that they add up to, for example, less than 180° , then we can assert that the triangle is not a Euclidean triangle. We may find that it is drawn on an hyperbolic plane, such as the inner surface of a bowl. By denying the fact that we are dealing with Euclidean geometry we protect the true mathematical proposition from falsification. Another example of an application of this argument is adding raindrops on a piece of glass, and observing that $1 + 1 = 1$. In this situation we would invoke the Shift in Meaning argument by

claiming that raindrops on a piece of glass to not instantiate the natural numbers (since Wilson's Stability condition is not met), and hence the observation cannot be considered to falsify the proposition $1 + 1 = 2$.

There are two important considerations regarding the Shift in Meaning argument. First, it is important to note that for true mathematical propositions it is not *in fact* possible to refute them using empirical evidence, however it is still *logically* possible that they be disconfirmed by experience. The proposition $2 + 2 = 4$ cannot be falsified by experience because the natural number structure is a feature of our world. However it is logically possible for a world to exist in which this proposition could be falsified by experience, for example a world consisting of only one atom. Nevertheless, given that our world is as it is, it is not factually possible for a disconfirming scenario for $2 + 2 = 4$ to exist. This is the same as claiming that mathematical truths are *physically necessary*, rather than being *logically necessary*. Castañeda, in his refutation of Gasking's attempt to show that the mathematics we use is not an objective feature of the world, demonstrates that we cannot use a different kind of mathematics and still describe the world. This notion, the physical necessity of mathematical truths, is discussed in greater detail in Chapter 8.

That mathematics is physically necessary, or built into the physical world, is the reason why true mathematical propositions cannot be disconfirmed by relevant experience, and false ones cannot be confirmed. For this reason we must have the Shift in Meaning escape clause, because when experience appears to contradict mathematical truth there has to be something wrong. However the second important consideration regarding this argument is that there are limits on the circumstances in which it may be applied, otherwise it would be circular. If we could invoke the Shift in Meaning argument whenever empirical evidence contradicts our preferred mathematical theory then mathematics would be entirely arbitrary, or at the very least highly subjective. Mathematics is an objective feature of the world, empirical evidence contributes to our knowledge of mathematical truths, however in evaluating

the evidence we must be able to determine which observations are relevant, and which ones we can disregard under the Shift in Meaning argument.

How do we determine whether or not we can invoke the Shift in Meaning argument in a particular situation in order that it not beg the question? There is some degree of subjectivity, as individuals must use their best judgement. The prime consideration pertinent to this question is the degree to which we are certain of the proposition that is potentially being falsified, as well as the nature of the physical system providing the falsifying observation. In the case of the above examples, we have a great deal of evidence for the sum of the angles of a triangle being 180° , including countless observations and also deductive proof. One counterexample is insufficient to disconfirm this belief, which is a good indication that the Shift in Meaning argument may be applicable. When we examine the surface on which the recalcitrant triangle is drawn and find that it is not flat, this finding supports the indication that these are reasonable circumstances to invoke the Shift in Meaning argument. Similarly, our evidence for $1 + 1 = 2$ is so convincing that we are reluctant to give it up based on a single noncompliant observation. If we closely examine the behaviour of raindrops on a piece of glass it is evident that they do not remain distinct, and this is sufficient evidence that we may use the Shift in Meaning argument. There is a degree of subjectivity involved in this decision process. It is similar to the process involved in evaluating disconfirming scenarios for scientific theories: we must consider what evidence we have in support of the theory, and also the extent to which the observations are accurate. However in the case of mathematical knowledge there is more agreement, especially at the basic level, as to which theories are and are not accepted. This results in more consensus than in most scientific areas as to when this escape clause may be invoked.

6.4 Role of Proof

The role played by formal proof is crucial in determining how we acquire mathematical knowledge. Foundationalism and the Euclidean ideal are based on the notion that for every true mathematical result there must be a rigorous deductive proof. While at first this seems reasonable, a closer inspection of the methods employed by mathematicians as well as any one of us that has some mathematical knowledge reveals that this is not necessarily the case.

First of all, I wish to introduce the notion that there are true mathematical results that have no proof. A probable example of such a result is the Riemann Hypothesis, which states that all the zeros of a certain (easily defined) function of a complex variable z lie on the line: real part $(z) = \frac{1}{2}$.²³ This hypothesis is widely believed to be true, and has never been proved false, but there is no known proof for it. While many mathematicians believe that it will one day be proved, it is conceivable that this will never happen. Gödel's second incompleteness theorem tells us that there are more true propositions in any mathematical system than there are provable propositions, in other words, that there are some true mathematical results for which there exists no proof. The result states that if consistent, arithmetic is not complete. This means that not all true propositions have proofs, hence there are always more true propositions than there are provable ones. Often a proof can be found by moving to a higher-order or meta-theory, or by introducing new axioms, but although proofs can then be found for previously unprovable propositions, new truths will be generated in the meta-theory for which there are again no proofs in the new theory.

This problem suggests our mathematics would be enriched if we accept mathematical conjectures that have been extensively tested and never disproved, as they may fall into this category of results. Our mathematics will be more powerful if it

²³ The line formed by plotting the complex numbers whose real component is $\frac{1}{2}$.

includes propositions whose truth we have good evidence for. Rejecting such propositions as false outright because of the lack of conclusive proof would weaken our theory. This means that the quasi-empirical method of knowledge acquisition is in fact an appropriate means for generating mathematical knowledge. Knowledge acquired in this way may not be conclusive but we can still use it for productive purposes.

Furthermore, the perceived role that deductive proof plays in our mathematical epistemology is sometimes questionable in itself. Irvine²⁴ draws attention to Russell's regressive method using Russell's example concerning the mathematical proposition " $2 + 2 = 4$." Every reasonable person with a basic grasp of arithmetic would hold this proposition to be true. While there does exist a deductive proof of $2 + 2 = 4$, it is unlikely that the average person in the street would be able to understand this proof, let alone reproduce it. This would indicate that the formal proof for $2 + 2 = 4$ does not have a significant epistemic function for most people who hold it to be true, and thus their justification must come from elsewhere. The proof for $1 + 1 = 2$ was only formalised less than a century ago,²⁵ but this does not mean that previously no one was justified in believing it to be true. The formal deductive proofs of these very elementary arithmetical results do not provide our justification for holding them to be true. The reason we believe $2 + 2 = 4$ to be a true statement is that we constantly observe physical models which confirm this claim. Whether we derived the proposition initially from our own experience or were informed of it by a teacher, it is the continued verification in our experience of physical objects that causes us to hold it with such conviction.

When a formal axiomatic system allows us deductively to prove $2 + 2 = 4$, a proposition we know (or at least strongly suspect) to be true, this regressively justifies

²⁴ Irvine [1989], p.308.

²⁵ The proof of $1 + 1 = 2$ appeared in Whitehead & Russell [1913], Part III, Section B, *110.643.

the axiomatic system. This result lends implicit support to other propositions whose primary justification is derivation from these axioms. Thus we construct a coherent web of interrelated beliefs, rather than building from a foundation of axioms. Russell explains that this process works the same way both in mathematics and in what he calls “the sciences of observation.”²⁶ He distinguishes between empirical premises and logical premises, where the empirical premises are what cause us to believe the proposition in question, and the logical premises are those from which it follows deductively. It is often the case, both in science as well as in mathematics, that the empirical premises and logical premises can diverge. Russell uses the example of $2 + 2 = 4$, explaining that “the proposition “2 sheep + 2 sheep = 4 sheep” was probably known to shepherds thousands of years before the proposition $2 + 2 = 4$ was first discovered; and when $2 + 2 = 4$ was first discovered, it was probably inferred from the case of sheep and other concrete cases.”²⁷ The logical premises of $2 + 2 = 4$, which are “certain principles of symbolic logic,”²⁸ do not have the same epistemic function; in showing that the proposition follows from logical premises we provide further justification for our belief, but the logical premises were not responsible for our original belief that $2 + 2 = 4$.

This process is much the same as how we obtain scientific knowledge. Russell explains that in every science we begin with a collection of propositions of which we are fairly certain, and seek to determine either what follows from them, or what propositions they follow from. We begin with empirical premises, and try to find logically simpler principles that our empirical premises follow from. The notion of logical simplicity of a proposition has to do with “the number of its constituents,”²⁹ so the proposition “ $2 + 2 = 4$ ” is logically simpler than “2 sheep + 2 sheep = 4 sheep”

²⁶ Russell [1907], p. 272.

²⁷ Russell [1907], p. 272.

²⁸ Russell [1907], p. 273.

²⁹ Russell [1907], p. 273.

because the latter has one more constituent, namely “sheep.” If we begin with empirical premises we can generalise to logically simpler propositions, and this is how we get the laws of science. Russell explains that two things are required for these laws to become as certain as the empirical premises: first, the empirical premises must be shown to follow from these laws; and second, we must demonstrate that we can *only* derive the empirical premises by using these laws. Often we have the first of these conditions but not the second, so the laws are only probable, rather than certain.³⁰ Many scientific theories are accepted as likely, even though we cannot show that they are *required* in order to explain our experience of the world. Mathematical theories, however, have a higher degree of certainty because we do need the proposition “ $2 + 2 = 4$ ” to be true in order to account for the empirical premise “2 sheep + 2 sheep = 4 sheep.” This shows why mathematical knowledge is often held to be more certain than scientific knowledge, even though we acquire them both in the same way.

Putnam, in his paper “What Is Mathematical Truth?”³¹ tells a story about what he calls “Martian Mathematics.” He asks the reader to imagine that we have made contact with another civilisation on Mars and, having learned their language with minimum difficulty, proceed to read all their literature throughout all fields. In this fantasy, the Martians’ mathematical text books bear a significant difference from our own: they include as assertions many of the results that our best mathematicians have tried and failed to prove (such as the Riemann Hypothesis). When we read the textbooks to learn their proofs of these results, we discover that they use “methods that are analogous to the methods of the physical sciences”³² in mathematics, and that the statements which are used to test hypotheses are products of proof and calculation

³⁰ The extent to which we are certain of the laws is limited to the degree of certainty with which we believe their empirical premises, but often laws of science do not attain even this degree of certainty.

³¹ Putnam [1975].

³² Putnam [1975], p. 51.

rather than being physical observations of the usual scientific type. In other words, in evaluating whether to accept a mathematical conjecture in the absence of a deductive proof, we can look at such factors as whether it can be used to prove desirable results, if it yields expected conclusions, and so on.

Consider, for example the Riemann Hypothesis: extensive computer testing has failed to find a counterexample, it has been used in the proof of many theorems and none of these consequences have been disproved, so we could say that it has been verified by quasi-empirical means. In Putnam's fantasy, the Martians consider quasi-empirical verification to be sufficient grounds on which to accept a mathematical result, even though it might be disproved at a later time. After all, this is a widely accepted method in the physical sciences. When it is claimed that the Martians do not understand our notion of 'proof' and are questioned on this matter, Putnam imagines that they reply as follows:

What you call 'proof' is simply deduction from principles that are (more or less) self evident. We recognise proof, and we value proof as highly as you do – when we can get it. What we don't understand is why you restrict yourself to *proof* – why you refuse to accept *confirmation*.³³

Putnam points out that Gödel's theorem shows that it is not necessarily the case that if mathematical statements are true then they are provable, a fact which lends further support to the validity of quasi-empirical verification as a method of justifying mathematical knowledge.

Putnam then goes on to claim that the Martians are in fact us: that we do use confirmation along with proof. He shows that we have used both empirical and quasi-empirical methods to generate axioms, giving examples from mathematical history of results that have been accepted based on quasi-empirical or empirical evidence. For example, Euler's proof that the sum of the series $1/n^2$ is $\pi^2/6$ was not really a proof at

³³ Putnam [1975], p. 52.

first. He used an analogy to arrive at $\pi^2/6$ and when mathematicians calculated the sum of $1/n^2$ to thirty or so decimal places they did not doubt that it was $\pi^2/6$. It was twenty years before Euler came up with a rigorous proof. Putnam stresses the similarity with empirical science: “intuitively plausible though not certain analogies lead to results which are then checked ‘empirically’. Successful outcomes of these checks reinforce one’s confidence in the analogy in question.”³⁴ The use of such methods reinforces the similarity between mathematics and other natural sciences.

So far I have presented five arguments to demonstrate that mathematical knowledge is justified by quasi-empirical means: (1) Lakatos’ comparison of the Euclidean ideal to a quasi-empirical account finds the latter to be a closer description of actual mathematical practice. (2) Following consideration of the debate between Casullo and Wilson over whether mathematical claims can be falsified by experience, I concluded that it is possible for a mathematical proposition to be falsified empirically, although it is physically impossible to refute a *true* mathematical claim in this way. (3) Gödel’s Second Incompleteness Theorem lends further support to the view of mathematics as quasi-empirical, since we cannot rely solely on deductive proof to generate mathematical knowledge if there are more true propositions of mathematics than there are provable ones. (4) Russell’s regressive method proved to be an accurate description of how we obtain both mathematical and scientific knowledge. (5) Finally, Putnam’s thought experiment about Martian mathematics demonstrates that we do use other means than pure deductive proof to generate mathematical knowledge, and that many of our axioms come from empirical or quasi-empirical inferences.

Having shown that mathematical knowledge is generated, at least in part, by quasi-empirical means, it now remains to determine the implications of this

³⁴ Putnam [1975], p. 56. See pp. 53-57 for Putnam’s other examples of empirical and quasi-empirical methods in mathematics.

conclusion, and there are two that I focus on. First, given that our mathematical knowledge comes from our experience of the world to some extent, the Argument from Quasi-Empiricism (as explicated earlier in this chapter) proves that the subject matter of mathematics is objective and a part of the physical world. If mathematical entities are in the physical world then epistemic structuralism is the best theory of mathematical ontology, because it accounts for their physicality but avoids identifying mathematical entities with specific physical objects. Identifying them with structures or patterns instead, we can account for their seemingly abstract nature and applicability in a wide range of physical systems. Using the formal sciences, which have similarities both with mathematics and with more traditional sciences, I will show that the quasi-empirical component of mathematical knowledge acquisition involves perceiving and reasoning about mathematical structures that we observe, and in this way we can find the necessity of mathematical truths in our observations of the physical world. The formal sciences are useful in this respect because the structures they study are perhaps more obviously physical than the structures of pure mathematics.

The second corollary of quasi-empirical methods contributing to mathematical knowledge is that mathematical knowledge becomes a lot more like any other kind of scientific knowledge. In particular, this means that it arises from the physical world and is fallible, rather than certain. I will use an argument from Tymoczko to show that computer proofs highlight the close resemblance between mathematical and scientific knowledge, and raise issues of fallibility in mathematics.

6.5 Formal Sciences Bridge the Gap

Having established that the process by which we acquire mathematical knowledge is quasi-empirical, the next task is to determine the implications of this conclusion. I hope to show that a quasi-empirical story about mathematical knowledge acquisition supports a structuralist account of mathematical entities. The

formal sciences are useful here because they illustrate very clearly the structuralist position with respect to mathematical entities, since they take structures from the physical world and reason about them in the abstract. The process is very close to how we do any kind of mathematics, except that the structures studied in the formal sciences are perhaps more clearly instantiated in the physical world, so that the gap between reality and the abstract system is smaller than in pure mathematics. In this way, the formal sciences can be thought of as narrowing the perceived gap between pure mathematics and empirical science, and in so doing, they illustrate the similarities between the methods of the two disciplines.

Franklin³⁵ gives an account of the formal sciences, explaining that they began as ‘operations research’ and are often very case specific. The formal sciences originated just before and during World War II. They began with operations research, starting as investigations into questions concerning the most efficient search patterns for submarines and the best size for convoys, and consisted of multidisciplinary scientific teams theorising abstractly about actual physical events. The methods of the formal sciences have now been extended to include problems such as task scheduling and bin packing, and usually the process involves searching through the possibilities for completing the task and using mathematical theorisation to rule out incorrect cases. Operations research is now considered to constitute a field of abstract study because, while it is concerned with physical systems, an abstraction of the system is used to make calculations and develop theories about the behaviour of the system. Included under the broad heading of operations research are subjects such as scheduling theory, communications networks and highway traffic. Control theory is another old branch of the formal sciences, and is concerned with adapting a given system in order to produce a desired result, by varying the settings. Network analysis,

³⁵ Franklin [1994].

or the mathematical study of flows using computer modelling, is an important discipline with many applications in telecommunications networks, as well as traffic control and information management. Branches of game theory analyse systems affected by the kinds of decisions that agents make in various competitive or cooperative environments, and these disciplines have significant applications in economics. All these disciplines fall under the classification of the formal sciences.³⁶

The significant distinguishing feature of the formal sciences for Franklin's purposes is that the structures they deal with are so clearly a part of the physical reality that it is evident that their results can be projected back onto the physical system. They are all relatively new mathematical theories that can be treated as abstract structures, but which closely mirror physical phenomena. It is unclear whether to classify the formal sciences as a branch of mathematics or of science, and this conundrum has two important effects, both of which support the view that mathematics is the science of structures. The fact that results in the formal sciences apply equally in both the abstract and the physical system is especially useful for structuralism because it gives an example of how truths about structures considered abstractly translate directly to truths about their physical instantiations. When Franklin claims that the structures of the formal sciences are particularly close to reality, this means that the abstract model (the concept) and the physical system have a great many features in common. Mathematical structures are just as much a part of reality as those studied by the formal sciences. The other effect of the formal sciences seeming to fall under both the categories of physical science and mathematics is that they can be used to show that there is no significant distinction between the two.

³⁶ See Franklin [1994], pp. 514-521 for a more detailed survey of formal sciences.

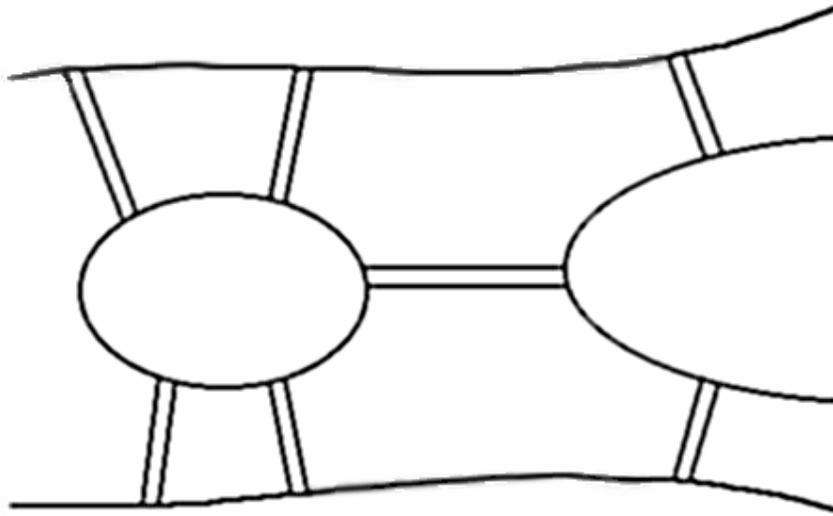


Figure 1: The Bridges of Königsberg.

A simple example which Franklin gives of such a mathematical system, but one which can be used to illustrate his point, is that of the Königsberg bridges. In this example there are seven bridges over the River Pregel at the city of Königsberg (see Figure 1) and, given certain elementary conditions, it is impossible to walk over each bridge exactly once. Euler proved this result mathematically, namely by abstracting the structure of the physical system and reasoning within the mathematical model that this abstraction produced. When his result is projected back onto the physical situation, this gives a physically necessary³⁷ truth about the world. While it could be argued that the arrangement of bridges is a purely contingent matter, and hence there is no necessity involved, Franklin asserts that “[t]he necessity lies in the connection between the bridges having the arrangement they do and the properties of paths

³⁷ Physical necessity is distinct from logical necessity; refer to Chapter 8 for a detailed discussion of the necessity of mathematical truths.

through them.”³⁸ Furthermore, he explains that the necessity must exist in the physical world, since the arrangement of bridges and the properties of paths are elements of a physical system, rather than an abstract model.

This brings out the important notion of physical necessity, which is discussed in greater detail in Chapter 8. However, it is important to note that the necessity of claims in the formal sciences arises out of the structural relations in the physical system, and that we are able to perceive these necessities. These examples bring out two significant aspects of necessity: first that there can exist necessity in the physical world, rather than it arising purely from the way we define our terms or from logical tautologies. Second, it is evident that we can *perceive* this necessity in the physical world. We shall see that the necessity of mathematical knowledge can be characterised as this kind of physical necessity.

A point which is worth emphasising here is not that mathematics can make necessary claims about reality, but that statements which hold within a mathematical system can be translated into statements about reality. The reason that the examples of the Königsberg bridges and the other formal sciences are striking is that the gap between the abstract model and the physical system is smaller than in pure mathematics. Another way of expressing this fact is that the abstract structure and the physical model have more features in common than in the case of other mathematical structures. The abstract model is a mental copy that we have of the structural aspect of the physical system. A structure from the formal sciences will leave out fewer characteristics of a physical system than a mathematical structure. Recall that abstraction, which allows us to extract the mathematical structure out of a system that we perceive, consists of simplification since we leave out the non-mathematical aspects of the object of our perception. In the case of the formal sciences we leave out

³⁸ Franklin [1994], p. 530.

less of what we perceive. Another way of putting this is that in the formal sciences we do not simplify as far as we do to obtain mathematical structures, because the structures that concern the formal sciences are metaphysically richer than most pure mathematical structures.

Consider, for example, a row of people which instantiates both an element of the natural number structure and also a queue. More features of the physical system are significant when we consider them as a queue than if we are looking simply for numerical properties. In both cases we can disregard or ignore many of the physical features of the people, such as height, eye colour, and so on. In order to consider the queue which they instantiate, we need to take into account how many there are, the order in which they are standing, and from which ends people join and leave the queue; however to consider the number exhibited by the row of people, we need only take into account how many there are. Thus the queue structure captures more features of the physical system than the numerical structure, hence the gap between the queue and its physical instantiation is smaller than between the number and its exemplification.

This example also overcomes one of the standard objections to the claim that mathematics can make assertions about the real world, which is based on the argument that mathematical truths, as opposed to scientific claims, are not derived from our experience of the physical world.³⁹ It has been claimed that the statement “ $2 + 2 = 4$ ” is true only because of the way that we have defined the terms “2,” “4,” “+” and “=,”⁴⁰ and consequently cannot say anything about relations between physical objects. There is no way that this objection can be applied to the case of the

³⁹ This comes from the Platonistic claim that mathematical truths are about Platonic Forms, which are not a part of the physical world. Hence our knowledge of mathematics does not come from our ordinary perceptual experience of the world, but from a faculty often called ‘mathematical intuition’ (cf. Gödel [1947]).

⁴⁰ This is the formalist account of mathematical truth, as argued by Hilbert, for example.

Königsberg bridges, because it is plainly obvious that the statement “it is impossible to cross each bridge exactly once” is not a conventional truth based solely on definitions. It is a clear example of how a statement of a truth about mathematical structure, expressed independently of physical instantiations, ‘hooks onto’ its physical instantiation. This is a truth about a mathematical structure that holds equally true in the physical model and the mathematical theory that captures it, and thus provides a strong case for the structuralist position.

Additionally, for the present concern, this example serves a useful purpose in blurring the distinction between natural science and mathematics. Many would claim that scientific facts are about things in the world, while mathematical facts are either about abstract, non-physical entities, or are purely manipulations of symbols. Statements in the formal sciences are unquestionably about things in the world, and yet they seem to be mathematical in nature too. The process of identifying the underlying structure that many physical systems have in common, for example in queuing theory, and reasoning abstractly in this mathematical model seems to be very similar to the way we do mathematics if we examine it closely. Thus we find that mathematical and scientific methods are more similar than most philosophies of mathematics would lead us to believe.

6.6 Computer Proofs

The four colour theorem (4CT) states that any map on a plane or a sphere can be coloured with no more than four colours in such a way that no neighbouring regions are the same colour. This result resisted proof until 1977, when Appel, Haken and Koch published a proof,⁴¹ and the problem is considered to be solved. However this is an especially interesting case, because the proof relies on computer results.

⁴¹ Appel, Haken and Koch [1977].

Tymoczko⁴² considers whether a computer proof can really be held to be a proof in the traditional sense, and what implications the acceptance of the solution to the 4CT has for mathematics as a whole.

A computer proof differs significantly from traditional mathematical proofs. Tymoczko identifies three main characteristics of proof: proofs are convincing, surveyable and formalisable. Since being convincing is a psychological notion that is highly subjective, the other two characteristics of proofs (that they are surveyable and formalisable) are taken to explain what makes proofs convincing. Surveyability has to do with the perspicuity of a proof: the proof is “an exhibition, a derivation of the conclusion, and it needs nothing outside of itself to be convincing.”⁴³ Surveyability is subjective, since it relates a proof to a given knower or mathematician, but I agree that it is a feature of proofs that contributes to their being convincing. Formalisability is related to surveyability, and most mathematicians and philosophers would agree that all surveyable proofs are formalisable.⁴⁴ The objectivity of mathematical proofs comes from their formalisability since a formal proof, a finite sequence of formulas of some formal theory that deduce the conclusion from the axioms using logical rules, is objective. Provided that there is no mistake in the application of rules there is some objective fact of the matter as to whether or not the proof is true, or holds. Hence formalisability also contributes to the convincingness of a proof, and is tied up with surveyability as formalising a proof makes it easier to survey, and it is generally assumed that any surveyable proof must be formalisable.

While the proof of the 4CT is considered to be formalisable, it is not surveyable, because it relies on an appeal to results produced by a computer.

⁴² Tymoczko [1979].

⁴³ Tymoczko [1979], p. 59.

⁴⁴ Intuitionists would not agree with this claim.

Tymoczko makes an analogy to another mythical community of mathematicians on Mars:

Martian mathematics, we suppose, developed pretty much like Earth mathematics until the arrival on Mars of the mathematical genius Simon. Simon proved many new results by more or less traditional methods, but after a while began justifying new results with such phrases as ‘Proof is too long to include here, but I have verified it myself.’ ... So great was the prestige of Simon... that the Martian mathematicians accepted his results; and they were incorporated into the body of Martian mathematics under the rubric ‘Simon says.’⁴⁵

The point of this story is that if we substitute “by computer” for “Simon says,” we have not made any great change to the logical structure of the proof of the results, however it is hardly a controversial claim that the appeal to Simon does not constitute a legitimate basis for proof. In order to be able to accept computer results as reliable elements of a proof without having to accept “Simon says,” we must show why computers are more dependable than Simon.

Accepting computer results as part of a mathematical proof is a form of appeal to authority, however computers are an authority that we trust because we have good evidence for their reliability. There are two aspects of computers that a computer proof relies on, namely the hardware and the software. When we make an appeal to computer results in a proof, our result depends on the reliability of the physical components of the machine, as well as the reliability of the program we use to generate the results. There is no way to give a traditional mathematical proof of the reliability of these two aspects of the computer, the best we have is empirical evidence for the truth of their reliability. The degree to which we are certain of the reliability of the hardware and software of the computer constitutes a maximum limit for how certain we can be of a mathematical theorem proven by a computer proof. This introduces a degree of fallibility into mathematical knowledge which was previously

⁴⁵ Tymoczko [1979], p. 71.

thought to be completely certain, yet the result of the 4CT is considered to have been proved.

The significance of computer proofs in revealing an empirical component in mathematical proof has been debated in recent years. Azzouni⁴⁶ claims that there *is* in fact a significant distinction between proofs that rely on computer generated results, and those that are surveyable in the traditional sense. He argues that a computer cannot provide a true proof, but can only supply empirical evidence for the existence of a proof. Consequently, he claims that mathematicians “have a more slender epistemic grip on such things [i.e. computer proofs] than they do on proofs that are actually possessed by the mathematical community.”⁴⁷

To support this claim, Azzouni offers a thought experiment in which a mathematician proves the existence of a counterexample to the 4CT. He assumes that the proof of the counterexample is carefully checked and found to be sound, and that years of scrutiny of the computer proof fail to turn up a flaw. In this situation, Azzouni claims there are two possible courses of action: the mathematical community could draw the conclusion that mathematics is inconsistent; alternatively, they could assume that the computers have made a mistake. Azzouni speculates that the mathematical community would likely reject the computers’ reasoning, and that this event would result in diminished faith in computer results thereafter.

This thought experiment does not suffice to prove that computer-generated proofs are significantly distinct from other types of mathematical proof. It is only speculation that the computer proof would be rejected rather than the proof of the counterexample. There is always the option to reserve judgement until such time as the examination of both the original (computer) proof and the counterexample turns up a flaw in one or the other. Further, the situation Azzouni described could arise in

⁴⁶ Azzouni [1994].

⁴⁷ Azzouni [1994], p. 170.

the case of *any* proof to which a counterexample was found. The same dilemma would exist, and each case would result in meticulous scrutiny of both the original proof and the proof of the counterexample. This situation is a possibility whether the original proof is computer-generated or not. In any case where a contradiction is found, all possible factors must be taken into account. The fact that the original proof is partially generated by a computer is only one such factor, and while this fact does give the original proof a higher degree of empirical content than a non-computer-generated proof, this is not the only relevant fact. Each individual proof must be examined on its own merits, and we cannot speculate that the computer-generated proof would be the more vulnerable in the majority of cases.

Burge⁴⁸ also challenges the claim that computer proofs make mathematical justification empirical. He uses a fairly broad definition of a priori knowledge, arguing that sense experience can still play a role in supporting a priori knowledge, but that the important question is where the “force of our warrant” comes from. He claims that while sense experience is responsible for our acquiring certain supporting information, it is not responsible for our acceptance of the belief. This is based on a distinction between “epistemic or justificational dependence” from other kinds of dependence, such as “causal or psychological dependence, or dependence for grasping intentional content or for learning.”⁴⁹ Burge claims that sense experience can contribute to belief acquisition without providing the normative force that causes us to hold the belief.

In the case of computer proofs, Burge maintains that the force of our warrant for holding the belief is a priori, rather than empirical. I will not go into his argument in great detail, as I would dispute some of his early assumptions, however his argument is not at odds with the claim that I am making: namely that mathematical

⁴⁸ Burge [1998].

⁴⁹ Burge [1998], p. 2.

knowledge, much like other kinds of scientific knowledge, has empirical content. In arguing that computer proofs result in knowledge that is a priori, just as he considers all mathematical knowledge to be a priori, Burge claims that sense experience and empirical beliefs have a role to play in all mathematical knowledge acquisition. Even in traditional mathematical belief acquisition we have assumptions about the good working order of our brain processes, and often too about the brains of others who supply proofs to us. There are a range of ways in which our sense experience of the physical world contributes to mathematical belief acquisition, and although Burge argues that this empirical component never supplies the force of our warrant for holding a mathematical belief, it is still the case that the process of forming mathematical beliefs does rely to some extent on empirical content. Indeed, in arguing that computer proofs remain primarily a priori in their justification, Burge provides evidence that the knowledge they produce is in essence no different from other mathematical knowledge, since *all* mathematical knowledge is impacted by empirical belief. This is exactly the point I wish to emphasize.⁵⁰

Traditionally mathematics was supposed to have no empirical content and to rely solely on proof, rather than on the type of experiments found in natural science, and this made mathematical truths more certain than scientific truths. However the reliance on computer results radically changes this understanding of mathematical knowledge, as there is a fallible empirical element to proofs by computer, which are nevertheless accepted as demonstrating the truth of a theorem. The reliance on the empirical truths of the dependability of the machine and the program used removes the divide between mathematics and science, since it shows that there can be a degree of experimentation and probabilistic reasoning in mathematics and this does not

⁵⁰ That empirical evidence impacts all mathematical knowledge is only part of Burge's argument, evidently. However I will not go into his full argument for the a priority of computer proofs, as I do not accept some of his assumptions, and here is not the place to dispute them.

invalidate the result. It may be the case that we are more certain of many mathematical beliefs than we are of other beliefs, but this does not mean that they belong in their own privileged category of entirely certain beliefs. Mathematical truths can involve an empirical component without being demoted to a separate category, and the use of computer proofs illustrates this undeniably.

The fact that proofs which involve an appeal to computer results, and hence a fallible empirical component, are so readily accepted as being just as good as more traditional mathematical results, supports the notion that the more basic mathematical beliefs involve an appeal to empirical truths and can be fallible. Even my belief that $2 + 2 = 4$ relies on empirical truths, because I can see that it is true every time I count some objects of my perception. If I had never added 2 objects and 2 more objects together and ended up with 4 objects, it is unlikely that I would believe $2 + 2 = 4$ with the certainty I do. Russell's regressive method of justification for the axioms of mathematics suggests that basic mathematical truths are supported primarily by their discovery and constant verification in our *experience* of the world (as Mill argued), and it is primarily our experience – our interaction with mathematical structures – which causes us to believe them. This being the case, it should not be surprising that a mathematical theorem proven using a computer is so readily accepted, because fallibility and reliance on empirical truths are already important factors in mathematical knowledge. The fallibility of mathematical knowledge and its reliance on empirical truths, especially as brought out by the use of computer results in mathematical proofs, is a further demonstration that mathematical knowledge is much like any other kind of scientific knowledge.

6.7 Formal Systems

Treating mathematical knowledge as a type of scientific knowledge supports my project because it is consistent with the notion that mathematical results are not absolute. Skolem's paradox tells us that set theoretical results are relative, and my

argument is that this points to a more general phenomenon, the inadequacy of formal systems in capturing reality. However just because formal systems cannot completely capture reality in its entirety, this does not mean that they cannot be useful in describing the world. In Lakatos' *Proofs and Refutations*⁵¹ we see how the concepts that we use are sharpened and developed in more detail as we investigate the things they apply to and extend their boundaries. In this way, quasi-empirical investigation is an important part of developing a theory. In order to improve any given theory we have to examine its applications, and by investigating where its applications fall short we can make adjustments that render the theory more powerful and accurate.

Formal systems always fall short (to some extent) of capturing reality completely. Is this because there is something about reality that makes it in principle out of our reach, that it is impossible for us to come up with a theory that captures reality entirely? Or is it just that we are finite creatures and so we just keep improving and improving our theories, and just like we cannot count to infinity so our theories cannot capture the enormous complexity of the world in its entirety? I think we cannot know which is the case, but I am inclined to say that there is nothing about the world that makes it impossible in principle for us to develop theories that entirely capture reality, however since we have limited access to the world and we are finite in nature, this becomes a practical impossibility.

The more important issue here is how we go about sharpening or improving formal systems. Lakatos gives a very good account in *Proofs and Refutations*, which demonstrates the importance of observation and interaction in the evolution of concepts. This suggests that the quasi-empirical method employed in the 'traditional' sciences is also an important part of mathematical knowledge acquisition. We start out with an intuitive concept that corresponds to a mathematical structure that we observe

⁵¹ Lakatos [1976].

in the world, and by pushing the limits of the concept we can sharpen it and improve it. In this way our theory about that mathematical entity evolves and becomes an increasingly powerful tool for describing a particular aspect of the world. The following chapter demonstrates the importance of mathematical concepts in the process of knowledge acquisition. This emphasises the theory-dependence of mathematical concepts, and supports the view that mathematical knowledge is not significantly different from any other variety of scientific knowledge.

CHAPTER 7: MATHEMATICAL CONCEPTS

This chapter examines the role of mathematical concepts in the process of mathematical knowledge acquisition. Epistemic structuralism is largely characterised by an epistemic approach to understanding mathematics, thus it is by examining the way we come to have mathematical knowledge that I hope to understand the nature of mathematical entities. Previous chapters have described the process of abstraction and argued that mathematical knowledge bears no significant difference from other kinds of scientific knowledge. It is a feature of scientific knowledge that the theoretical concepts under investigation evolve along with the theory. It is argued in this chapter that the same is true of many mathematical concepts; they are dependent on the theory with which we describe them, and our descriptions of the world are characterised by the concepts we use to understand aspects of it. In mathematics, as with the rest of science, we find that theory and concepts are inextricably intertwined. It is also argued that our references to mathematical concepts are incomplete and hence fallible, but this does not preclude us from acquiring increasingly accurate knowledge of mathematical reality.

7.1 Justification of Mathematical Knowledge

In the previous chapter I argued that mathematical knowledge can be considered to be a species of scientific knowledge. Breaking down the division between mathematics and science, and between analytic and synthetic knowledge,¹ we start to get a unified picture of our knowledge about the world. On this account, all knowledge is fallible and all knowledge is subject to the same kind of evaluation. If we reject the idea of mathematical knowledge as purely analytic and independent of

¹ Quine argued against the legitimacy of the analytic-synthetic distinction; see Quine [1951].

experience, and admit that there is at least a quasi-empirical element to mathematical knowledge, then the dividing line between mathematical knowledge and other natural scientific knowledge disappears, as argued in the previous chapter. Our knowledge of mathematics can now be treated in the same way as any other knowledge we have about the world, since it is an empirically-driven (at first) investigation into mathematical structures, an objective feature of the world.

In the following chapter we shall see in what way mathematical knowledge is fallible, but here I shall examine the various ways in which mathematical knowledge can be justified, since this will further illuminate the nature of mathematical concepts. Epistemic structuralism and an account of abstraction show how mathematical beliefs are acquired, but we must also examine their justification to show why these beliefs constitute knowledge. The traditional Euclidean picture of mathematical knowledge involves deduction from axioms which are innately known necessary truths. However, given Quine's argument against a priori knowledge and the arguments for basing our mathematical knowledge on our experience of the world and the process of abstraction, we must abandon this ideal and examine other ways in which our mathematical beliefs can be justified.

Brown's book *Philosophy of Mathematics: an introduction to the world of proofs and pictures* aims to show that pictures can constitute our primary proof for many of the truths of mathematics, and that deductive proof is not the only method of acquiring mathematical knowledge. He gives an example of Bolzano's "purely analytic proof" of the intermediate zero theorem and two other related theorems. Brown provides the following proof for the intermediate zero theorem which is not exactly Bolzano's proof of 1817, but is "in the modern spirit which he created":

Theorem: If f is continuous on the interval $[a, b]$ and f changes sign from negative to positive (or vice versa), then there is a c between a and b such that $f(c) = 0$.

Proof: Assume (with no loss of generality) that $f(a) < 0 < f(b)$. Let $S = \{x: a \leq x \leq b \ \& \ f(x) < 0\}$. This set is not empty, since a is in it; and it is

bounded above by b , so it has a least upper bound, c . There are three possibilities.

$f(c) < 0$. If this is true there is an open interval around c , i.e. $(c - \delta, c + \delta)$, in which $f(x) < 0$, for all x in the interval including those greater than c . This contradicts the assumption that c is an upper bound.

$f(c) > 0$. If this is true there is an open interval around c , i.e. $(c - \delta, c + \delta)$, in which $f(x) > 0$, for all x in the interval including those less than c . But that's impossible since c is the least of all the upper bounds, so that $f(x) < 0$ for all x less than c .

$f(c) = 0$. The other two possibilities being ruled out, this one remains. And so, the theorem is proved.²

If deductive proof is the only way that we can be sure of a mathematical truth, then it follows that before Bolzano's proof we could not be certain of the intermediate zero theorem. However Brown shows two other methods that could convince us of the truth of this theorem. The first is illustrated in Figure 1:

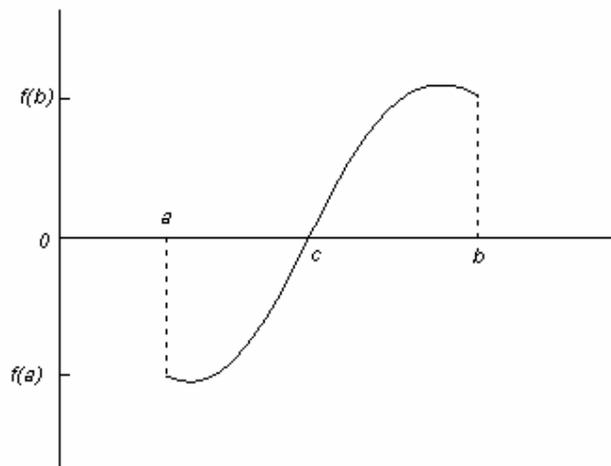


Figure 1: Intermediate Zero Theorem

We can see that if $f(a) < 0$ and $f(b) > 0$, then if f is a continuous function on the interval $[a, b]$ it must cross the x -axis somewhere between a and b , and this crossing

² Brown [1999], p.26.

marks the value of c . This diagram is just as convincing as the analytic proof, if not more so. We can see instantly by looking at the picture that the curve must cross the x -axis, and thus there must be a zero between a and b . The analytic proof can not be considered to increase our certainty of the truth of the intermediate zero theorem, since once we have seen the picture we cannot doubt the theorem. A similar diagram can be used to illustrate the intermediate value theorem, which is a generalisation of the intermediate zero theorem.

There is another closely related theorem which states that if f and g are both continuous on the interval $[a, b]$ and $f(a) < g(a)$ and $f(b) > g(b)$, then there is a c between a and b such that $f(c) = g(c)$. This theorem can be demonstrated by a thought experiment, or visualisation of a physical problem. Brown³ gives the example of a mountain climber who starts at the base of a mountain at noon, reaching the top at 6pm, and takes the same path down the next day, again starting at noon and reaching the bottom at 6pm. The question is whether there is a time during the second afternoon at which she was at the same point on the path as she was at the same time the previous day. The answer is yes, and this is clearly illustrated by considering two hikers, one at the bottom heading up (the journey the original hiker did on the first day) and one at the top heading down (the journey she did on the second day). It is obvious that the two will meet at some point on the path, and the time that they meet is the time the original hiker would have been at the same point on both days. Brown's claim is that this thought experiment proves the theorem just as conclusively as an analytic proof.

There are various ways of interpreting the fact that both the picture and the analytic proof (as well as the solution to the thought experiment) achieve the aim of convincing us of the theorem. I agree with Brown's assessment, namely that the proof

³ Brown [1999], pp. 27-28.

is confirmed by the theorem. While some would argue that Bolzano's proof confirmed a theorem that was previously believed, but not known to be true, it is a more plausible explanation that we already knew the theorem was true, and that the proof merely showed that analytic proof could lead us to this true conclusion. The confirmation of the proof by the theorem that we can directly perceive to be true is a good demonstration of the reliability of the method of analytic proof. This is an example of the regressive method of proof first outlined by Russell,⁴ in which we believe the premises of an argument because they are used to prove consequences which we know to be true. In this case we believe that the reduction of geometry to algebra and arithmetic does not involve falsification, since an analytic proof can be used to demonstrate known truths. This means that we can confidently believe other results obtained in this way. As Brown points out, “[w]e can draw the moral quickly: *Pictures are crucial*. They provide the *independently-known-to-be-true* consequences that we use for testing the hypothesis of arithmetization.”⁵ This shows that pictures provide our primary justification for some truths of mathematics, and are an essential component in our acquisition of mathematical knowledge.

Picture proofs and analytic proofs are not the only way that mathematical knowledge can be justified. Brown lists six methods that we can use to acquire mathematical knowledge: proofs, intuition, induction, hypothetico-deductivism, pictures and diagonalisation.⁶ Proofs are used as a tool for convincing people of the truth of a mathematical conjecture, and finding a proof increases our belief that the conjecture is true. Intuition and pictures are related, because intuition can be considered a non-rigorous way of extracting mathematical truth from our experience of the world, and seeing pictures is one form of experience. So when we acquire

⁴ Russell [1907].

⁵ Brown [1999], p. 29.

⁶ Brown [1999], p. 180.

mathematical knowledge from pictures we are using intuition to gain mathematical knowledge from a particular observation, namely that of the picture. Hypothetico-deductivism and diagonalisation are just different ways of finding proofs.

The methods used to acquire mathematical knowledge can be divided into rigorous and non-rigorous methods. Non-rigorous methods are reading off pictures (as in the intermediate zero theorem), inductive generalisation from experience, and other less specific ways of intuiting truths about mathematical structures we experience, which we may lump together under “intuition.” Rigorous methods are typically used to confirm mathematical truths and present them in such a way as to convince others. These are various forms of proof, and include Brown’s categories of hypothetico-deductivism and diagonalisation. As Brown notes, there is no single method by which we acquire mathematical knowledge. There are several different ways that our mathematical beliefs can be justified, and proof is only one of them.

Mathematical truths were not all originally arrived at by deduction from hypotheses. Generally we start from a conjecture, or a belief that we suspect to be true. This suspicion may come from intuition, observation, or just some sort of a “hunch.” A mathematician will then search for a proof which serves to increase his or her own degree of belief in the conjecture, and will also convince others of its truth. In many cases pictures might be the origin of a conjecture, and as in the example from Bolzano above, an analytic proof is provided as a kind of formality to convince others of its truth. Traditionally pictures have not been considered to constitute adequate justification for mathematical knowledge, but this example demonstrates that the picture is just as convincing as the analytic proof.

The fact that picture proofs can be as convincing in justifying mathematical knowledge as deductive proofs has two important corollaries. The first is that it shows that mathematical knowledge can be justified by means other than deductive proof, particularly observation, and if pictures can justify mathematical truths it would suggest that other physical exemplifications of mathematical structures can do so as

well. The second corollary of this wider view of justification for mathematical knowledge is that it supports epistemic structuralism by removing any significant distinction between mathematical knowledge and other kinds of knowledge that we have about the world. This is crucial in absorbing our mathematical knowledge into a unified epistemology that can account for *all* our knowledge about the world, which is preferable to having a separate epistemological theory for our mathematical knowledge. This is related to the argument from uniform semantics, which favours a philosophy of mathematics that does not require a distinct semantics to account for our mathematical language. The application of mathematical knowledge to the physical world shows that the truths of mathematics are facts about the world, thus our epistemology should be able to account for mathematical truths in the same way as any other truths. This is one of the strong advantages of epistemic structuralism over other philosophies of mathematics that require a special epistemology for mathematical knowledge. Furthermore, the role of background assumptions, concepts and theory involved in mathematical knowledge now becomes crucial.

7.2 Importance of Mathematical Concepts

Mathematical knowledge does depend on human thought in a certain sense, and this has to do with the representations we use to describe the mathematical structures we perceive in the world. Brown explains that mathematical knowledge is fallible,⁷ not in the sense that our beliefs are outwardly wrong, but that we can have correct beliefs about the wrong⁸ concepts. For example, there are many representations that we can use for the natural numbers. The structure itself is an

⁷ This is only one sense in which mathematical knowledge is fallible; a more thorough discussion of fallibility of mathematical knowledge occurs in Chapter 8.

⁸ For a concept to be wrong means that it does not accurately capture the mathematical structure in question. Some concepts are better or more appropriate than others, but a wrong structure can yield mistaken results. Such concepts can often be confirmed by quasi-empirical methods.

objective feature of the physical world, but in order to reason about it we need some way of representing it. Generally we tend to use the Arabic numerals, but we could, if we chose, use the Roman numerals, set theory, or even a completely arbitrary representation. For example we could give the numbers random names so that instead of counting 1, 2, 3, 4, ... we count Bob, Sally, Frank, Jim, ... and so on. Each of these representations is a human (mental) construction and is dependent on us because it exists only in our minds, however what each of them represents is the natural number structure, an objective feature of the world. We are free to choose whichever representation of the numbers we wish, but when we use it to express mathematical truths and to reason about them, it is likely we will find that some representations work better than others. The reason that the Arabic numerals are so widely used to express the truths of arithmetic is because they work so well.

Brown argues that the reason we find Arabic numerals so useful in describing truths about the natural numbers is because the notation contains information about the thing that it stands for. For instance we can tell at a glance which of two numbers is the larger, and if a number is divisible by ten we can recognise this instantly by the fact that it ends in a zero. A numeral tells us a lot more about the number it represents than names of other objects do, such as “electron.” Brown points out that “[t]he crucial feature that is built into both Roman and Arabic numerals is their recursiveness.”⁹ He explains that this feature of the notation is what led formalists to do away with numbers altogether in their ontology, and make do with just the numerals. But as Brown points out, this is a case of putting the cart before the horse, since the Arabic numerals were developed deliberately so as to mirror the recursiveness of the numbers. On this view the numbers exist independently of the notation, and the reason the notation is so useful is that it captures significant features

⁹ Brown [1999], p. 80.

that the numbers possess. In this way, use of different concepts to understand, explain and reason about mathematical facts can help us understand different aspects of mathematical reality.

This fact is clearly brought out by Brown's discussion of various forms of notation used in knot theory, each of which bring out certain properties of knots, particularly the fact that "*notations reveal properties, they do not create them.*"¹⁰ A knot is something that we are all familiar with from ordinary life, and in mathematics a knot is defined as a closed, non-intersecting curve in space. Brown explains that the main problems of knot theory involve classifying and distinguishing different knots. Knots may be thought of as a piece of string with the ends joined (such as a child may use to play the game "cat's cradle," or an electrical extension cord plugged into its own end). The more common mathematical representation of knots is a two-dimensional projection of such a piece of string, as this can be drawn on a page. Given a projection of a knot, there are three allowable transformations that may be performed on it. These are known as Reidemeister moves.¹¹ Two knots that may appear different are said to be equivalent when one can be deformed into the other without cutting the curve or allowing it to intersect with itself. If one projection of a knot can be transformed into another by a sequence of Reidemeister moves, then this is one way of showing the two projections to be equivalent, namely that they are different projections of the same knot.¹²

Brown goes on to discuss two different kinds of notation used to describe and classify knots. The first is Dowker notation, which involves labelling the crossings of knots as ordered pairs and ultimately lets us express a knot as a sequence of even numbers. Given such a sequence (or set) we can draw a projection of the knot

¹⁰ Brown [1999], p. 89.

¹¹ These were introduced in an example in Chapter 3.

¹² See Brown [1999], pp. 81-86 for a more detailed description of these issues in knot theory.

described, because the Dowker sequence contains all the information we need. Figure 2 shows a knot labelled with Dowker notation. This notation is useful for classifying knots, and it has some interesting features. For example any crossing labelled by a pair of consecutive numbers can be eliminated by a Reidemeister move.

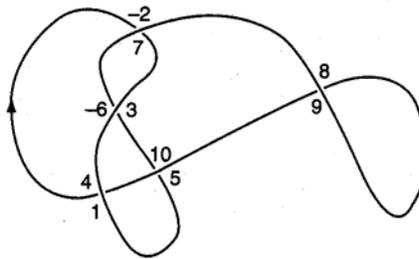


Figure 2: Knot labelled with Dowker notation¹³

A different kind of notation in knot theory is Conway notation. This notation is based on what is known as a tangle, which is “any region of the projection plane which is surrounded by a circle in such a way that the knot crosses the circle exactly four times.”¹⁴ This notation is based on the number of twists, and the directions in which the lines are twisted. A tangle is labelled by numbering each twist that it contains, so an example of how Conway notation can be used to express a tangle would be $3\ 2\ -4$. Figure 3 shows a picture of this tangle.

¹³ From Brown [1999], p.84.

¹⁴ Brown [1999], p. 84.

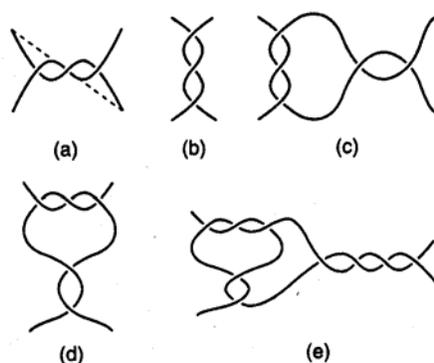


Figure 3: Construction of a $3\ 2\ -4$ tangle¹⁵

We can construct a continued fraction from these numbers, and it turns out that if two tangles are equivalent they will have the same associated continued fractions.

Yet another way to represent knots is by polynomials. There are various ways to do this, but Brown focuses only on one of the simpler methods. There are a series of rules for describing knots as polynomials: the first rule states that the unknot (a loop with no knots in it) is the number one; further rules specify methods for constructing polynomials of crossings and various links. It is not really necessary to go into further detail here, what is important is that it is possible to calculate a unique polynomial for any given knot. Interestingly, it turns out that polynomials of knots are invariant under Reidemeister moves, meaning that equivalent knots have the same polynomial.

Each of these three representations bring out different interesting aspects of knots, but they all describe the same structure. For any given knot, we could express it in either Dowker or Conway notation, or as a polynomial, and each representation would give us different information about it. However there is just one thing in the world that we are describing, and this does not change depending on which notation we use. Just as we may use “4,” “3 + 1,” “IV” or “||||” all to describe the number four,

¹⁵ From Brown [1999], p. 85.

the same truths hold about the knots we are describing using the various notations. Although there is a single structure in the world, there are several different representations available to us.

Brown uses knot theory as a paradigm example of the more general fact that for all mathematical objects there are different types of representation, and different representations reveal different aspects of the system in question. For example, writing a polynomial as an equation brings out its degree, which is not always evident if the polynomial is represented by a diagram. However a diagram will tell us the zeros (solutions) of the polynomial directly, just by looking at the curve and reading off where the curve crosses or touches the x -axis. If we wanted to find the zeros based on the equation of the polynomial there would be a certain amount of calculation involved, so these different ways of representing the same thing (the polynomial) reveal different properties of the thing in question.

This is not true only of polynomials, for example we could use different methods to represent the natural number structure in order to reveal different properties of the numbers. In our standard base-10 notation we can easily see if a number is divisible by 5: its last digit will be 0 or 5. If we were to use hexadecimal¹⁶ notation we would lose this facility, but it would be replaced by others. Hexadecimal notation was brought into use by early computer programmers because it is easy to translate hexadecimal numbers into binary numbers: each digit becomes a four-digit “bit”. This makes it an effective way of abbreviating computer code that is in binary format. To abbreviate binary code in standard base-10 notation would be a cumbersome process and hardly worthwhile, however hexadecimal notations are extremely useful for this purpose. So while hexadecimal and decimal (base-10)

¹⁶ Hexadecimal notation is base-16 notation: 10 is represented by ‘a,’ 11 by ‘b,’ ... 15 by ‘f’ and 16 by ‘10.’

notations bring out different properties of the natural number structure, they both describe the same objective thing in the world.

7.3 Brown's Argument from Notation

These examples concerning notation are introduced to support an argument that Brown makes in favour of Platonism. The argument is based on the claim that notations reveal, rather than create, properties. It can be set out as follows:

- Premise 1: Mathematical entities can be described by a variety of notations.
- Premise 2: Different notations reveal different properties.
- Premise 3: If different notations reveal different properties of the same thing, then the thing that they describe must exist independently of the notation.
- Conclusion: \therefore Mathematical entities are independently existing, purely abstract (Platonic) entities.

There is nothing in this argument to support the conclusion that mathematical entities are purely abstract, in the Platonic sense. Instead, the argument only supports the claim that there exists (“behind” the representations) an objective reality which contains the mathematical entities and all their properties; it does not follow that this objective reality is abstract, independent of space and time, eternal and unchanging. We can denote the same object by the phrases “the couch” and “the largest piece of furniture in the living room.” Both refer to the same item and reveal different properties of that item: the first connotes a long cushioned bench-like object used for sitting on, and the second tells me that it is an unspecified piece of furniture that is larger than every other piece of furniture in the living room. This suggests that there exists an object that possesses these characteristics, and that it exists independently of my describing or thinking about it. However it does not imply that the object described is a purely abstract entity, existing outside of space and time.

Based on the argument Brown makes, mathematical entities could just as well be physical structures as Platonic forms. The fact that we have knowledge of different properties of mathematical entities, and that we can come up with notations that capture these properties, suggests that we are acquainted with the objects that they represent. Since we have different notations by which we refer to mathematical entities, it certainly makes sense that there is something objective in the world that we are describing with these various notations. This provides evidence for the existence of mathematical structures as a part of the world that we can experience, and then seek to describe.

If mathematical entities were abstract in the Platonic sense, it becomes more difficult to account for the different representations we have for them, compared to if they are structures in the world that we can experience. This comes down to the Problem of Access (as Brown calls it), which was discussed in an earlier section and is a serious problem for Platonism. However, the epistemic structuralist account is supported by the varying notations used in mathematics. If someone were to observe that the knot in their shoelace exhibited the same structure as a knot in the garden hose, he might seek a way to describe the common structure in itself, independently of these two instantiations. It is perfectly reasonable and highly likely that different people would come up with different notations. This is due to the fact that different people think differently and do things in different ways, and also because they might focus on different aspects of the structure. To use the example of polynomials, the graphic representation reveals the zeros of the function, whereas an equation tells us its degree. Each of these properties is deducible from either notation, but which property we wish to reveal more obviously is a subjective issue. However the structure being described and the properties it possesses remain objective, no matter which representation we use.

If mathematical entities are structures instantiated in physical systems, then it follows that we would have various ways of describing them that draw attention to

different properties. The fact that this is the case supports epistemic structuralism. Why should we have just one way of referring to mathematical structures? We rarely if ever have just one way of referring to something that we experience. Discounting different languages, we still find that there are a great variety of ways that we can refer to things we observe, depending on what we want to show about them. “The green couch” draws attention to the colour of the piece of furniture, while “the big couch” conveys an impression of its size. If I am making the reference to something that I have observed, both its size and colour are a part of the object that I’m referring to, the object designated by “the green couch” and “the big couch” is the same object, but attention is focused on one of its attributes. This is analogous to the example with polynomials: if I draw a graph of a polynomial I am showing its zeros and gradient, whereas by writing down the polynomial’s equation I reveal its degree, although all this information is contained in both representations. This supports epistemic structuralism since it indicates that mathematical entities are objective features of the world that we can perceive.

It should be noted here that the argument from notation supports the structuralist aspect of my ontology of mathematical entities, but does not provide conclusive evidence for the physicalist aspect. The fact that there are many different ways of referring to mathematical entities indicates that they are objective and not dependent on human thought. By stating that $1 + 4 = -3 + 8 = +\sqrt{25}$ and so on we can reveal different structural properties of the number 5. This argument supports structuralism, and while it provides no evidence for Platonism, nor does it imply a specifically anti-Platonistic structuralism. The argument from notation proves that mathematical entities are real and independent of us, and suggests that we have epistemic access to them, but it is neutral with respect to their ontology. It is not a conclusive argument in favour of my view, but it provides support for an aspect of it. This argument must be considered in conjunction with my earlier arguments against

Platonism and in favour of other aspects of epistemic structuralism in order appropriately to characterise my view on mathematical entities.

When we refer to mathematical entities we are relating a fact about the world that is independent of us. The fact that there are some assumptions¹⁷ underlying the mathematical concepts we use does not make the truth that is represented any less objective. However we must choose some language or notation to characterise our (non-Platonistic) abstract conception of the physical mathematical structure, and often the choices we make affect the expression of mathematical truths. Most of our mathematical truths are expressed as claims about uninstantiated structures, however the reality that they are describing and that makes them true is the physical reality of the world. In the process of abstraction from physical reality to concepts of mathematical structures we must leave out some aspects of the physical system. We then often extend the resulting concept. Some ways of performing this process might yield more effective mathematical concepts that give us a better grasp of mathematical reality. Use of different abstract concepts to describe mathematical reality brings out different features of that reality.

7.4 Theory Dependence of Mathematical Concepts

Lakatos, in his book *Proofs and Refutations* gives an example that illustrates how mathematical concepts are dependent on theory. This is an important consideration in an investigation into the nature of mathematical knowledge, since the relationship between the theories we use to describe mathematical structures and the concepts that these theories yield determines what we know about mathematical reality and how we express that knowledge. Epistemic structuralism claims that our

¹⁷ For example, the stability and correct counting conditions for the natural numbers, as explicated by Wilson [2000].

mathematical knowledge is based on concepts that we acquire by abstraction, so it is important to examine how the concepts are refined and developed, given an observation. Once we observe a physical system and isolate the mathematical aspect using abstraction, this gives us a mathematical concept. The theory that we develop about the mathematical concept shapes the concept itself, and the concepts that we use determine which kind of mathematical theories we will have.

The example Lakatos gives takes the form of a classroom discussion, in which a teacher presents a conjecture to some students.¹⁸ They are initially convinced of the truth of the conjecture but soon come up with counterexamples. The example concerns Euler's conjecture that for all polyhedra there is a relation between the number of vertices (V), the number of edges (E) and the number of faces (F), namely $V - E + F = 2$. A proof is given, but then a number of counterexamples to the theorem are raised. It becomes evident that the truth of the theorem and the scope of a proof depend significantly on our understanding of the concepts involved. Lakatos demonstrates that our understanding of concepts evolves during the course of finding a proof and refuting counterexamples.

When we first state Euler's conjecture we have a naïve conception of what a polyhedron is. As we construct a proof we come across counterexamples that challenge the conjecture, and our concept of a polyhedron is sharpened and made more specific in order to maintain the truth of the conjecture. The first counterexample Lakatos provides is a solid bounded by a pair of nested cubes, so it is a cube with a smaller cube cut out from its centre. For this solid $V - E + F = 4$, so it does not conform to Euler's conjecture. This counterexample prompts the revision to the concept of a polyhedron as a "surface consisting of a system of polygons."¹⁹ Further counterexamples give rise to further restrictions of what we may classify as a

¹⁸ Lakatos [1976], p. 6.

¹⁹ Lakatos [1976], p. 14.

polygon: a pair of tetrahedra that are joined by an edge or joined by a vertex lead us to specify that exactly two polygons must meet at every edge, and it should be possible to get from the inside of any polygon to the inside of any other polygon by a path that does not cross an edge at a vertex. Kepler's star-polyhedron gives rise to the clarification that each face should be a polygon such that exactly two edges meet at every vertex and the edges of a vertex have no points in common except the vertices. There are a few more modifications to the concept of a polyhedron, but the method is clear. Lakatos initially refers to this method as 'monster-barring', since each of the solids that violate Euler's conjecture are considered to be monsters that should not be classified as polyhedra. However this method serves to clarify and more precisely specify the concepts that we use.

The process of proof and refutation that Lakatos captures in this example takes us from the naïve conception of polyhedra to the theory- or proof-generated concept of a "simple polyhedron." This is the result of all the restrictions and specifications that have been introduced in order to rule out each of the counterexamples. There are several different ways of looking at this process. It could be argued that each of the 'monsters' attempts to stretch our concept of a polyhedron, and in barring these monsters we are simply preserving the original concept. On this view, our original conception of a polyhedron (which Lakatos does not specify in any exact terms, other than a solid bounded by polygons) was too narrow, and the process of monster-barring enables us to increase our knowledge by finding out more about what polyhedra should be like. However there is an alternate view which states that when we introduce more specifications in the face of counterexamples we are contracting our original naïve concept of a polyhedron.

Either way, the proof-generated "simple polyhedron" is a more specific but more narrow concept than our original naïve understanding of polyhedra. The latter concept is one that we gained directly by abstraction from our perceptions of the world. However the concept of a "simple polyhedron" with specifications such as "not

more than one face meeting at an edge,” and so on, is a much more highly evolved concept. This is a concept that we likely would not have come up with just by observing the world and abstracting from our perceptions. Lakatos claims that “each proof yields its characteristic proof-generated concepts,”²⁰ and that these replace the naïve concepts with which we begin our investigations into the truth of a conjecture.

7.5 Incompleteness of Reference

This example of the process of proof and refutation involving the sharpening and narrowing of our concept of polyhedron shows how mathematical concepts (namely our conceptions or ideas of mathematical structures) often do not completely capture mathematical reality. In the following chapter, we will see that this is a source of fallibility in mathematics. When we initially recognise a mathematical structure in the world we begin to develop a concept of it. This (mental) concept is initially a naïve concept, which is then refined and clarified as we utilise the concept in mathematical processes. The method of proof and refutation is a way for us to get a better grip on the mathematical entity as we learn more about it and refine our concept of it. Thus we develop a more specific idea of mathematical structures, and increase our knowledge of mathematical truths. Whenever we prove something about a mathematical entity we acquire more information about the structure itself, and we can specify it more completely. Thus we describe mathematical entities more successfully as our concepts of them are sharpened, and we become better at capturing mathematical reality. However this is an ongoing process, and we can never completely capture or describe a mathematical entity.

²⁰ Lakatos [1976], p. 90.

Twentieth century mathematical logic has demonstrated that mathematics is incomplete.²¹ This is an effect of our inability to refer completely to mathematical structures. In Lakatos' example of a simple polyhedron, we start with a naïve concept which is then refined and narrowed by a process of proof and refutation. After a while, our concept of a simple polyhedron contains a lot more information about the structure being referred to, however there is no end point to this process. We can make a referring term or a concept more complete, but it will never be entirely complete (because of our human limitations). Just as axiomatised set theory is unable to capture our intuitive notion of set theory (as revealed by Skolem's paradox), this should not be a surprising outcome, since the same principle applies to all reference, including ordinary non-mathematical reference.

Consider, for example, my concept of a particular car. I know the car by acquaintance and can successfully refer to it, however my reference to it remains incomplete. If I say "the blue car over there" I will usually make a successful reference, however my reference to the car is incomplete because my concept of the car is incomplete. I am acquainted with the car, but not acquainted with it in its entirety. I might know its make and model, number plate, colour and how many doors it has, however I do not know its age, how many cylinders it has, the chemical composition of its paint, and so on. No matter how thoroughly I investigate and how much information I acquire about the car, my concept of it will always be incomplete because there will always be something that I do not know about it (for instance the number of molecules in the gear stick). Even ordinary objects are so complex that we can never have complete concepts of them, so our reference to them is always incomplete. The same is true of mathematical objects.

²¹ See Boolos and Jeffrey [1974] for proofs of Gödel's incompleteness theorems.

Luckily, incompleteness of reference does not mean that we cannot refer at all. It is possible to refer successfully without having a complete concept of the object being referred to. Incompleteness of reference means that there will always be undecidable sentences about the objects of reference. For example, I can successfully refer to a person without knowing the number of cells in their body. My lacking this knowledge does not prevent me from referring to the person in their entirety (although not completely).²² In fact, I am unable to determine how many cells are in the person's body and, even if I could, this fact would quickly change. This means I cannot determine whether the sentence "that person has 10^{14} cells in their body" is true or false, hence for my purposes it is an undecidable sentence. If a reference is incomplete then it follows that not everything is known about the object of reference. References that we make to objects in the world are incomplete because of the complex nature of objects, so there will always be undecidable sentences. This is true both of ordinary everyday objects as well as mathematical objects, and is supported by the proven incompleteness of mathematics.

In the following chapter I shall evaluate the purported certainty of mathematical propositions, and argue that our mathematical knowledge is fallible. One sense in which mathematics is fallible arises from the view that mathematical concepts are theory-dependent, and we may refine our theories to capture mathematical reality with increasing accuracy, depth and scope. If a mathematical theory is not complete then some mathematical claims arising from that theory are subject to revision. Thus mathematical truths are not logically certain, nor do they express necessary truths. However since mathematical theories describe mathematical structures that are an objective feature of the world there is some necessity about

²² I have referred to the person in their entirety because I have successfully referred to the whole person (not just, for example, their height, or their age). This is not the same as making a complete reference to the person.

them. I will develop the notion of *physical necessity* and explain why this is an accurate characterisation of mathematical knowledge, as well as expanding on the senses in which we can consider mathematical knowledge to be fallible.

CHAPTER 8: CERTAINTY, NECESSITY AND FALLIBILITY

As far as the propositions of mathematics refer to reality, they are not certain, and as far as they are certain, they do not refer to reality.¹

Epistemic structuralism, by taking an epistemically driven approach to discovering the nature of mathematical entities, removes the distinction between mathematical knowledge and other kinds of knowledge about the world. One potential objection to this approach arises from the notion that mathematical truths are privileged over all other kinds of truths, that they are indubitable. In this chapter it is argued that mathematical truths are *not* privileged in this way, and that mathematical knowledge is indeed just like any other knowledge about the world.

Mathematical propositions are commonly thought to be both certain and necessary; these are two different claims. Certainty is an epistemological measure, it is an expression of how sure we are of a given belief. I shall argue that the truths of mathematics are not completely certain, rather they are fallible. Necessity, on the other hand, is a metaphysical notion, an expression of a way the world has to be. Mathematical statements are often supposed to express necessities about a realm of abstract entities. My claim is that the necessity of mathematical truths is a type of physical necessity, and this is really the only kind of necessity there is in our world.

8.1 Certainty

There are two kinds of certainty: psychological and logical. Psychological certainty is like a quotient. It comes in degrees, and can be thought of as a measure of belief. The greater my certainty in proposition p , the more strongly I believe p . Psychological certainty is not always rational, and does not have to depend on

¹ Einstein [1954], p. 233.

evidence, although it is often connected to evidence. It is mental rather than factual, and is specific to a particular subject who holds the belief in question.

The other kind of certainty is logical certainty. A statement is logically certain if and only if it is logically necessary. There are various hypotheses about the nature of logical certainty. I will briefly mention three of the main ones. The first, less controversial, account is that logically certain statements are those whose contradictions are self-contradictory, although this account does not give much insight into the nature of logical certainty. A second, more in-depth explanation is that logically certain propositions are true in all possible worlds, a view put forward by Lewis.² This is a controversial point of view, since it requires a plausible account of what possible worlds are and in what sense they can be said to exist. A third account of logical certainty is based on the notion of analyticity. There is a view that considers analytic truths to be logically certain and known a priori, while synthetic truths required verification through experience and thus could only be known a posteriori. This sharp distinction has been broken down to some extent in the second half of the twentieth century, as explorations of the dichotomy have encountered difficulties. One of the most significant difficulties has been Quine's argument that there are no genuine analytic truths that are not vacuous, and that all logically certain statements are devoid of factual or empirical content.³

Although the nature of logically certain statements is a contentious issue and one that is not my subject here, it is possible to show that mathematical truths are not certain in the logical sense. The negation of a mathematical truth is not self-contradictory. It may be wrong, but it is not a contradiction in itself. It is logically possible that $2 + 2 \neq 4$. For instance, in a world containing only two molecules it is true to say that $2 + 2 \neq 4$, because it is not possible to add two more items to any two

² Lewis [1973].

³ Quine [1951].

items. However in our world it is *false* to say that $2 + 2 \neq 4$, because the structure of Peano arithmetic is built in to our world. Later sections in this chapter demonstrate that the sense in which mathematics is certain is a physical sense,⁴ so it is only in our particular physical world that our mathematics is necessary (recall that necessity is the metaphysical notion corresponding to certainty, which is either logical or psychological). However, although the truths of mathematics are not logically certain, we attribute to them a high degree of psychological certainty.

There are truths of science of which we are just as (psychologically) certain as the mathematical truths to which we are most faithfully attached, and we are far more certain of these scientific facts than of some more complex mathematical claims. For example, I am close to completely certain that $2 + 2 = 4$, barring any wildly improbable scenario in which I live in a mathematically different world and I am being deceived by some Cartesian-style demon.⁵ If I step outside and watch drops of water falling from the sky, hear them tapping on the roof of the car outside my building and feel the top of my head becoming wet, then I would be just as certain that it is raining now in Vancouver as I am that $2 + 2 = 4$. However I might also believe the proposition that it is raining now in Vancouver with less certainty, if my belief was based on less convincing evidence (for example hearing a distant rain-like noise, and having a brief glimpse of grey cloud around the edge of a curtain); whereas there is no foreseeable situation in which my certainty in the belief that $2 + 2 = 4$ would diminish.⁶ Still, under the right conditions and given enough information I can

⁴ While the truths of mathematics are only certain in a physical sense, we shall see that they are physically necessary.

⁵ There are some other cases in which I might relinquish some basic arithmetical beliefs, for instance if we adopt a modulo arithmetic. However a standard arithmetic has proven reliable and consistent with our experience thus far and we are likely to remain fairly certain of standard arithmetical truths unless we are given very good reason to suspect that a different arithmetic will fit our experience of the world better.

⁶ Assuming, of course, that my mental faculties remain acute and are not interfered with by senility or some other mental disorder.

hold these beliefs with equal conviction. Furthermore, I can be much more certain about simple observational and scientific claims than I can about more complex claims of mathematics.

One reason that we are inclined to give privileged status to many (and especially the most basic) mathematical truths is that they are constantly verified in our experience of the world.⁷ Another reason is that they are metaphysically much simpler than many other propositions with which we are faced. The process of abstraction which allows us to acquire generalised knowledge about mathematical structures leaves out many of the complicating features in the physical world. This lets us focus only on the mathematical aspects of a physical system. Mathematical structures are comparatively not so complex (for example, compare the structure of a sphere with the structure of an atom). This simpler way of looking at features of the physical world gives us a very high degree of certainty. However the certainty that we gain is still psychological certainty, which comes in degrees or is like a quotient (albeit a very high quotient in this case). Mathematical truths still may not be considered to be logically certain.

8.2 Physical Necessity

Moving now from the topic of certainty to the metaphysical notion of necessity, it is useful to introduce the notion of physical necessity (*de re* necessity) as distinct from necessity belonging to linguistic propositions or claims (*de dicto* necessity).⁸ Physical necessities are physical phenomena that are characteristic of how the world is, and we tend to ascribe a high degree of certainty to propositions about

⁷ This is Mill's view of why we consider mathematical beliefs to be more certain than other beliefs.

⁸ *De re* necessity means "necessity of things," while *de dicto* necessity means "necessity of words." The difference between these kinds of necessity can also be expressed as the difference between the propositions "some *a* is (necessarily an *F*)" and "necessarily (some *a* is an *F*)." The former is a claim about a thing, saying of that thing that it is necessarily an *F*; the latter expresses a necessary truth, namely that there is an *a* that is an *F*.

physical necessities. These kinds of claims are not necessary propositions in the strict logical sense, since they are contingent on how the world is. Their necessity arises from laws of nature that are true in our world. The necessity of this type of claim lies in the connection between certain physical facts, and the predicted outcome. Claims that can be considered physical necessities may be regarded as a species of ‘hypothetical necessities’, since they are hypothetically true. The hypothesis on which they are contingent is that the physical world is as it is, they are dependent on certain facts about the world, namely laws of nature. While the physical world does not have these specific laws of nature *by logical necessity*, there are many physical necessities that are contingent on these laws. These necessities may be considered hypothetical because their necessity is contingent on the laws of the physical world.

This hypothetical necessity is what distinguishes physical necessity from factual contingency. A factual contingency is just a fact that happens to be true, such as the fact that I am wearing a blue sweater today. I could have chosen a grey one, but since I did not then it is a fact that I am wearing the blue one today.⁹ This is a different category of claim from physical necessity, an example of which is “a ball bearing placed on a smooth, unimpeded ramp (on Earth) will roll down it.” Assuming that we have all the relevant information about the situation (for example that the ball bearing is not coated in heavy glue), the ball bearing will roll down the ramp as a matter of necessity. It does not have the option to rest near the top of the ramp, because given the way that the world is, it is a necessity that the ball bearing must roll down the ramp. The necessity lies in the connection between the physical world being as it is and the outcome, given the facts of the situation. Hence it is a physical necessity.

⁹ The claim that the sweater I chose this morning is contingent may be refuted by a deterministic metaphysics, however for present purposes I am assuming that we have at least some sort of freedom to make choices.

Factual contingencies are true of the world, but could have been otherwise without altering the nature of the world.

The formal sciences are useful in showing how there can be necessity in the physical world, when it can seem that everything physical is contingent. The example of the bridges of Königsberg¹⁰ is a clear example of a physical necessity. The claim “it is impossible to walk across each bridge exactly once” is a necessary truth, a physical necessity. Given that the bridges are arranged as they are and that the laws of nature are as they are, this statement is necessarily true. It does not qualify as a strict necessity, because it is contingent on how the bridges are arranged (and also on how matter is allowed to behave). There is nothing necessary about the physical arrangement of bridges, this is simply a brute fact, a contingent fact that is true of the world. However given that this is a fact about the world, the claim that it is impossible to walk across each bridge exactly once becomes a necessity. The necessity is in the physical world, rather than being an abstract notion, and it is a necessity that we can experience or perceive directly from our interactions with the physical world. This notion of physical necessity will be useful to show why we are so certain of the truths of mathematics.

Frege gives the example of the Pythagorean theorem as a Platonic truth. The reason he considers it to be Platonic is that the thought it expresses is objectively and eternally true. Platonists consider that an objective and eternal truth has to be about abstract entities, since it is independent of contingent physical facts. This view leads to all the problems of knowledge and reference associated with Platonism. A more viable option is to consider the Pythagorean theorem to be a feature of our physical world, a way the world is. It is a truth in our world that whenever a right-angled triangle is drawn on a flat surface, the square of the hypotenuse is equal to the sum of

¹⁰ Recall details from Chapter 6.

the squares of the other two sides. This is a necessity which is embedded in our physical world, there is no good reason why it has to be about abstract entities. This kind of truth is an example of a physical necessity, it is a feature of our world, a necessary truth given that our world is as it is. Considering mathematical truths to be physical necessities not only explains why we hold them with such certainty, it is a view which demystifies the processes by which we can formulate mathematical beliefs and acquire mathematical knowledge.

8.3 Necessity of Mathematics

There is a sense in which the statements of mathematics are necessary, and that is that the truths of mathematics follow necessarily from the nature of mathematical structures. This is because the subject matter of mathematics, the structures we study in mathematics, are part of the world, and so our mathematical truths are about the structures that are given to us by abstraction from our perceptual experience of the world. Since these structures are as they are, they have certain features or properties which we can discover, and these are the truths of mathematics. Since these are truths about structures that are built into the way the world is, these truths are physically necessary. We are very confident of the truths that we know about mathematical structures, because they are truths about the world that we have discovered in a similar way to how we discover scientific truths, but they are often (although not always) much simpler and more accessible than scientific facts. Other sciences often involve a lot more guesswork as to the nature of the structures that they deal with. For example the structure of a sphere is a lot more accessible, simpler and much easier to understand than the structure of the atom. Rather than constructing complicated experiments to test hypotheses we have about atoms, we can perceive spheres (or approximations to spheres) in a multitude of instantiations all around us, and we can interact with these exemplifications of the structure in order to uncover its various features. One reason mathematical truths seem to have a privileged status of

indubitability is their metaphysical simplicity which allows us to know them with clarity; another is that (at least on a basic level) the mathematics we have is the *only* mathematics which can describe the world.

This fact is brought out in a response to a paper in which Gasking seeks to argue the opposite. In his paper “Mathematics and the World,”¹¹ Gasking examines the nature of mathematical propositions and argues that it is possible to use other mathematical rules than the ones we have, and still get around in the world without problems. He supports this claim with various examples of ‘queer’ arithmetics, and shows how they can be used for practical tasks in the physical world, such as tiling a room. He argues that even given the way we count and measure, we could use different mathematics and still describe the world accurately if we are prepared to change our physics. The mathematics we choose does have something to do with how ‘neat’ a physics it lets us have, but this is more psychological than to do with the external world. Castañeda, in “Arithmetic and Reality,”¹² argues against Gasking, claiming that his argument is based on false assumptions and mistaken reasoning. He makes several interesting objections to Gasking’s assumption that mathematical propositions and mathematical rules are the same thing, and that these are incorrigible statements and hence cannot describe the world. However the part of his counter-argument which is of greatest relevance to my present concern is his argument against Gasking that even if we were prepared to alter our physics, there could be no ‘queer’ arithmetic which describes the world.

Castañeda considers what a queer arithmetic is: it is a consistent system of arithmetic that is non-trivially different from our present system. A trivial change of our system of arithmetic is something like changing the symbol for “7” to “*,” or deciding only to use even numerals, so that we can get back to the original arithmetic

¹¹ Gasking [1940].

¹² Castañeda [1959].

by substituting a “7” for every “*” we see, or by substituting $n/2$ for every number n . Such a change does not produce a queer arithmetic, because the arithmetic itself is unchanged, it is merely expressed in different symbols. This is essentially a claim that the mathematical structure is the same, but we describe it in a different way. We could switch to Roman numerals and our mathematics would be unchanged in any significant way. Castañeda’s argument is that the truths of mathematics are necessarily about the structures that they describe, that no other structures could form the subject matter of mathematics. This makes sense because these structures are a part of the world. However we could use different concepts to explain truths about mathematical structures, and this is where the subjective and fallible element of our mathematical knowledge comes from.

Castañeda lays out the five Peano postulates and five counting principles,¹³ claiming that in order to be a queer arithmetic, a system must involve a change in one or the other. He systematically goes through the Peano postulates showing that if it was discarded, the resulting arithmetic would be either only a trivially changed arithmetic, an inconsistent one, or one in which it is impossible to count things. The exception is the fifth postulate, in the case of which abandoning it gives us the transfinite numbers which are not used in counting, so the resulting arithmetic is a wider arithmetic than the original one, but contains the ordinary arithmetic. He goes through a similar process for the counting principles, showing for each that if we give it up we get an arithmetic which is only trivially different from ordinary arithmetic, or one which is incapable of counting. Thus he concludes that “[e]very queer arithmetic, i.e., one which does not turn out to be either an incomplete or a trivially different or changed arithmetic, is inconsistent or inadequate to help us get on in the world in all situations.”¹⁴

¹³ Castañeda [1959], p. 414.

¹⁴ Castañeda [1959], p. 415.

Castañeda's argument is intended to prove that Peano arithmetic is logically necessary, however he succeeds in proving only that it is a built-in feature of *our* physical world. He points out that Gasking's argument comes down to a claim that only by changing the laws of physics could we successfully use a queer arithmetic, for "it is not possible to isolate both counting and measuring from the properties of physical objects."¹⁵ His argument, outlined above, against the possibility of using a queer arithmetic does not demonstrate the *logical* necessity of the arithmetic that we currently use. The failure of a queer arithmetic is manifested in its inability "to help us get on *in the world* in all situations"¹⁶ (my emphasis), but he does not allow for the possibility of a world so different from ours that it lets us use such a queer arithmetic, or even an inconsistent one.

It is logically possible that there exists a world containing only five atoms, and that the physical structure of the world is somehow cyclic so that an appropriate arithmetic for describing the mathematical structure of this world would be a modulo arithmetic in which 1 is the successor of 5. Castañeda considers abandoning the fourth Peano postulate (that 1 is not the successor of any number) and concludes that this would make it impossible to deal with practical problems such as tiling a room. That is true, because *given that our world is as it is* we could not use any other arithmetic with success. However our arithmetic is not logically necessary, it is logically possible that our world could have been so differently structured that we would need a completely different arithmetic in order to describe the mathematical structures of the world. As Castañeda demonstrated, we could only use a *trivially* different arithmetic and still successfully describe the world. It might use different concepts and be an entirely different theory, but upon closer inspection we would find it to be a

¹⁵ Castañeda [1959], p. 410.

¹⁶ Castañeda [1959], p. 415.

description of the same structure. The concepts that we use are discretionary, the mathematical structure they capture is a feature of the physical world.

This argument shows that the truths of mathematics are necessary, however this necessity is physical necessity, it is dependent on the world being as it is. There are objective mathematical truths that are independent of our knowledge or perception of them, since there are mathematical structures that are aspects of objective physical reality. The necessity of mathematical truths is dependent on the physical fact that the structures are as they are, in the same way that the necessity of the claim that it is impossible to cross every bridge of Königsberg exactly once is dependent on the physical arrangement of the bridges. The truths of mathematics are contingent on the way the world is (namely what mathematical structures it contains), but they are necessary because the world is as it is. Given that the natural number structure is a physical feature of our world, when you put one marble next to another marble you will necessarily end up with two marbles. These truths, these structural relations, are what we try to capture in our theoretical conceptualisations of mathematical structures. If we have performed an abstraction which captures all the mathematically significant features of the physical system, then our mathematical system will be useful in describing the world. When we abstract from physical systems that exhibit mathematical structures we are trying to capture all the mathematical features in our abstract system, and this includes the necessity. Our mistaken belief that mathematical truths are necessary truths, and absolutely or logically certain comes from the high degree of psychological certainty we ascribe to them.

Mill¹⁷ gave a good account of why we consider our mathematical beliefs to be certain. He argued that all mathematical knowledge could be reduced to induction, since geometry and arithmetic were only deductive in so far as they followed from

¹⁷ Mill [1872], Book II, Chapter VI.

their definitions or axioms. He considered that these first principles were generalisations from experience, truths which had their basis in our observations of the world. For example the claim that $3 = 2 + 1$ can be thought of as a definition of the number three, but Mill understood this as the articulation of a physical fact about all collections of three things, rather than a purely logical fact. We originally became acquainted with this physical fact through experience, and since it is constantly verified in our perceptions of the world we become increasingly certain of it.

For Mill, it is the constant confirmation through experience which lends the axioms of geometry and arithmetic their high degree of certainty. These truths are justified in the same way as any other knowledge we have, but the empirical evidence is so strong that we believe them with a high degree of certainty. He rejects the idea of mathematical knowledge being justified by some other faculty which can intuit a priori truths, since there is nothing to be gained by postulating such a faculty, and the certainty of mathematics can be explained in terms of empirical justification. While our mathematical beliefs are only as necessary as the generalisations (or abstractions) from experience from which they are derived, these generalisations are among the most highly confirmed among all our empirical beliefs. Mill considered that this explains why we believe mathematical truths with such certainty, and this is what leads us mistakenly to consider them to be analytic truths. This explanation accounts for the high degree of psychological certainty with which we hold many mathematical beliefs, and also fits with a quasi-empirical account of mathematical knowledge acquisition and the epistemic structuralist position that this thesis explicates.

8.4 Fallibility of Mathematics

If we accept realism about mathematical entities and the existence of an objective mathematical reality as a part of the physical world, then it follows that our mathematical knowledge will be fallible. There are three sources of fallibility in our

knowledge of the truths of mathematics: two of these were discussed by Brown (one of which is highlighted by Skolem's paradox), and another was brought out by Mill's account of mathematical knowledge.

Brown¹⁸ outlines two main ways in which our mathematical beliefs can be fallible. The first, fairly trivial way is that we can have false mathematical beliefs which are the result of mistakes in calculations. Our most basic arithmetical beliefs such as $2 + 2 = 4$ are certain to a very high degree, and this is what causes us (mistakenly) to consider them to be necessary truths. The reason we are so certain of these beliefs is that they are constantly verified by our experience of the world, and this lets us recognise that we have captured a physical necessity. However, we are not as certain of all mathematical beliefs. Suppose I perform a long multiplication algorithm and determine that $701 \times 678 = 475,278$. I am fairly certain that it is a true statement, but not as certain as I am that $2 + 2 = 4$. The former proposition cannot be verified directly by my experiences with instantiations of the natural number structure since the numbers involved are too big, and while I have performed a calculation using a method that I know to be reliable, I cannot be certain that I have not made a mistake. Empirical confirmation of this claim would have to rely on counting, which is also a potential source of error. However, I do know that there is some objective truth about the matter, that there is a unique correct answer to 701×678 . This example reflects both the objectivity of mathematical truth and the fallibility of our mathematical beliefs.

There is another way that mathematical beliefs are fallible which was also mentioned by Brown, namely the use of wrong or misleading concepts to describe mathematical truths. This kind of fallibility comes from the fact that mathematical concepts are revisable, in the same way that scientific theories are revisable. Set

¹⁸ Brown [1999].

theory provides a good example here: the current view of set theory is that an iterative version is the correct one. However, until Russell's paradox was discovered at the beginning of the twentieth century, a formulation which can be referred to as naïve set theory was thought to be correct.¹⁹ The discovery of the paradox led to the revision of the concept of set, and now the predominant view is that the naïve conception of set is wrong. Thus at least some truths of this set theoretical system must also be wrong. If this could happen to another formulation of set theory, it could conceivably be the case that the predominant contemporary conception of set theory might one day be similarly overturned.

This kind of fallibility in mathematics is also brought out by Skolem's paradox, which tells us that formal axiomatisations fail to capture our intuitive notion of set theory. The relativity of set theoretical notions, revealed by Skolem's paradox, shows that some claims in set theory are relative to the given interpretation. This lack of absoluteness is an example of the failure of a formal axiomatisation to capture our intuitive understanding of a mathematical structure. In this way, the concepts we use to express mathematical truths fall short of capturing every aspect of the mathematical structure in question. This is analogous to the inability of formal languages to capture fully our use of natural language. The result of this phenomenon is the kind of fallibility that Brown mentions, which arises from discrepancies between the concepts we use to describe mathematical structures, and mathematical reality. This aspect of the fallibility of mathematics is also revealed by Lakatos in *Proofs and Refutations*, which traces the process of defining and refining the concept of a polyhedron. Lakatos demonstrated the difficulty in capturing mathematical reality by following the process by which our concepts approach mathematical reality. The process of mathematical discovery brings our concepts closer to the actual mathematical entity they refer to,

¹⁹ See Boolos [1971] for a description of these two conceptions of set.

but it seems to be an ongoing process. This means that there will always be a gap between mathematical reality and the concepts we use to describe it, and this gap is a source of fallibility in our mathematical knowledge.

From Mill's account of mathematical knowledge as based on generalisations from our experience of the world, we find another source for potential fallibility in our mathematical knowledge. If $3 = 2 + 1$ is a definition of 3, based on all groups of 3 things that we have observed and consistently reconfirmed by our experience, then we can only be as certain of $3 = 2 + 1$ as we are of the observations on which it is based. Descartes showed us that we can always doubt our perceptions, although we generally consider most of them to be fairly reliable in the right circumstances. However the account of our mathematical beliefs as based on abstraction from our perceptions gives another way that our mathematical knowledge is fallible, even though the fallibility is negligible. This shows that the truths of mathematics cannot be considered completely certain, nor logically certain.

Mathematical knowledge is not in fact privileged over all other kinds of knowledge. This misconception is a product of the constant verification of our basic mathematical beliefs through our daily experience, and the ontological simplicity of our mathematical knowledge as compared to our other knowledge about the world.²⁰ An earlier chapter compared mathematical knowledge to scientific knowledge, demonstrating that the use of computer proofs in mathematics introduces a degree of fallibility into mathematics that is similar to the kind of fallibility that exists in the other sciences. The mathematician has to make assumptions about the integrity of the hardware of the computer and the results are dependent on the computer program doing what it is supposed to do and being bug-free. These are not facts that can be verified with absolute certainty, yet the results derived using the computer are taken to

²⁰ Our more basic mathematical knowledge is the most striking in its simplicity; of course, higher mathematics describes much more complex structure.

be true mathematical results. The use of the computer not only introduces a fallible empirical element into mathematical justification, but the fact that the result is readily accepted into our body of mathematical knowledge suggests that previously there was some degree of fallibility in our knowledge of mathematics, so the use of computer proofs did not cause a significant shift in the nature of mathematical knowledge.

Having claimed that mathematical knowledge is fallible and thus neither entirely certain nor expressive of logically necessary truths, we still have an intuitive belief that the truths of mathematics are more certain than any other kind of knowledge we have. In order to account for this intuition, we need to consider the ways in which our mathematical beliefs are justified. Since the truths of mathematics are justified in a variety of ways, this makes them more certain than other empirically-based beliefs. In an earlier section we saw that there are a number of ways that our mathematical beliefs can be justified, including (but not limited to) proof, observation, intuition, induction and so on. While this supports the absorption of mathematical knowledge into a unified epistemology which accounts for all our knowledge about the world, the variety and consistency of justification that we have for our mathematical knowledge adds to the certainty with which we believe mathematical truths. Take for example the intermediate zero theorem. We saw in an earlier section that this theorem can be verified by looking at a picture, but there is also an analytic proof which supports it. This justification and re-justification increases the certainty with which we believe the theorem. The reason we are so certain of the claim “ $1 + 1 = 2$ ” is that every time I put on one shoe and then another I end up wearing two shoes, every time I stand next to one other person there are two of us standing there, and so on. Perhaps I have read and grasped a proof that $1 + 1 = 2$. The variety and constancy of justification serves to reinforce our belief in mathematical propositions, and this is why we believe in them with such a high degree of certainty.

CONCLUSION

This thesis has developed a version of mathematical structuralism to explain the nature and justification of mathematical knowledge. Epistemic structuralism has similar elements to both Hellman's modal structuralism and a theory that takes mathematical structures to be Aristotelian universals. A significant way in which epistemic structuralism differs from other philosophies of mathematics is that its primary concern is with epistemology rather than ontology. This gives it the advantage of integrating mathematical knowledge into a unified naturalistic epistemology.

A serious problem in the philosophy of mathematics concerns finding a solution to the problem Benacerraf identified in his [1973]. Maintaining a standard theory of truth and reference, the traditional ontology for mathematics is Platonism. Mathematical terms are taken to refer to independently existing abstract entities, and the statements of mathematics express truths about these Platonic entities. The problem with this approach is how to account for our knowledge of mathematics, since knowledge requires some causal interaction with the subject. Since Platonic entities are ideal forms that are external to the space-time manifold we inhabit, they are causally inert, so it is difficult to account for our knowledge about them.

This dilemma leaves two routes available: if Platonism is to be retained, then the problem of the lack of a causal connection must be addressed, either by refuting the requirement for a causal connection in truth and reference, or by explaining how we can interact causally with Platonic forms; alternatively Platonism can be refuted, and a naturalistic philosophy of mathematics adopted that accounts for mathematical knowledge using classical theories about truth and reference.

Gödel, Maddy and Brown are among the philosophers that argue for Platonism. Gödel postulates a faculty of mathematical intuition that is mysterious and could not be explained by natural science, but is analogous to sense perception. Maddy¹ follows Gödel in postulating this faculty of mathematical intuition, but she provides an explanation for how it might function, given what we know about the mechanics of other forms of sense perception. Maddy's approach is compatible with the requirement for a causal connection between us and mathematical entities, in order for there to be mathematical knowledge. By contrast, Brown claims that there can be knowledge with no causal connection, and uses this claim to defend a modern version of Platonism. It has been argued in this thesis that Brown's argument against any causal requirement in knowledge formation is based on flawed premises. Explaining the faculty of mathematical intuition in terms of the physical sciences appears to be the only way to salvage Platonism.

Another approach to the dilemma that Benacerraf identifies is to deny Platonism and the independent abstract existence of mathematical entities. Structuralism is one of the more recent theories to adopt this approach, and finds its origins in Benacerraf's [1965]. This paper argues that since there is no unique best way to conceive of numbers as either sets or objects, we have no way of choosing between competing concepts of the natural numbers. Hence the only defining characteristics that numbers possess are their relations to each other. This conclusion led to the structuralist view of numbers, namely that they are recurring structures or patterns. Similar structuralist arguments can also be constructed for other mathematical entities.

Several varieties of mathematical structuralism have arisen. Generally these are theories of ontology: they describe the nature of mathematical entities, argue that

¹ Maddy [1980].

they are structures, and explicate the nature of mathematical structures. This thesis postulates a distinct variety of structuralism – epistemic structuralism – that takes a different approach. Epistemic structuralism takes as its starting point our mathematical knowledge, and by examining how we acquire mathematical knowledge an ontology of mathematical structures is developed. This puts the epistemic structuralist view at a distinct advantage compared to many other philosophies of mathematics, since it does not require a separate epistemology for our mathematical knowledge. The same epistemology accounts for our knowledge of mathematics, and also all our other knowledge about the world.

Another advantage of epistemic structuralism is that it is a form of realism, for which there are at least three persuasive arguments. One is the indispensability argument of Quine and Putnam, which claims that we are ontologically committed to mathematical entities since they are indispensable in science. Another argument for realism with respect to mathematical entities is the argument from uniform semantics. This argument is based on the desirability of having one theory of semantics that accounts for all discourse, both mathematical and non-mathematical. Brown also makes an argument from notation, which claims that the various ways we refer to mathematical entities suggests that they exist independently of us.

These arguments for mathematical realism have traditionally led to Platonism about mathematical entities, which has the problem of accounting for our mathematical knowledge, given the causal inertia of Platonic entities. This is what Brown refers to as the Problem of Access, and gives rise to the dilemma that Benacerraf's [1973] identified. As well as being a serious problem for Platonism, this difficulty is also an obstacle for what is known as *sui generis* structuralism. This term is applied to the version of structuralism put forward by Shapiro, and this theory suffers the traditional difficulties of a Platonistic ontology for mathematics. Epistemic structuralism does not have these problems, since mathematical structures are taken to be primarily physical. Hence we can have causal access to them, and mathematical

knowledge can be explained using the same naturalistic epistemology that we use for all other kinds of knowledge.

Additionally, the epistemic structuralist approach of locating mathematical entities primarily in the physical world avoids the difficulty of showing why mathematics is able to describe the world as accurately as it does. If mathematical structures are just one kind of system in the world that we are trying to understand and explain, then it is evident that truths about mathematical structures will have explanatory power in the physical world. There are several arguments in favour of considering mathematical knowledge to be in a similar category to other kinds of scientific knowledge. Examination of mathematical practice and the role of proof suggests that actual mathematical methods are closer to a quasi-empirical model than an axiomatic-deductive one. Arguments from the formal sciences, which blur the distinction between mathematics and natural science, and from the use of computer proofs in mathematics, provide further support for the claim that mathematical knowledge should not be placed in a distinct category from other scientific knowledge.

The epistemic structuralist approach takes mathematical structures to have no significant differences from other systems that are studied by the natural sciences. Among the implications of this approach are that the notion of mathematical knowledge as privileged and entirely certain must be relinquished. A degree of fallibility is introduced into mathematical knowledge. The main source of fallibility arises from the epistemic structuralist claim that every mathematical theory is a formal system that is an attempt at capturing some aspect of the world. Gödel's theorem, Skolem's paradox and other limitative results have shown us that a formal system will never be able completely to capture the aspect of the world that it describes. However it has been argued that this fallibility is not as problematic as it might seem, and it does not preclude mathematics from having strong explanatory power and hence being a very useful tool. As well, the incompleteness of

mathematical knowledge facilitates its integration into a complete epistemology, since incompleteness of language and reference is a wider phenomenon common to all disciplines that make use of formal systems to capture reality.

This thesis has not attempted to provide an exhaustive proof of epistemic structuralism as the correct theory to explain mathematical knowledge (whether for all of mathematics or some select areas). Such a project would require close and careful characterisation of various branches of mathematics and a very different approach to the one taken. Given the variety of perspectives in the philosophy of mathematics it is debatable whether any particular theory could be proven to be correct. This thesis does not claim to have provided conclusive evidence for epistemic structuralism. Rather, several of the predominant alternatives have been considered, and it has been argued that epistemic structuralism is a strong contender among them.

Additionally, in order to focus on explicating the epistemic structuralist approach, several important issues (such as the Aristotelian theory of universals) have been bracketed and put aside for present purposes. However these are still important considerations and will need to be considered in a comprehensive evaluation of epistemic structuralism. This thesis does not claim to have defended epistemic structuralism exhaustively.

Epistemic structuralism is proposed as a plausible theory for explaining the nature of mathematical knowledge and providing an account of how we acquire mathematical knowledge, while maintaining a feasible ontology of mathematical entities. It has been argued that this theory strikes a viable balance between the desired elements for a philosophy of mathematics: its ready absorption into prevailing naturalistic epistemic and semantic schemes; consistency with mathematical practice; applicability to different areas of mathematics; and the lack of any requirement for postulation of non-natural powers or entities. As such it has a distinct advantage, since many theories have strengths in some of these areas but are found wanting in others. Indeed, it is acknowledged that the ontology of mathematical entities that arises from

epistemic structuralism is rather cumbersome. However it is argued that this slight ungainliness is outweighed by its advantage in providing an epistemically driven account of the nature of our knowledge of mathematics.

BIBLIOGRAPHY

- Akiba, Ken [2000], "Indefiniteness of Mathematical Objects," *Principia Mathematica*, **8**, 26-46.
- Appel, Kenneth & Haken, Wolfgang in conjunction with Koch, John [1977], "Every Planar Map Is Four Colorable," *Illinois Journal of Mathematics*, **21**, 429-567.
- Armstrong, D.M. [1992], "Properties" reprinted in Mellor & Oliver [1997] 160-172.
- Azzouni, Jody [1994], *Metaphysical Myths, Mathematical Practice: The Ontology and Epistemology of the Exact Sciences* (Cambridge, New York: Cambridge University Press).
- Bacon, John [1995], *Universals and Property Instances: the Alphabet of Being* (Oxford, New York: Blackwell).
- Benacerraf, Paul & Putnam, Hilary (ed.) [1964], *Philosophy of Mathematics*, (Englewood Cliffs, New Jersey: Prentice-Hall).
- Benacerraf, Paul & Putnam, Hilary (ed.) [1983], *Philosophy of Mathematics*, Second Edition (Cambridge: Cambridge University Press).
- Benacerraf, Paul [1965], "What Numbers Could Not Be," reprinted in Benacerraf & Putnam [1983] 272-294.
- Benacerraf, Paul [1973], "Mathematical Truth," *Journal of Philosophy*, LXX, 19, 661-679.
- Bernays, Paul [1935], "On Platonism in Mathematics," reprinted in Benacerraf & Putnam [1983] 258-271.
- Bigelow, John [1988], *The Reality of Numbers* (Oxford: Clarendon Press).
- Black, Max [1971], "The Elusiveness of Sets," *Review of Metaphysics*, **24**, 614-636.
- BonJour, Laurence [1985], *The Structure of Empirical Knowledge* (Cambridge, Massachusetts, and London, England: Harvard University Press).
- Boolos, George S. & Jeffrey, Richard C. [1974], *Computability and Logic*, Third Edition (Cambridge: Cambridge University Press, 1989).
- Boolos, George S. [1971], "The Iterative Concept of Set" reprinted in Benacerraf & Putnam [1983] 486-502.

- Boyer, Carl B. [1989], *A History of Mathematics*, Second Edition (New York, Toronto: John Wiley & Sons, Inc.).
- Brown, James Robert [1999], *Philosophy of Mathematics: an Introduction to the World of Proofs and Pictures* (London and New York: Routledge).
- Burge, Tyler [1998], "Computer Proof, Apriori Knowledge, and Other Minds," *Philosophical Perspectives*, **12**, *Language, Mind, and Ontology*, Tomberlin, James E (ed).
- Burgess, John P. & Rosen, Gideon [1997], *A Subject with No Object: Strategies for Nominalistic Interpretation of Mathematics* (Oxford: Clarendon Press).
- Carnap, Rudolf [1931], "The Logicist Foundations of Mathematics," reprinted in Benacerraf & Putnam [1983] 41-52.
- Castañeda, Hector-Neri [1959], "Arithmetic and Reality," in Benacerraf and Putnam [1964] 404-417.
- Casullo, Albert [1988], "Necessity, Certainty, and the A Priori," *Canadian Journal of Philosophy*, **18**, 43-66.
- Collins, H.M. [1992], *Changing Order: Replication and Induction in Scientific Practice* (Chicago and London: University of Chicago Press).
- Colyvan, Mark [2001], *The Indispensability of Mathematics*, (New York: Oxford University Press).
- Detlefsen, Michael [1986], *Hilbert's Program* (Norwell, MA: D. Reidel Publishing Company).
- Duhem, P. [1904-5], *The Aim and Structure of Physical Theory*, Second Edition, Philip P. Wiener (trans.) (New York: Atheneum, 1962).
- Einstein, A. [1954], *Ideas and Opinions* (New York: Crown).
- Field, Hartry H. [1980], *Science Without Numbers* (Oxford: Basil Blackwell).
- Franklin, J. [1994], "The Formal Sciences Discover the Philosopher's Stone," *Studies in History and Philosophy of Science*, **25**, 513-533.
- Frege, G. [1884], *The Foundations of Arithmetic*, Second Revised Edition, trans. J.L. Austin (Oxford: Basil Blackwell, 1978).
- Gähde, Ulrich and Stegmüller, Wolfgang [1986], "An Argument in Favor of the Duhem-Quine Thesis: From the Structuralist Point of View" in Hahn & Schilpp [1986] 117-136.

- Gandy, Robin [1973], “‘Structure’ in Mathematics,” in Robey [1973] 138-153.
- Gasking, D.A.T. [1940], “Mathematics and the World,” reprinted in Benacerraf and Putnam [1964] 390-403.
- Gillies, Donald [1993], *Philosophy of Science in the Twentieth Century: Four Central Themes* (Oxford and Cambridge: Blackwell).
- Gödel, Kurt [1947], “What is Cantor’s continuum problem?” reprinted in Benacerraf & Putnam [1983] 470-485.
- Goldman, Alvin [1973], “A Causal Theory of Knowledge,” *Journal of Philosophy*, **70**, 357-372.
- Haack, Susan [1978], *Philosophy of Logics* (Cambridge: Cambridge University Press).
- Hahn, Lewis Edwin and Schilpp, Paul Arthur (ed.) [1986], *The Philosophy of W.V. Quine* (La Salle, Illinois: Open Court).
- Hand, Michael [1989], “A Number is the Exponent of an Operation,” *Synthese*, **81**, 243 - 265.
- Hand, Michael [1993], “Mathematical Structuralism and the Third Man,” *Canadian Journal of Philosophy*, **23**, 179-192.
- Hardy, G.H. [1978], *Ramanujan: Twelve Lectures on Subjects Suggested by his Life and Work*, Third Edition (New York, N.Y.: Chelsea Publishing Company).
- Harsanyi, J.C. [1983], “Mathematics, the Empirical Facts and Logical Necessity,” *Erkenntnis*, **19**, 167-192.
- Hart, W.D. (ed.) [1996], *The Philosophy of Mathematics* (Oxford: Oxford University Press).
- Hellman, Geoffrey [1989], *Mathematics without Numbers: Towards a Modal-Structural Interpretation* (Oxford: Clarendon Press).
- Hellman, Geoffrey [1990], *Modal-Structural Mathematics*, in Irvine (ed.) [1990] 307-330.
- Hellman, Geoffrey [2001], “Three Varieties of Mathematical Structuralism,” *Philosophia Mathematica*, **9**, 184-211.
- Hilbert, David [1925], “On the infinite,” reprinted in Benacerraf & Putnam [1983] 41-52.

- Hinckfuss, Ian [1993], "Suppositions, Presuppositions, and Ontology," *Canadian Journal of Philosophy*, **23**, 595-618.
- Irvine, A.D. (ed.) [1990], *Physicalism in Mathematics* (Boston: Kluwer Academic Publishers).
- Irvine, A.D. [1989], "Epistemic Logicism & Russell's Regressive Method," *Philosophical Studies*, **55**, 303-327.
- Irvine, A.D. [1990], "Nominalism, Realism & Physicalism in Mathematics" in Irvine (ed.) [1990] ix-xxvi.
- Islam, Amitavo [1996], "Mathematics as Science" (University of Sydney), by personal correspondence.
- Jané, Ignacio [2001], "Reflections on Skolem's Relativity of Set-Theoretical Concepts," *Principia Mathematica*, **9**, 129-153.
- Kanigel, Robert [1991], *The Man Who Knew Infinity: A Life of the Genius Ramanujan* (Toronto: Collier Macmillan Canada).
- Keränen, Jukka [2001], "The Identity Problem for Realist Structuralism," *Philosophia Mathematica*, **9**, 308-330.
- Kitcher, Philip [1983], *The Nature of Mathematical Knowledge* (New York, Oxford: Oxford University Press).
- Kleene, Stephen Cole [1959], *Introduction to Metamathematics*, (Amsterdam: North-Holland Publishing Co.).
- Körner, Stephan [1960], *Philosophy of Mathematics* (London: Huchinson & Co).
- Kouremenos, Theokritos [1995], *Aristotle on Mathematical Infinity* (Stuttgart: Franz Steiner Verlag).
- Kripke, Saul [1977], "Speaker's Reference and Semantic Reference," *Midwest Studies in Philosophy*, **2**, 255-276.
- Kripke, Saul [1980], *Naming and Necessity* (Cambridge, Massachusetts: Harvard University Press).
- Lakatos, Imre [1976], *Proofs and Refutations: The Logic of Mathematical Discovery* (Cambridge, New York: Cambridge University Press, 1994).
- Lakatos, Imre [1978], "A Renaissance of Empiricism in the Recent Philosophy of Mathematics?" reprinted in Tymoczko [1998] 29-48.

- Lehman, Hugh [1979], *Introduction to the Philosophy of Mathematics* (Oxford: Basil Blackwell).
- Leibniz, G.W. [1765], *New Essays on Human Understanding*, Abridged edition, trans. and ed. P. Remnant & J. Bennett (Cambridge: Cambridge University Press, 1982).
- Lenzen, Victor F. [1938], "Procedures of Empirical Science," *International Encyclopedia of Unified Science* **1** (Chicago, Illinois: University of Chicago Press).
- Lewis, David [1973], *Counterfactuals* (Oxford: Basil Blackwell).
- Maddy, Penelope [1980], "Perception and Mathematical Intuition" reprinted in Hart [1997] 114-141.
- Maddy, Penelope [1984], "Mathematical Epistemology: What is the Question?," *The Monist*, **67**, 46-55.
- McGee, Vann [1997], "How We Learn Mathematical Language," *The Philosophical Review*, **106**, 35-68.
- Mellor, D.H. & Oliver, Alex (ed.) [1997], *Properties* (Oxford: Oxford University Press).
- Mill, John Stuart [1872], *A System of Logic*, ed. J.M. Robson (Toronto and Buffalo: University of Toronto Press, 1973).
- Morton, Adam & Stich, Stephen P. (ed.) [1996], *Benacerraf and his Critics* (Oxford: Blackwell Publishers).
- Oddie, Graham [1990], "Verisimilitude by Power Relations," *The British Journal for the Philosophy of Science*, **41**, 129-135.
- Parsons, Charles [1977], "What is the Iterative Conception of Set?" reprinted in Benacerraf & Putnam [1983] 503-529.
- Parsons, Charles [1990], "The Structuralist View of Mathematical Objects," *Synthese*, **84**, 303-346.
- Potter, M.D. [1993], "Iterative Set Theory," *The Philosophical Quarterly*, **43**, 178-193.
- Putnam, Hilary [1975], "What is Mathematical Truth?" reprinted in Tymoczko [1998] 49-65.
- Putnam, Hilary [1983], *Philosophical Papers, Volume 3: Realism and Reason* (Cambridge and New York: Cambridge University Press).

- Quine, Willard Van Orman [1948], "On What There Is" in Quine [1953] 1-19.
- Quine, Willard Van Orman [1951], "Two Dogmas of Empiricism" in Quine [1953] 20-46.
- Quine, Willard Van Orman [1953], *From a Logical Point of View*, Second Edition (New York and Evanston: Harper and Row, 1961).
- Quine, Willard Van Orman [1981], *Theories and Things* (Cambridge, Massachusetts and London, England: The Belknap Press of Harvard University Press).
- Quine, Willard Van Orman [1990], *Pursuit of Truth* (Cambridge, Massachusetts and London, England: Harvard University Press).
- Resnik, Michael D. [1975], "Mathematical Knowledge and Pattern Cognition," *Canadian Journal of Philosophy*, **5**, 25-39.
- Resnik, Michael D. [1997], *Mathematics as a Science of Patterns* (Oxford: Clarendon Press).
- Robey, David (ed.) [1973], *Structuralism: an Introduction* (Oxford: Clarendon Press).
- Rosenkrantz, R.D. [1980], "Measuring Truthlikeness," *Synthese*, **45**, 463-487.
- Russell, Bertrand [1907], "The Regressive Method of Discovering the Premises of Mathematics," reprinted in Russell [1973] 272-283.
- Russell, Bertrand [1912], *The Problems of Philosophy* (London: Oxford University Press, 1964).
- Russell, Bertrand [1973], *Essays in Analysis*, Douglas Lackey (ed.) (London: George Allen & Unwin).
- Shapiro, Stewart [1989], "Structure and Ontology," *Philosophical Topics*, **17**, 145-171.
- Shapiro, Stewart [1997], *Philosophy of Mathematics: Structure and Ontology* (New York: Oxford University Press).
- Shapiro, Stewart [2000], *Thinking about Mathematics: The Philosophy of Mathematics*, (New York: Oxford University Press).
- Steiner, Mark [1975], *Mathematical Knowledge* (Cornell: Cornell University Press).
- Steiner, Mark [1989], "The Application of Mathematics to Natural Science," *Journal of Philosophy*, **86**, 449-480.
- Steiner, Mark [1983], "Mathematical Realism," *Noûs*, **17**, 363-385.

- Steiner, Mark [1998], *The Applicability of Mathematics as a Philosophical Problem* (Cambridge, Massachusetts: Harvard University Press).
- Stenius, Erik [1974], "Sets," *Synthese*, **27**, 161-188.
- Thorne, Kip S. [1994], *Black Holes and Time Warps: Einstein's Outrageous Legacy* (New York, New York: W. W. Norton & Company, Inc.).
- Tymoczko, Thomas (ed.) [1998], *New Directions in the Philosophy of Mathematics* (Princeton, New Jersey: Princeton University Press).
- Tymoczko, Thomas [1979], "The Four-Color Problem and its Philosophical Significance," *Journal of Philosophy*, **76**, 57-83.
- van Dalen, Dirk and Ebbinghaus, Heinz-Dieter [2000], "Zermelo and the Skolem Paradox," *The Bulletin of Symbolic Logic*, **6**, 145-161.
- van Dantzig, D. [1955], "Is 10^{10} a Finite Number?," *Dialectica*, **9**, 273-277.
- Wetzel, Linda [1989], "That Numbers Could be Objects," *Philosophical Studies*, **56**, 273-292.
- Whitehead, Alfred North & Russell, Bertrand [1913], *Principia Mathematica* (Cambridge, Cambridge University Press), Volume II Second Edition, 1959.
- Wilson, Jessica M. [2000], "Could Experience Disconfirm the Propositions of Arithmetic?," *Canadian Journal of Philosophy*, **30**, 55-84.
- Woods, John, Irvine, Andrew and Walton, Douglas [2000], *Argument: Critical Thinking, Logic and the Fallacies* (Toronto: Prentice-Hall).