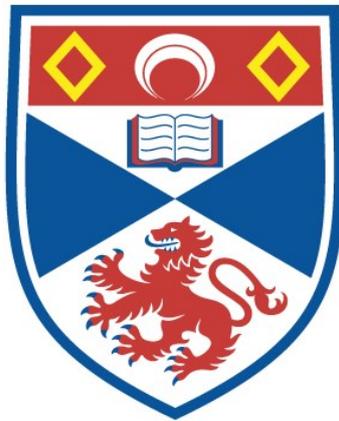


**PROOF, RIGOUR & INFORMALITY:
A VIRTUE ACCOUNT OF MATHEMATICAL KNOWLEDGE**

Fenner Stanley Tanswell

**A Thesis Submitted for the Degree of PhD
at the
University of St Andrews**



2017


**Full metadata for this item is available in
St Andrews Research Repository
at:
<http://research-repository.st-andrews.ac.uk/>**

**Please use this identifier to cite or link to this item:
<http://hdl.handle.net/10023/10249>**

This item is protected by original copyright

PROOF, RIGOUR & INFORMALITY:
A VIRTUE ACCOUNT OF MATHEMATICAL
KNOWLEDGE

Fenner Stanley Tanswell



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of PhD
at the
University of St Andrews

Date of Submission
31st August 2016

Declaration

1. Candidate's declarations:

I, Fenner Stanley Tanswell, hereby certify that this thesis, which is approximately 65000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September 2013 and as a candidate for the degree of Doctor of Philosophy in September 2013; the higher study for which this is a record was carried out in the University of St Andrews between 2013 and 2016.

Date signature of candidate

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date signature of supervisor

Date signature of supervisor

3. Permission for publication:

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or

research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis: No embargo on print copy. No embargo on electronic copy.

Date signature of candidate

 signature of supervisor

 signature of supervisor

Abstract

This thesis is about the nature of proofs in mathematics as it is practiced, contrasting the *informal proofs* found in practice with formal proofs in formal systems. In the first chapter I present a new argument against the Formalist-Reductionist view that informal proofs are justified as rigorous and correct by corresponding to formal counterparts. The second chapter builds on this to reject arguments from Gödel's paradox and incompleteness theorems to the claim that mathematics is inherently inconsistent, basing my objections on the complexities of the process of formalisation. Chapter 3 looks into the relationship between proofs and the development of the mathematical concepts that feature in them. I deploy Waismann's notion of *open texture* in the case of mathematical concepts, and discuss both Lakatos and Kneebone's dialectical philosophies of mathematics. I then argue that we can apply work from *conceptual engineering* to the relationship between formal and informal mathematics. The fourth chapter argues for the importance of *mathematical knowledge-how* and emphasises the primary role of the *activity of proving* in securing mathematical knowledge. In the final chapter I develop an account of mathematical knowledge based on *virtue epistemology*, which I argue provides a better view of proofs and mathematical rigour.

Acknowledgements

I am very grateful for the extensive support, encouragement, feedback and friendship I've benefitted from in the last three years in producing this thesis. First of all, I would like to thank the family of Caroline Elder for the generosity in providing the Caroline Elder PG scholarship in her memory. Thanks also for the further SASP scholarship I received, as well as travel funding from the Indo-European Research Training Network in Logic.

Thanks to my supervisors Aaron Cotnoir and Patrick Greenough, who have both been fantastic throughout and have had a huge influence on my thinking and writing in this thesis. Their comments and feedback have been invaluable and led to massive improvements in the work presented here.

My particular interest in proofs and the philosophy of mathematical practice started under the supervision of Benedikt Löwe while doing the Master of Logic in Amsterdam. In particular, the first two chapters were strongly influenced by Benedikt and the work I did there for my master's thesis, so many thanks.

I have also had direct feedback on previous drafts from Noah Friedman-Biglin, Bruno Jacinto, Josh Habgood-Coote, Ásgeir Berg Matthíasson, Alex Yates, Caroline Torpe Touborg, Mark Bowker, Ryo Ito, Patrik Hummel, Brian King, Claire Field, Morgan Thomas, Chris Sangwin and Colin Ritberg. Many thanks to all.

I would also like to thank Jody Azzouni who was a great help for chapters 1 and 5, and who provided excellent comments and discussion on the overgeneration problem which led to a more robust objection. Thanks also to Robert Thomas, Holger Andreas and Peter Verdée, who provided useful feedback on chapters 1 and 2 towards publication. Thanks too to three anonymous referees in writing reports on earlier versions of the first two chapters.

Thanks to Stephen Read for extra guidance and discussion on the philosophy of mathematics, as well as for running the primary seminars I have been

involved with in Arché. Thanks also to Lynn Hynd for the organisational brilliance in keeping Arché running, and to the other staff who ran the philosophy department over the last three years: Laura, Shaun, Emmy, Lucy, Rhona and Katie.

The work in chapter 4 formed part of the content of a joint Prize Seminar for St Andrews undergraduates with Josh Habgood-Coote, so thanks to him for the huge amount of feedback and help on knowing-how. Thanks also to the students who attended, asked insightful questions and pushed us on many philosophical points.

The work in chapter 5 was presented alongside Colin Rittberg's work on theoretical virtues for set theory theory as a three-lecture series in Aberdeen. Thanks to Colin for extensive discussion in the run up to these and for inviting me to Brussels twice. Anti-thanks to ISIS who tried to stop us.

Thanks to audiences at the many talks I've given in St Andrews and Stirling, as well as audiences and organisers in Umeå, Bristol, Munich, Cambridge, Storrs, New Delhi, Paris, Brussels and Aberdeen. Thanks also to the wider St Andrews community, members of the philosophy department and Arché, for being great people.

I spent two months in 2014 at the University of Connecticut, with gratitude to the Arché Travel Fund which paid for it. I would like to thank those who I interacted with while there and who made me incredibly welcome: Jc Beall, Dave Ripley, Morgan Thomas, Nathan Kellen, Hanna Gunn, Andrew Parisi, Damir Dzhafarov, Reed Solomon, Brown Westrick, Daniel Silvermint, Keith Simmons, Nate Sheff and Suzy Killmister. I am also really thankful to the many people who welcomed me into their community in Coventry CT, with special thanks to Holly Parker and Joram Echeles.

Thanks to Aikman's and Rendez-Vous for the good living.

I am surrounded by friends who have made doing a PhD possible. In addition to those mentioned above, many thanks to Joe Slater, Laurence Carrick, Mat Chantry, Ally Holmes, Calum McConnell, Carley Hollis, Elspeth Gillespie, Alex Purser, Beth Allison, Bella Bandell, Ste Broadrick, Paul Black, Alper Yavuz, Marissa Wallin, Gabriela Rino Nesin, Maja Jaakson, Aadil Kurji, Nathaniel Forde, Jessica Olsen, Hraban Luyat, Ravi Thakral, Ryan Nefdt and the many others I have failed to mention. You're all top bananas.

Most of all, I am indebted to my family, in particular mum & Chris, dad and Thecla: thank you all. Finally, thanks to Lea who has given me confidence, patience, love and support over the last three years.

Publication

A version of chapter one has appeared in *Philosophia Mathematica* in 2015 with the title “A Problem with the Dependence of Formal Proofs on Informal Proofs”. This is in the bibliography under (Tanswell 2015).

A version of chapter two is forthcoming in the collection *Logical Studies of Paraconsistent Reasoning in Science and Mathematics*, edited by Holger Andreas & Peter Verdée as part of Springer’s *Trends in Logic* series, with the title “Saving Proof from Paradox: Gödel’s Paradox and the Inconsistency of Informal Mathematics”. This is in the bibliography under (Tanswell forthcoming).

Contents

Introduction	9
0.1 The Twin Faces of Orthodoxy	10
0.2 The Philosophy of Mathematical Practice	14
0.3 Outline	16
1 A Problem with the Dependence of Informal Proofs on Formal Proofs	20
1.1 Introduction	20
1.2 Minimal Desiderata	23
1.3 Azzouni’s Derivation-Indicator View	25
1.4 A Dilemma	28
1.5 Agent-Independent Derivation-Indicators	29
1.6 Agent-Dependent Derivation-Indicators	37
1.7 Conclusion	39
2 Saving Proof from Paradox	41
2.1 Introduction	41
2.2 Formal and Informal Proofs	43
2.3 Gödel’s Paradox and Beall’s Argument	45
2.4 Priest’s Argument	48
2.5 Formalising Mathematics	50
2.5.1 A First Option	50
2.5.2 A Second Option	52
2.6 On Mathematical Super-Theories	53
2.7 Fragmented Formalisations	56
2.8 On The Formal and The Informal	57
3 Mathematical Concepts	59
3.1 Introduction	59
3.2 Waismann and Open Texture	61
3.3 Lakatos and <i>Proofs & Refutations</i>	67
3.3.1 A Briefing on the Proof and Responses to Counterexamples	68

3.3.2	Concepts and a Dialectical Philosophy of Mathematics	74
3.3.3	Fallibilism and Formality	78
3.4	Kneebone on Mathematics	84
3.5	The Open Texture of Mathematical Concepts	88
3.6	Conceptual Engineering in Mathematics	94
3.7	Haslanger's Manifest and Operative Concepts	96
3.8	Scharp on Replacing Defective Concepts	103
3.9	The Concepts of Mathematics	108
3.10	Conclusion	112
4	Mathematical Know-How	113
4.1	Introduction	113
4.2	Rav and <i>Pythiagora</i> the Oracular Computer	115
4.3	Knowing-How and Knowing-That	121
4.4	Mathematical Know-How	133
4.5	Löwe & Müller on Mathematical Skills	138
4.6	Proving in Action	144
4.6.1	Larvor's Inferential Actions	144
4.6.2	Proving as an Activity	149
4.7	Conclusion: Knowing How to Prove It	153
5	A Virtue Approach to Mathematical Epistemology	157
5.1	Introduction	157
5.2	The Moderate Proposal	160
5.3	Virtue Epistemology	162
5.3.1	Virtue Reliabilism	163
5.3.2	Virtue Responsibilism	164
5.3.3	Hybrid Virtue Approaches	166
5.3.4	Epistemic Vices	167
5.4	The Strong Proposal	169
5.5	Virtues and Proving	174
5.6	Rigour	179
5.7	Mathematical Understanding	182
5.8	Case Study: Mochizuki and the <i>abc</i> Conjecture	187
5.9	Conclusion	193
	Conclusion	196
	Bibliography	199

Introduction

Some will think that a mathematical argument either is a proof or is not a proof. In the context of elementary analysis I disagree, and believe instead that the proper role of a proof is to carry reasonable conviction to one's intended audience. It seems to me that mathematical rigor is like clothing: in its style it ought to suit the occasion, and it diminishes comfort and restricts freedom of movement if it is either too loose or too tight.

— George F. Simmons (Simmons 1991, p. xi)

Proofs have many roles and functions in mathematics and beyond. They are one of the central and long-standing ways to establish mathematical truths and eradicate any rational doubt. They can also serve to convince a stubborn-but-fair audience and to explain why some mathematical statement is true. They contain the methods, techniques and know-how of mathematics. They make explicit the logical dependencies of mathematical propositions. They can educate students to understand particular areas of mathematics and demonstrate acceptable inferential actions at work. They can be used to test the boundaries of the concepts they use and point to improvements that might be made. They can step beyond those boundaries and lead us into paradox. They can tell an engaging mathematical story and be elegant and beautiful. Finally, they make for a good topic for three years of philosophy PhD research.

This thesis is about *proofs* as found in all mathematical settings, from napkin scribbles to classrooms, and from journal articles to computer-coded derivations. Naturally, proofs in these different settings vary a great deal

in how they are presented, and this is something which will be of major importance throughout the thesis. In particular, there is a difference to be made between the idealised notion of proof, where each step is explicit and endorsed by some underlying formal system, and the proofs that mathematicians generally produce in practice. Following the literature, I will therefore distinguish between these as *formal proofs* and *informal proofs*. I will save an exploration of this distinction for the start of the first chapter.

In the coming five chapters I will be looking at many of the functions of proofs in turn and in relation to one another. Each of the functions for proofs will have associated guiding philosophical questions we will want answers to. Amongst these, though, three topics come to the fore as those which form the central problems of the thesis and relate to many of the above roles for proofs. First of all, there is the question of the *formality* or *informality* of proofs and how this affects whether the proofs are correct and rigorous. Secondly, I shall consider the question of what exactly *mathematical rigour* is. Thirdly, the question of the role of proofs and proving in our possession of *mathematical knowledge* and in obtaining that knowledge.

The answers that I shall be giving to these questions will stand in contrast to two strands of orthodox thinking in the philosophy of mathematics, which will be worth setting out before we begin with the main body of the thesis.

0.1 The Twin Faces of Orthodoxy

In (Kitcher & Aspray 1988), Kitcher & Aspray offer an “opinionated introduction” to the philosophy of mathematics, dividing up contemporary thinking into two streams of thought: the mainstream tradition and the maverick tradition. The first of these is taken to be following up on the Fregean quest for foundations, and is characterised in (Ernest 1997) as best represented by the work found in the edited volume (Benacerraf & Putnam 1983) and will be discussed now. Kitcher & Aspray take the second stream to originate with Lakatos, and has since broadly grown into the *philosophy of mathematical practice*. We will return to that shortly, and Lakatos in detail in chapter 3.

I want to actually separate out two general families of views captured under the heading of the mainstream, orthodox tradition for philosophy of

mathematics. The first set of views, which I shall simply call the *Traditionalist* approach, concerns the kind of questions we are interested in about mathematical knowledge, focusing on the a priori and privileging knowledge of mathematical truths. The second view is the associated take on mathematical proofs, which I shall refer to as the *Formalist-Reductionist* view, which wants proofs to primarily be understood in terms of formal derivations. This latter umbrella term is taken from (Antonutti Marfori 2010, p. 262) and is intended to capture the focus on formal proofs of the approach while avoiding conflation with the formalist position in the philosophy of mathematics. Let us consider these twin faces of orthodoxy in turn.

The first implicit trend in the philosophy of mathematics, which I have termed the ‘Traditionalist view’, is about the kind of questions that philosophy of mathematics is interested in, in particular in relation to mathematical knowledge. Primarily, these seem to be about the *a priori* nature of mathematical knowledge, what enables us to possess it generally and what enables us to have *access* to the objects of mathematics to allow for such knowledge to begin with. There are great historical precedents for these questions, as these seem to be the kind of questions which concern almost all major figures in the historical consideration of the nature of mathematics, from Plato to Descartes, through to Kant and beyond. As a result of the emphasis on the ontology of mathematics and the problem of how we access it in order to have mathematical knowledge, the Traditionalist tends to mainly focus on mathematical truths, or how it is we can come to know true mathematical facts and what they consist in. Such a prioritisation of truths tends to come at the expense of systematic discussion of the nature of proof, which is downplayed as explainable via the Formalist-Reductionist view which we will turn to shortly. The Traditionalist picture generally tries to account for mathematical truth then offer a subsidiary picture of proof as delivering these truths unto us. It is this trend of Traditionalism which is the target of Rav’s argument that it is proofs rather than truths which are the primary bearers of mathematical knowledge found in (Rav 1999) which I shall be looking at in chapter 4.

Now, if you were to press someone who fell broadly into the Traditionalist camp on the importance of proof, the likely response would be to admit that there may be issues pertaining to proofs, such as how they deliver a priori or deductive knowledge. Indeed, the Traditional picture is that proofs provide

justification for the truth of the theorems proved, and that they do so in a special way which elevates mathematical knowledge obtained via proofs to some especially robust form of knowledge. Such robustness is normally taken to be the result of the deductive nature of mathematical proofs, where giving a deductive proof rules out the possibility of error or omission.

The problem comes when we try to apply this Leibnizian ideal of proof to proofs as found in practice. Real proofs involve general proof techniques and strategies which may not be easily reducible to anything formal; they may suppress or omit many steps that would be found in a strict derivation; they assume varying degrees of knowledge on the part of reader; they can use diagrams and natural language explanations; and they may be more aimed at understanding why a claim is true than filling in trivial details. The gap between the two notions here is often defined in terms of *formality*; that the ideal of proving consists in formal proofs in formal systems, while those found in practice are called informal proofs. The question then is how these informal proofs, which constitute the majority of proofs given in mathematics, meet the standards of rigour expected of mathematics as the science of deduction.

Our second face of the orthodox view, the Formalist-Reductionist view, takes a strong stance on the issues of rigour and formality: they take it that good, correct proofs—that is, proofs which are sufficient to secure the special kind of mathematical knowledge— just are formal proofs in formal systems. However, more needs to be said to avoid the following troublesome triad:

1. Mathematicians have mathematical knowledge gained from correct proofs.
2. Correct proofs just are formal proofs.
3. Mathematicians only rarely know or encounter any formal proofs.

The first is taken for granted, since proofs are central to the methodology of mathematics, while the third is a fact of mathematical practice which is easily observed and corroborated. The second item is the Formalist-Reductionist assumption, which is meant to complement the Traditionalist view by answering how proofs deliver a priori and deductive knowledge. Yet, taken together there is the clear difficulty that formal proofs are meant to be

the main source of the special kind of knowledge, while simultaneously not being accessed by the mathematicians who have the mathematical knowledge via proofs.

The extended answer of the Formalist-Reductionist view is that informal proofs can be justified as good, correct or rigorous via *corresponding* formal proofs. Thus, the initial restriction of correct proofs only including formal ones is relaxed to allow some of the informal proofs to be correct too, so long as they are related to formal correspondents in the right sort of way. There are many ways this can be filled out into a full account: the connection between informal proof and formal proofs could be abbreviation/approximation (Mac Lane 1986, p. 377), formalisation, derivation-indication (Azzouni 2004a), Carnapian explication (Sjögren 2010)¹, picking out a logical form, using the informal proof to convince the reader of the existence of a formal proof “in the right way” (Burgess 2015) or even that the informal proof needs to be sufficient for a hypothetical ‘midwife’ logician to generate the formal proof (Steiner 1975). While the details of these accounts may vary dramatically, the key thrust of the approach will be to save the informal proofs of mathematical practice as properly deductive and rigorous through their formalisations to formal proofs. Formal proofs have all of the desirable properties, the idea goes, since by squeezing out all gaps we leave no room for mistakes or omissions, so proofs are fully rigorous according to our settled logical rules of inference and deductive from the explicitly identified axioms and granted premises. Informal proofs are somehow matched up to these ideal, logical, formal proofs and can thus inherit, via what Rav calls ‘Hilbert’s Bridge’ (Rav 1999, p. 12), the same justification, rigour and epistemic status.

So how do our two characters, the Traditionalist and the Formalist-Reductionist, relate to one another? Certainly nothing so strong as one entailing the other, since one can have Traditionalist pre-occupations while rejecting the formalistic conception of proofs, while one can believe that only formal proofs are correct proofs while being markedly unconventional in your other philoso-mathematical investigations. Nonetheless, the two positions do interrelate and provide mutual support. For one thing, we saw that the Traditionalist needs a story to tell which fits with the a priori and special nature of mathematical knowledge. The Formalist-Reductionist

¹Carnap’s notion of explication will be discussed in chapter 3.

picture at least seems to fit this role without obviously upsetting any of the requirements the Traditionalist would have for such a story. In fact, the Formalist-Reductionist view is often the result of Traditionalist motivations towards a special degree of infallibility that we want from mathematical knowledge, translating to proofs needing to be formal in order to rule out errors, illicit assumptions or gaps of reasoning. While the Traditionalist position is deeply ingrained in the history of the philosophy of mathematics, Formalist-Reductionist thinking is a much more recent trend, focusing as it does on the relatively recent notion of formal proofs, seeing formal proofs as the ultimate solution to these ongoing questions of rigour and fallibility.

My answers will go strongly against the Formalist-Reductionist positions, arguing against it explicitly in the first two chapters. The Traditionalist is a less concrete opponent, and I don't think that all of what has come before under this heading is flawed. Nonetheless, I shall be looking at a broader class of questions about mathematics, following the approach of the *philosophy of mathematical practice*. Let us introduce this now.

0.2 The Philosophy of Mathematical Practice

The second stream of thought in modern philosophy of mathematics, according to Kitcher & Aspray, is the “Maverick” tradition following the work of Lakatos.² A great number of thoughts come out of the Lakatosian approach, which will be covered in some detail in chapter 3. But while it is fun to see oneself as a maverick, the rough amalgamation of researchers favouring this approach is not generally ostracised nowadays and the distinction between the two is not always so clear³, so we may call the approach the study of *mathematical practice*, with the philosophical wing being the *philosophy of mathematical practice*.⁴ For now, I will set out some of the main works that fall under this heading, though for a fuller and more fine-grained discussion of the history of the philosophy of mathematical practice see (Van Bendegem

²Lakatos starts this tradition with the famous phrase “You can be my wingman anytime”. I was unable to trace this quotation, but see (Güz & Eismann 1986) and (Lakatos 1976, fn. 2, p. 62).

³For example, a glance at the membership list for the *Association for the Philosophy of Mathematical Practice* will reveal a great many traditional philosophers of mathematics.

⁴In (Löwe 2016), Löwe discusses to what extent this does form a unified school of thought or community, as well as the place of philosophy amongst the various disciplines taking an interest in mathematical practice.

2014).

Naturally, there were forerunners to the mathematical practice approach that Lakatos is seen as the early champion of. Later, in chapter 3, I will discuss G. T. Kneebone who did pre-empt some of the particulars of the Lakatosian approach but seems to have been widely overlooked. However, there are also others who are well-known precursors such as George Pólya, who wrote about heuristics and mathematical reasoning, such as in (Polya 1945), and who is acknowledged as one of the main inspirations for Lakatos alongside Hegel and Popper. Another work that goes in this direction is Hardy's *A Mathematician's Apology* (Hardy 1929), which presents mathematics from a mathematician's point of view. Other major works include Davis & Hersh's *The Mathematical Experience* (Davis & Hersh 1981), Corfield's *Towards a Philosophy of Real Mathematics* (Corfield 2003) and the collected volume (Mancosu 2008). We may consider Maddy's *second philosophy* project as a similar practice-oriented approach (see Maddy 2007). One notable absence from the list is the later Wittgenstein, who seems to have had very little direct influence on the literature, with the exception of (Avigad 2008).

The methodological issue that arises is how exactly to engage with mathematical practice. The researchers one finds at connected conferences span a diverse set of fields, which might provide a similar span of answers drawing on philosophy, history, sociology, maths education, psychology and ethnography. For one thing, we can rely on reports from practicing mathematicians and their reflections on what mathematics is, such as those of Hardy, Maddy, William Thurston (Thurston 1994), Yehuda Rav (Rav 1999) or Tim Gowers (Gowers 2006). The careful use of case studies and historical episodes is also common, mirroring their general use in the philosophy of mathematics, examples found in (Van Kerkhove & Van Bendegem 2008), (Muntersbjorn 2003) and (Schlimm 2011). Sociological studies can be carried out, such as is done in (MacKenzie 2001). Empirical data collection can also shed light on what is going on in mathematics, as is done in (Martin & Pease 2013), (Inglis & Aberdein 2015), (Mejia-Ramos & Weber 2014) and (Geist, Löwe & Van Kerkhove 2010). All of these combine to paint a rich picture of what mathematics looks like in practice, and one which provides a stark contrast with the traditional picture found in the philosophy of mathematics.

Proofs in mathematics form one of the central topics in the study of

mathematical practices, and the opposition to the Formalist-Reductionist approach to proofs a unifying theme. Lakatos argues this point strongly throughout (Lakatos 1976), as do Davis & Hersh (Davis & Hersh 1981), Thurston (Thurston 1994) and Rav (Rav 1999, 2007). Additionally, further arguments are made in (Robinson 1991, 2005), (Hersh 1993, 1997), (Detlefsen 2008), (Antonutti Marfori 2010), (Larvor 2012, 2016), and (de Toffoli & Giardino 2014), amongst others. It is this literature to which the arguments of the first two chapters can be added.

The approach I will be taking to proofs in practice in the coming five chapters will be a philosophical one, but drawing on quite a range of literature on mathematical practices and several case studies and examples. The intention is to demonstrate that there is a great deal of relevant work from philosophy that can be deployed in the debate about the nature of proofs, in rejecting the Formalist-Reductionist picture, and in replacing it with a better account which does accord well with practice.

0.3 Outline

Let us run through what is to come in the five chapters in turn.

Chapter 1 takes on one prominent version of the Formalist-Reductionist position, Jody Azzouni's derivation-indicator account. I give a new argument that this cannot work. The target for the argument is the correspondence supposedly holding between informal proofs and their formal counterparts. I present the dilemma of whether this correspondence is inherent in the proof and independent of any particular agent who is engaging with it, or alternatively whether the correspondence is agent-dependent and may fluctuate between different mathematicians. The latter is a poor fit for the Formalist-Reductionist view, I argue, so they must commit to the first horn of the dilemma. However, this fails due to an overgeneration problem, where a given informal proof might be formalisable to too many substantially different formal proofs. This is problematic because it undermines the prospect of the questions of rigour and correctness being resolved by the dependence on the associated formal proofs, which was the main motivation. We need another set of answers as to why the informal proof only matches to rigorous and correct formal proofs, meaning that the original Formalist-Reductionist answer has not made any progress towards an answer.

In chapter 2, I take on a related set of arguments about the inconsistency of the concept of proof and the inconsistency of the whole of mathematics. The first argument comes from both Jc Beall and Graham Priest and is based on Gödel's Paradox, the sentence which says of itself that it is informally unprovable. They suggest that this shows mathematics to be inherently inconsistent, but I respond that while this may be a paradox, it doesn't have the devastating effects on mathematics they claim. The second argument is a more advanced version of the first and is due to Graham Priest, depending on Gödel's First Incompleteness Theorem. Briefly put, he argues that the theorem proves the existence of a true but unprovable statement in mathematics. But mathematicians agree that all of mathematics can be formalised, so therefore we prove the statement and a contradiction, meaning that the formal system of mathematics is inherently inconsistent. I consider the relationship between formal and informal proofs, arguing that there are multiple reasons that the formalisation move is dubious. First of all, there is the re-emergence of the difficulties from chapter 1, where there may be many different ways to formalise mathematics and no unique right answer. Furthermore, we might grant that mathematics can all be formalised, but even then it requires the extra fact of putting it all into a single formal system, which we have no reason to accept. My suggestion is that alongside the usual two axes of consistency and completeness, we can add a third dimension of formality which explains where arguments such as Priest's go astray.

In chapter 3, I consider the nature of mathematical concepts and how this relates to the formal/informal axis. I begin with a discussion of the notion from Friedrich Waismann of *open texture*, where concepts are not fully delimited for all applications. Next, I go into a detailed investigation of Imre Lakatos's *Proofs and Refutations*, looking at how conceptual development in mathematics is coupled to proving activities. I then show that some key ideas found in Lakatos were also anticipated by G. T. Kneebone's earlier work, especially the advocacy of a dialectical philosophy of mathematics. Bringing these topics together I show that there is open texture to be found in mathematical concepts too. However, spurred by a criticism from David Corfield, I turn to the question of how the sketched approaches to mathematical concept development can account for the modern mathematical usage of formal methods and results. For this, I look at two recent

proposals from conceptual engineering. First, I deploy Sally Haslanger's distinction between manifest and operative concepts to account for the gap between what mathematics looks like in certain areas and how mathematicians reflectively describe it, using the example of set-theoretic foundations. Second, I see how Kevin Scharp's replacement strategy for defective concepts might also be effective in the mathematical realm.

The final two chapters deal with the role of proofs in mathematical epistemology. Chapter 4 is specifically about knowledge-how in mathematics. I begin with a discussion about Yehuda Rav's argument of *Pythiagora* the oracular machine that it is knowledge of proofs which is of primary interest in mathematics. I agree with the general point that knowledge of methods, techniques, concepts, interrelations and strategies are all important, but offer an alternative interpretation of the thought-experiment which puts these alongside more traditional concerns. From here I discuss the knowing-how and knowing-that distinction as found in the epistemology literature, including Gilbert Ryle's original arguments, Jason Stanley & Timothy Williamson's modern revival of the debate, and the more recent position taken by Jennifer Hornsby and David Wiggins, on which knowledge-how and knowledge-that are in practice closely related and interdependent. Following this I run through several examples of how this position is reflected in the realm of mathematical knowledge. I spend some time critically discussing the previous stance taken by Benedikt Löwe and Thomas Müller that mathematical knowledge of some theorem is context-dependent, in that it involves having the contextually-specified skills required to prove it. Finally, I emphasise the perspective of proving as an activity put forward by Brendan Larvor, and argue that this is what needs to be focused on in an epistemological theory for mathematics.

In chapter 5, I set out one way of giving an account of mathematical knowledge using the resources of virtue epistemology. The suggestion will be that both of the main strands of virtue epistemology, reliabilism and responsibilism, can be adopted to give full and rich accounts of mathematical knowledge. This approach is previously unexplored and has advantages in fitting well into the broader picture of knowledge through proving that I have been drawing. Furthermore, I argue that for the responsibilist mathematical rigour can be seen as a virtue possessed by mathematical agents, which gives a new perspective on what rigour amounts to. I finish on a discussion of

the recent controversy surrounding Shinichi Mochizuki's proof of the *abc* conjecture, showing that virtue-theoretic terminology is already being used by mathematicians in relation to mathematics.

Chapter 1

A Problem with the Dependence of Informal Proofs on Formal Proofs

[S]hould a naturalist who had never studied the elephant except by means of the microscope think himself sufficiently acquainted with that animal? [...] The logician cuts up, so to speak, each demonstration into a very great number of elementary operations; when we have examined these operations one after the other and ascertained that each is correct, are we to think we have grasped the real meaning of the demonstration? [...] Evidently not; we shall not yet possess the entire reality; that I know not what which makes the unity of the demonstration will completely elude us.

— Henri Poincaré (Poincaré 1907, p. 21)

1.1 Introduction

We can distinguish two types of proof: *informal proofs* and *formal proofs* (or *proofs* and *derivations*). On the one hand, formal proofs are given an explicit definition in a formal language: proofs in which all steps are either axioms or are obtained from the axioms by the applications of fully-stated inference rules. On the other hand, informal proofs are proofs as they are written and

produced in mathematical practice. They may make assumptions about the intended audience's background knowledge and ability to follow lines of reasoning, skip over tedious or routine steps and make reference to semantic properties and properties of mathematical objects¹ without stating these fully. They also are not confined to formal languages: though mathematical symbolism may be used, natural language, diagrams and mixed-mode explanations are freely employed too.

While formal proofs, in our sense, may be defined mathematically in any number of ways², informal proofs are much harder to pin down precisely. One way to identify the objects of discussion is to give a general description of what they are like, as is done frequently:

[...] what we do to make each other believe our theorems [...] [an] argument which convinces the qualified, skeptical expert. (Hersh 1997, p. 153)

[...] a kind of meaningful narrative [...] more like a story, or even a drama, conveyed to us in language calling on our semantic and intuitive understanding. (Robinson 1991, p. 269)

[...] a conceptual proof of customary mathematical discourse, having an irreducible semantic content [...] (Rav 1999, p. 11)

[...] a sequence of thoughts convincing a sound mind. (Gödel 1953, p. 341)³

A proof of a theorem in mathematics is what we require to convince ourselves and others of the truth of the statement made by the theorem. (Feferman 2012, p. 371)

However, the real problem is not giving such a general description of what informal proofs are like, but it is rather to sort those informal proofs which are correct and rigorous from those which are not.

¹As seen through the 'Plato-tinted spectacles' described in (Buldt, Löwe & Müller 2008).

²Avoiding, for the purposes of this work, the need to fully get to grips with what it means to be formal. For work towards this see (MacFarlane 2000; Dutilh Novaes 2011).

³When I presented this in a talk, Stephen Read pointed out that the sound mind in question probably refers to Gödel's mind.

While we may associate deductive reasoning and logicity with formal proofs in formal systems, actual mathematics is regularly presented informally using informal proofs.⁴ This challenges any proponent of an account of philosophy of mathematics to also give an account of how proving, as it is practised, relates to the idealised notion of formal proofs. There are many routes to take for such an account, from Lakatosian dialectics (see Lakatos 1976), to be discussed in chapter 3, all the way to denying that any mathematics took place before Frege. In this chapter I want to focus on just one family of responses, wherein the rigour and correctness of informal proofs is taken to be dependent (in some sense) on associated formal proofs. As in the introduction, we call this family of views the *Formalist-Reductionist* approach.⁵ There are a number of different connections that informal proofs can be argued to have to their formal counterparts: reductions, logical forms, explications, abbreviations, sketches, formalisations, etc. In this chapter I will look at one particular proposal by Azzouni: that informal proofs indicate underlying formal proofs.⁶

I will begin by laying out some desiderata that any successful Formalist-Reductionist account of informal proofs must meet. I will then explain Azzouni's view of informal proofs, focusing on the particular connection between formal and informal proofs that is posited and how well Azzouni's view would meet the given desiderata. In section 4 I present a dilemma, asking whether the link from informal proofs to underlying formal derivation is an agent-independent one or whether it is dependent on the agent who is presenting the proof. I take Azzouni to need the former in order to be successful in obtaining his brand of formalism-reductionism, but in section 5 I will criticise this horn of the dilemma based on a problem of *overgeneration*. Azzouni can avoid this problem if he adopts the second horn of the dilemma,

⁴What is 'actual mathematics'? The intended answer here is *mathematics as it is practiced* but this is only enlightening in that it points to further questions that need to be addressed, concerning which parts of mathematical practice are relevant. For the purposes of this thesis I take actual mathematics to simply be that published in mathematics journals, presented at conferences and taught in mathematics classes. A number of interesting discussions of this question can be found in (Mancosu 2008).

⁵To emphasise, I specifically avoid calling this simply formalism because the Formalist-Reductionist stance is broader and may encompass positions that would traditionally fall outside of the formalist school of thought. For instance, logicism is guided by the thought that mathematics is part of logic and its approach to proof would fall under the heading of Formalist-Reductionist, while being deeply opposed to formalism on other grounds.

⁶The bulk of this position is given in (Azzouni 2004a) and (Azzouni 2005a).

but in section 6 I will argue that this is not compatible with Azzouni's theory. I shall therefore argue that the account is deficient in dealing with the various desiderata it is aiming to address. Finally, I will conclude that the fundamental difficulty that prevents Azzouni's account from being successful is one that is a general roadblock to successfully providing a Formalist-Reductionist account of informal proofs.

1.2 Minimal Desiderata of a Formalist-Reductionist Account of Informal Proofs

In this section I shall lay down the minimal aims that a Formalist-Reductionist account should achieve in dealing with the problem of informal proofs. By making these intentions clear from the outset, we will be able to see where conflicts arise.

We can begin with two desiderata that were already mentioned:

(Rigour) To give an account of how informal proofs are (or can be said to be) rigorous through their connection to formal proofs.

(Correctness) To distinguish correct informal proofs from incorrect ones i.e. the connection should only link informal proofs that are correct to the formal proofs that justify them.

The first of these is precisely the challenge the Formalist-Reductionist faces in arguing that informal proofs can be rigorous if they are connected to formal proofs in the right kind of way. The second adds to this the need to properly distinguish the correct informal proofs from incorrect ones. One could interpret this as the intention not to overgenerate through the posited connection: it would be undesirable for the link matching informal proofs to formal proofs to also associate *flawed* informal proofs with justifying formal proofs.

Since informal proofs arise from mathematical practice and the way in which we engage with and do mathematics, another desideratum is the following:

(Agreement) To explain how, in practice, mathematicians manage to consistently converge and agree on the correctness of informal proofs. (Additionally, to give an account of informal proofs

that were conceived of long before we had a sufficiently strong account of formal proofs to support them.)

The main part of **(Agreement)** is to actually engage with informal proofs as a social phenomenon; to explain how and why the mathematical community has employed informal proofs, as well as how the underlying link the Formalist-Reductionist account argues for relates to this practice. The addendum presses the requirement further, asking for the account to also explain how the cumulative nature of mathematics fits with the fact that fully formal proofs, of the sort required by the Formalist-Reductionist, are a rather recent discovery. A requirement like this is to avoid the immediate objection that might be raised: that formal proofs cannot underwrite informal ones, because historically we have been using the latter far longer.

Now I will impose a stronger demand on the Formalist-Reductionists, the demand that their account doesn't simply state what the link is between formal and informal proof (abbreviating, indicating, logical form etc.) but that instead it gives some substance to the link.

(Content) To show how the content of an informal proof determines the structure of the formal proof(s) it maps to.

A reason that informal proofs do present a substantial difficulty is that, in many ways, they are and appear quite different to any formal proofs. In answering such a difficulty, then, saying that the relation between them is of a certain kind is the easy part; showing that it is so is much harder. What the account needs to provide is an explanation of how exactly the informal proof can be used to pick out some formal proof or proofs. The picking out must surely (and at least partially) follow the content of the informal proof, so the account needs to tell us about how this content determines the structure of the formal proof that is associated with it.

We may elaborate the above further, to require a response to the particular tricky cases:

(Techniques) To provide an explanation of apparently inherently informal techniques.

A main example of what is required here is dealing with diagrams in mathematics. A legitimate response is to argue for some kind of eliminability

thesis for diagrams: that all diagrams must be eliminable from proofs entirely. Of course, such an argument would need to be given to complete the account, and may bring additional commitments. Other examples are proofs using symmetry, or the ‘untraversed gaps’ described in (Fallis 2003).

1.3 Azzouni’s Derivation-Indicator View

Azzouni’s derivation-indicator view of mathematical practice, as presented in (Azzouni 2004a, 2005a), takes the link between informal and formal proofs to be that informal proofs *indicate* underlying formal proofs.⁷ In his own words:

I take a proof to indicate an ‘underlying’ derivation... Since (a) derivations are (in principle) mechanically checkable, and since (b) the algorithmic systems that codify which rules may be applied to produce derivations in a given system are (implicitly or, often nowadays, explicitly) recognized by mathematicians, it follows that if proofs really are devices mathematicians use to convince one another of one or another mechanically-checkable derivation, this suffices to explain why mathematicians are so good at agreeing with one another on whether some proof convincingly establishes a theorem. (Azzouni 2004a, p. 84)

The focus here is very much on answering (**Agreement**), dealing with the general social conformity regarding good and bad proofs. However, it is clear that for Azzouni this is closely linked to (**Rigour**) and (**Correctness**) in that the link will explain the agreement in terms of informal proofs being correct or rigorous due to underlying formal proofs.

An interesting aspect of Azzouni’s view is that the formal proofs are defined more liberally than usual. He takes them to be located within ‘algorithmic systems’, which are not restricted in the ways we generally take formal proofs to be:

I’ve already stressed that ‘algorithmic systems’ are restricted neither to a particular logic, a particular subject-matter, nor even

⁷Although it should be noted that Azzouni has largely dropped the ‘indicating’ terminology in later developments of the view in (Azzouni 2005b) and (Azzouni 2009) for reasons we will see in section 1.4.

to an explicit language (as opposed to something diagrammatic or pictorial). What is required is that ‘proofs’, however these be understood, are (in principle) mechanically recognizable. (Azzouni 2004a, p. 86)

This has already been criticised (see Rav 2007), with a response from Azzouni in (Azzouni 2009), so I shall not take up this discussion here. However, in the present context the motivation for this view should garner at least some sympathy, for Azzouni is explicitly trying to leave open a straightforward route to meeting the demands of (**Techniques**), in particular those regarding diagrams as used in mathematics. This focus on diagrammatic reasoning becomes clearer if we note Azzouni’s reference to another of his papers analysing diagrammatic reasoning in Euclid’s *Elements* (Azzouni 2004b), suggesting that he believes diagrammatic proofs do not always need to be informal, so long as they are given a mechanically checkable structure.⁸ More on Azzouni’s views of diagrams in mathematics can also be found in (Azzouni 2013).

Turning now to the question of (**Content**) and how exactly it is that derivation-indication links informal proofs to formal ones, Azzouni does not argue that each informal proof is underwritten by some unique formal proof in one algorithmic system. That would, he claims, be implausible as an account of mathematical practice because in reality mathematicians are not held to one specific inference system. Furthermore, if an account did limit mathematicians to one specific formal system it would be open to objections based on incompleteness phenomena.⁹ Instead, in Azzouni’s view each informal proof relates to a family of formal proofs which are located in a number of different algorithmic systems.

It doesn’t much matter where in the family of algorithmic systems we take ‘the’ derivation indicated by a proof to be located... since algorithmic systems embedded in one another are so embedded to conserve derivational results, we can take the derivation indicated to be one located in any algorithmic system within which the result occurs and is surveyable. (Azzouni 2004a, pp.

⁸Understanding formal proofs as mechanically checkable ones takes one of the stances on the debate over what it means to be formal found in (Dutilh Novaes 2011).

⁹Discussion of the relationship between incompleteness, consistency and proof will be the subject of chapter 2.

93–94)

The conservativity requirement holding between algorithmic systems in which ‘the same’ formal proof is located comes closely coupled with a *translation* of the ideas up and down systems:

Indeed, provided one is very strict about concept-individuation conditions, what can be claimed is that the new systems come with all-new concepts—and the old ones have simply been stipulatively identified with (some of these) new concepts. Such a stipulative identification of concepts that proves valuable is innocuous solely because of the cumulative way that algorithmic systems are embedded in one another: none of the old results regarding the old set of concepts are jettisoned—new material has only been added. (Azzouni 2004a, p. 98)

Azzouni rightly observes the need to deal with **(Techniques)** and, specifically, that many informal techniques do not seem to point directly to something formal. The particular example Azzouni gives is using symmetry, i.e. doing one part of a proof and then observing that another part is proved symmetrically. What is understood is that the part of the proof already given could be easily edited and adjusted to give the other part, though the exact details of such an adjustment are never given. The solution he offers is that in the course of informal proofs mathematicians may be using ‘meta-level’ reasoning, which means that the system(s) that the indicated derivation is located in will be ‘larger’:

When formalized as a derivation, such a proof will necessarily contain metamathematical elements which naturally drive it into the form of a derivation in a system strictly larger than one about, say, the objects officially under study. Mathematicians automatically ascend to a discussion of what can be taken to be properties and relations of the relations and properties of the objects they are proving results about. (Azzouni 2004a, p. 94)

Here, his discussion of how to deal with the case of symmetry additionally reveals some of the main evidence of what his view on **(Content)** is. It appears that aside from these tricky cases of meta-level reasoning and the like, the actual link from informal proofs to formal ones will usually be a

straightforward ‘filling in the gaps’-type process. However, in developing the view further in later work, Azzouni explicitly moves away from this ‘filling in the gaps’ account to a more sophisticated picture separating the way we come to understand informal proofs (through ‘inference packages’) from the way that corresponding formal proofs are determined (see Azzouni 2005b, p. 40). For simplicity I focus on the earlier account, but believe that the problems I raise for it will remain problematic for the mechanisability requirement still present in the more sophisticated picture.

1.4 A Dilemma

It is now time to start exploring the relation of derivation-indication more thoroughly. A particularly weak understanding would be to see informal proofs as a kind of time-saving communicative device, allowing mathematicians to quickly transfer formal proofs by indicating them to one another using informal proofs. However, this is not Azzouni’s intended meaning; he, in fact, explicitly rules out the idea that mathematicians need to be aware of the underlying derivations (“I should add that it isn’t a requirement on ‘indicating’ that mathematicians, generally, be aware that their proofs indicate derivations.” (Azzouni 2005a, fn. 16) or similarly (Azzouni 2009, fn. 17)). So if not this, what is meant by indicating? Since the general intention is to give an account of **(Agreement)**, **(Correctness)** and **(Rigour)**, it appears that what is required is that indication is some kind of dependence relation, but what properties it should have is just one of many questions that must be faced to complete the account.

The particular question I propose to press for this account is the following: is derivation-indication agent-dependent or agent-independent?¹⁰ Since, in essence, it is a *proof* that indicates a derivation it is relevant to ask who the supposed agent in this dilemma is. The proposal is that, on the one hand, the dependence link could be argued to not involve any kind of agent (say mathematician, student, listener, reader or anyone else that is involved in the particular instance of the proof). On the other hand, the agent-dependent horn of the dilemma suggests that the link from informal to

¹⁰This question is very close to the question of whether formalisation is a process that varies with the agent performing it, like Carnap’s notion of explication (Carnap 1945) or whether it instead is a process of revealing the ‘deep structure’ of the target phenomenon.

formal proof may not be fully present in the proof itself, but instead something over and above generated by the practice of proving or formalising i.e. something that is added by some involved agent.

In what follows I will examine the two horns of this dilemma, arguing that Azzouni is proposing an agent-independent link between formal and informal proof. However, I will contend that taking this horn will not be successful, based on a problem of the informal proofs corresponding to multiple, non-equivalent formal proofs. Taking the link to be agent-dependent, I argue, is not an escape option for the Formalist-Reductionist though, because doing so fails to satisfy the original motivations of the Formalist-Reductionist enterprise.

1.5 Agent-Independent Derivation-Indicators

In this section I will consider the agent-independent horn of the dilemma, investigating the correspondence it posits between informal and formal proofs in order to show the ways in which this correspondence cannot support the answers to the various desiderata set out above.

Let us consider the following question: does each informal proof relate to just one unique formal proof or to many of them? We have already seen that for Azzouni each informal proof relates to a whole family of derivations, due to the fact that he believes that ‘the’ formal proof is located in a range of algorithmic systems and, strictly speaking, these are different proofs.¹¹ The question can be reissued in these terms, though: for some given informal proof, is there a unique formal proof relative to each algorithmic system it appears in? In all cases where Azzouni touches on the issue, he seems to want each informal proof to pick out one unique formal proof per algorithmic system, within the upper and lower bounds.¹²

Let us think about this kind of uniqueness, since proof identity conditions are central to the problem I raise this section.

We already saw in section 1.3 that on the Formalist-Reductionist picture the informal proofs depend on formal proofs to be able to meet the

¹¹This is because a formal proof is relative to a formal system and language.

¹²It should be noted that Azzouni, despite attempting to deal with some of the key issues of mathematical practice and informal proof, is never particularly explicit about the answers to these questions. Dealing with the various options for what he can and may want to mean is precisely the current undertaking.

various desiderata laid out above. In the light of this, the identity of proofs is highly relevant because it affects which formal proof(s) an informal proof depends on, and consequently impacts how well the desiderata are met. For example, if an informal proof does not depend on a unique formal proof (per algorithmic system) but instead depends on multiple, non-identical or non-equivalent formal proofs, then this could lead to further difficulties, say, in satisfying **(Rigour)** and **(Correctness)**. For it is the underlying formal proofs that are meant to be ensuring the rigour and correctness of informal proofs, but if there are multiple different formal proofs simultaneously being depended upon this undermines the effectiveness of the explanation the derivation-indicator account gives. For example, what is there then to stop an informal proof from corresponding to both one correct and one incorrect formal proof? The point is that if it is the case that the informal proof does not uniquely determine which formal proof it depends on, then the dependence is far weaker than is required to actually satisfy the desiderata. Once it is conceded that there are multiple different, non-equivalent formal proofs underlying some informal proof, we can immediately ask why it is these particular ones that are selected and what ensures that it is only correct and rigorous formal proofs that are picked out. Now, if we need an extra step to clarify why the informal proof only corresponds to just those formal proofs which do ensure rigour and correctness, then it is this additional step that is doing all of the philosophical work and the account given has failed to properly answer the questions posed.

If the underlying formal proof is unique in some sense, then it seems the structure of the formal proof could, perhaps, be related to the content of the informal proof and avoid this underdetermination. Such considerations are also clearly present in Azzouni's theory: the fact that he writes of the underlying formal proof in the singular¹³, even when it is in fact located in different algorithmic systems with different languages, does not appear to be accidental. Of course, we did see that this required two extra components. Firstly, the moves 'upwards' had to be conservative of the derivational results, to make sure 'the' formal proof is still present as one extends the system. Secondly, we need to be able to identify proofs up and down systems to ensure they are still the same in this crucial sense. As we saw above, this is achieved by stipulatively identifying concepts between formal systems. I

¹³As evidenced by the quotations in section 1.3.

shall return to these moves once the worry has been further articulated.

In the remainder of this section I will show that the lack of a unique determination of the formal proofs an informal proof depends on does indeed occur and that as a result the problem just described applies to the derivation-indicator view on the agent-independent horn of the dilemma. I believe the other horn does not suffer this same problem, as will be discussed in section 1.6. Another way of describing this problem is as an *overgeneration* problem.¹⁴ The idea is not that informal proofs are too resistant to formalisation but instead that they are not resistant enough. There are multiple, equally legitimate formal proofs corresponding to any given informal proof and it is this multitude which throws doubt on there being any deep philosophical significance to the correspondence at all.

When proposing this overgeneration worry for Formalist-Reductionist views, something that has often been brought up in response is whether or not the difference between the various proofs is *substantial*. The thought is, presumably, that if the type of difference between the various formal proofs is only minor or insubstantial, then the proofs may be essentially the same and so the overgeneration problem loses its bite. However, I do not find this distinction particularly helpful in avoiding the problem for two reasons. Firstly, while being essentially the same may hold for two formal proofs with only some minor change, making lots of minor changes could add up to a substantial change quite easily. Secondly, I don't believe there is any robust way of separating the variations between formal proofs into substantial and insubstantial ones, but rather think that there is a whole spectrum of potential variations, ranging from very minor differences all the way to having no commonalities at all. Nonetheless, I will accept the distinction for the sake of argument and proceed to why I think there will be both the smaller and the more substantial variations between the formal proofs that some given informal proof will depend on.

Given some informal proof, it is straightforward to see that there must be a selection of formal proofs that it corresponds to just from the minor and insubstantial variations that can be introduced. Examples I have in mind are variable-renaming; changing the order of independent lemmas; switching

¹⁴I have adopted this terminology from Etchemendy's attacks on the reductive view of logical consequence in (Etchemendy 2008). Furthermore, I believe the spirit of the arguments Etchemendy gives is very close to those deployed here.

between inter-definable logical constants; changing the order you prove biconditionals (i.e. starting right-to-left or left-to-right) etc. Of course, the kind of changes that are minor will depend on the particular proof, since at times these rather innocuous differences can be relevant (or even crucial) to the success of the proof. This not only supports my claim that the distinction between minor and substantial differences is not a robust one, but also the more general argument I am making that even the minor differences can potentially cause problems for the agent-independent position.

Now Azzouni has essentially two options. He can stick to his guns, as it were, and insist that for any given informal proof there is just one formal proof per algorithmic system, in which case he fails to capture basic intuitions about formal proof identity concerning these minor variations, say, as well as being exposed to a worry about the arbitrariness of the particular proof that underlies the informal proof. Alternatively, he can accept that there is instead some *equivalence class of formal proofs* in each system matching up to any informal proof. In this case, for some given informal proof and an appropriate algorithmic system, there is a class of formal proofs that the informal proof indicates. It seems obvious that Azzouni should take the latter option; given that he accepts inter-system identity of proof, intra-system identity does not appear to be any more problematic.

However innocuous intra-system identity may seem, it is in fact deeply problematic, even in cases of insubstantial variation. To begin, a concern is that even though we have seen some suggestions for the acceptable minor variations listed above, if the minor variations do still keep the given formal proof ‘essentially the same’, then we would certainly like a more complete description of the kind of variations that are acceptable. With this comes the further need to justify such choices and convince us that adding up the differences will not eventually amount to a more substantial change. Considering the huge variety of systems that we could be talking about here, these demands will not realistically be met. The rhetorical point, though, is that the granularity of the notion of proof identity in play will have a bearing on how well the theory holds up under scrutiny.

Even if there are answers to the questions of the previous paragraph, this does not settle the matter concerning proof identity. Azzouni’s theory, for good reason, identifies proofs between different algorithmic systems via the stipulative identification between concepts and conservative translations

between the systems. Again, if we were dealing with just a single formal proof in each algorithmic system then this process might work, but if there is an equivalence class of formal proofs underlying some informal proof in each algorithmic system, then once again there are technical issues that must be addressed. Even when the variations are minor relative to some particular algorithmic system, those differences could be exacerbated and enlarged by the translation between systems. Formal proofs that were essentially the same (in the sense of being in the same equivalence class) in one system could, for all we know, be translated to proofs that are no longer the same according to the equivalence conditions in that other algorithmic system. Suppose we have two formal proofs P and Q in algorithmic system A that are both in the equivalence class underlying some informal proof, then translate them in Azzouni's sense to some other algorithmic system B . There is no guarantee that the translations $t(P)$ and $t(Q)$ will be in the equivalence class for the informal proof in system B .

There are two ways that one might try to avoid this concern of identity and translation: by appeal to conservativity and stipulative identity. Conservativity ensures that no results are jettisoned when moving between algorithmic systems, so we are safe in the knowledge that whatever we have a proof for in the weaker system will also have a proof in the stronger one. Yet this is certainly not enough to avoid the problem, since the way it is posed does not require the result to disappear, rather that the translation may take minor differences and make them substantial in the translation process. This can certainly happen if the result is still present in the new system. The fact that the identification between systems is stipulative can also not do any work here, because as we saw above the stipulative identity is only argued to be innocuous thanks to the conservativity. Now I argue that when it comes to formal proof identity, the stipulation of identity might not be innocuous (in that substantial differences might creep into proofs during translation) and that conservativity does not allow a way out of this fact, therefore making use of stipulative identity would beg the question.

So much for minor variations; what of more substantial ones? Are there ways in which the underlying formal proofs can differ which amount to significant and sizeable differences? I believe that there certainly are and will now give an example where this can be seen. First, though, I want to give some thought as to what 'substantial' differences could be like. There is a

sense in which the type of differences is constrained by the informal proof that the formal proofs all correspond to, yet this constraint does not, I argue, prevent substantial differences from appearing. The two most straightforward places to see this are in the treatment of mathematical objects and the mathematical dependencies a theorem has. Firstly, the treatment of the objects of an informal proof have to give some formal reconstruction of the objects in terms of relevant properties (at least those that are used in the proof). How the objects are represented in the formal system, then, will affect how the formal details of the proof go. Even for these details there may be multiple different ways to do things (totally ignored in the informal proof). Together we get different formal constructions (which will only have to overlap in some crucial properties) with different technical details. Of course, differences in the representation will have knock-on and snowballing effects the further through the proof we go, as the different representations and details of the formal proofs cascade along. After all, the exacting nature of formal proofs brings with it a delicate balance that must be maintained for the proof to be correct. Secondly, by representing the proof in different places, the mathematical dependencies that the proof has will be altered to support the type of specific inferences that may be made in that system. From all of these factors, the appearance of variations between the formal proofs that are substantial should be expected.

Let us flesh this out with a concrete example. The one I have in mind is that of the mutilated chessboard.¹⁵ The statement and proof are the following:

An ordinary chess board has had two squares—one at each end of a diagonal—removed. There is on hand a supply of 31 dominos, each of which is large enough to cover exactly two adjacent squares of the board. Is it possible to lay the dominos on the mutilated chessboard in such a manner as to cover it completely? (Black 1946, p. 157)

It is impossible ... and the proof is easy. The two diagonally opposite corners are the same color. Therefore their removal leaves a board with two more squares of one color than of the

¹⁵This example is central in (Robinson 1991) and can also be found in (Black 1946) and (Gardner 1988).

other. Each domino covers two squares of opposite color, since only opposite colors are adjacent. After you have covered 60 squares with 30 dominos, you are left with two uncovered squares of the same color. These two cannot be adjacent, therefore they cannot be covered by the last domino. (Gardner 1988, p. 28)

This example has intentionally been chosen as one which is intuitively correct, rigorous and understandable but also has a great deal of freedom regarding the underlying formal derivations that Azzouni's theory is committed to. Another advantage of this proof is that it has been formalised a number of times in different systems as a good example of informal reasoning that is tricky to capture formally. The standard way that the various attempts tend to approach the problem is reconstructing it set-theoretically, after this was issued as a challenge in (McCarthy 1995), with such attempts found in (Bancerek 1995; Rudnicki 1995; Subramanian 1994). These represent the board as sets of co-ordinates and then define an adjacency relation on the sets of co-ordinates with a tiling making use of this relation. However, this is not the only approach that can be taken, as is demonstrated by Paulson in (Paulson 2001), who instead makes use of inductive definitions for the set of dominoes and the tiling, which then allows crucial properties to be proved by rule induction. Furthermore, in (Subramanian 1996), Subramanian takes another distinct approach, in terms of states (of the chess board) and actions (of placing dominoes on the board), with a focus on modelling finite state machines.

Inspection of these various different formal versions of the same informal proof shows that substantial differences do appear. Let us consider these differences in turn.

Firstly, the different representations of the various objects of the proof lead to major differences in the way that the main steps in the proof are formalised. The representation obtained by defining tilings in the direct set-theoretic way, inductively or over states, are different ways of giving formal accounts of the various parts of the informal proof. Taking these different approaches actually results in changing the significant steps in the proof. For instance, in (Paulson 2001) the rule induction does a great deal of the work for proving various facts, while the (Subramanian 1996) version has to add a large number of additional, almost trivial facts to make the proof go through.

Secondly, these different representations also lead to different requirements for how the technical details in between the major steps are cashed out, with no obvious way to translate these between the systems. On the first approach, a number of extra details are added to pick out the right set-theoretic facts and make these particular details form into a correct proof. On the second, the inductive definitions are added to the set-theoretic components to allow several of the sub-proofs to be completed in substantially different ways. The third approach has to add lots of extra details to ensure that the numbers of coloured squares are well-behaved moving between states.

Thirdly, invoking such different representations means that the proof is dependent on different mathematical facts and background assumptions, which is to say that we arrive at a proof with different dependencies. In our chess board example we can point to the difference between a reliance on set theory alone, the soundness of rule induction and the logic of the state machines, as well as a whole selection of minor discrepancies.

Finally, since the formal proofs are all in different systems, we see the additional problem that each system adds to the formal proof a number of system-specific artefacts. These seem to be totally without correlates, should we want to attempt to translate between the systems.

All of these are examples of substantial and sizeable differences in formal derivations corresponding to the informal proof of the mutilated chess board example. As such it should now be clear that substantial variations do exist in underlying formal proofs and that the problem of overgeneration described above is in full force against the derivation-indicator account.

One common response that Azzouni would usually have open is, interestingly, not available in this case. The response would be that Azzouni could give up the need for translations between systems and simply rely on the fact that for each relevant algorithmic system there exists a formal proof that corresponds to the informal one. So long as this property of existence is preserved when moving between algorithmic systems, as Azzouni holds is constantly happening in mathematical practice, the derivation-indicator account can be maintained.¹⁶ The reason this path is not open on this occasion is that my objection is one of overgeneration. In effect, I am in-

¹⁶Thanks to an anonymous referee for this suggestion as to how the objection may be avoided.

sisting that this type of mere existence of corresponding formal proof is not sufficient for the correspondence to provide adequate solutions to the problems of **(Rigour)** and **(Correctness)**. Furthermore, giving up the translations between systems also concedes the point that there are multiple different and non-equivalent formal proofs underlying an informal proof, leading back into the need for an additional explanation of how it is that the derivation-indication only picks out correct proofs etc. Such a retreat does not, therefore, avoid the problems we have seen.

As a final point against the agent-independence of underlying formal proofs, I add that the problems of the identity and uniqueness of formal proofs strike me as the easier ones to handle compared to questions about the identity of informal proofs. There are very strong intuitions concerning which informal proofs are the same and which are not (an issue that is, for example, important to properly crediting mathematicians for their new discoveries). Presumably, in trying to examine mathematical practice, at least some attention should be paid to ideas of informal identity. An even broader line of difficulties would emerge from this though, concerning whether informal proofs which are informally identical should indicate the same classes of formal proofs and if not, why not.

All of the above follows Azzouni along the agent-independent horn of the dilemma. Taking the other horn would make matters like this far easier to deal with, since which equivalence class of formal proofs (both inter- and intra-system) underlies the informal proof would depend on the particular agent and circumstances of the informal proof. Unfortunately, the second horn cannot be what Azzouni wants because it does not suffice to establish the Formalist-Reductionist claims, as I will argue in the next section.

1.6 Agent-Dependent Derivation-Indicators

So let us consider the other horn of the dilemma, which has it that the formal proof(s) underlying any given informal one are agent-dependent and supplied over and above what is already present in the proof itself. This horn would yield great benefits: there would be readily available practical evidence that proofs can be linked to formal derivations from the field of Formal Mathematics, in which there is an ever-growing collection of computer-checkable

formal counterparts for well-known mathematical proofs.¹⁷ At least on the surface, the success of this grand formalisation project should add great credence to the idea that informal proofs can be linked to derivations. Formal Mathematics is very clearly agent-dependent, with different mathematicians converting different informal proofs to equally different derivations (as we saw in the mutilated chess board example). In this section I will make the case that the agent-dependent horn of the dilemma is, unfortunately, not available to Azzouni or other Formalist-Reductionists.

Firstly, as we have seen, Azzouni insists that the agents need not be aware of the indicated derivation that underlies the informal proof they are communicating. This in itself seems to put a stop to agent-dependence for Azzouni, for if the formal proof depends on agents who have no access to the formal proof there is little hope of success in this direction. Furthermore, one of the main desiderata for Azzouni, that of **(Agreement)**, would be left in a far more precarious position. For the social agreement on what constitutes a correct proof is explained in terms of the indicated derivations, but if the link is now agent-dependent then there is no given reason why any two people will have the *same* class of derivations underlying the informal proof. In this case, mathematics could then end up as a lot of talking past one another.

The original reasons for wanting to reject the notion that mathematicians are aware of the underlying formal derivations are good ones. Firstly, this simply does not match up to the reality of mathematical practice. Secondly, this would fail to answer the clause of **(Agreement)** which asks for an explanation of mathematics done long before there were formal proofs in mathematics. Finally, formal proofs for mathematics tend to be long and unwieldy therefore not the kind of thing that are ‘easy’ to know. In (Pelc 2009), it is argued that the formal counterparts to informal proof of theorems that have already been proved may very well not just be currently inaccessible to us, but beyond the physical limits of our universe to ever check.¹⁸ By avoiding having the mathematicians aware of the underlying derivations, Azzouni will be sidestepping these three concerns. Except, if Azzouni were to now take the second horn of the dilemma then these worries would be back with a vengeance. For in that case he would need to

¹⁷In mechanical proof-checkers such as *Coq*, *Mizar*, *Isabelle*, etc.

¹⁸See also (Boolos 1987) for another unwieldy formal proof for a clear informal one.

explain how the link from informal to formal proofs can be agent-dependent while the agent may nonetheless have no access to the formal proof, which is precisely the type of worry he was attempting to sidestep.

Generalising somewhat to other Formalist-Reductionist arguments, there is an even more crucial reason that they should not want to accept that their link is agent-dependent. This is that whatever the posited link may be from informal proofs to their formal counterparts, this link is one of *dependence*. The entire project is aimed at explaining the utility of informal proofs in terms of formal proofs, with their philosophically more straightforward account of logic and deductive reasoning. The desiderata of **(Correctness)** and **(Rigour)** can be tackled by taking advantage of the dependence of informal proofs on formal ones to import the story of rigour or correctness we have for the latter. However, if we make this link agent-dependent, then the clear waters are muddied once again by the complicated relationship that the posited link has with the mathematicians themselves. For the entire point of the undertaking is to resolve the tricky problem of the practical, real-life side of mathematics in the philosophically simpler terms of formal derivations. If the posited link is agent-dependent, the very difficulty that we were resolving simply re-emerges at another level. In short, the attempt to answer the problem of informal proofs in mathematical practice will find itself once again dealing with the practical difficulties of formalisation. This should be unacceptable to any Formalist-Reductionist account.

1.7 Conclusion

Having seen that the agent-dependent approach is not compatible with the Formalist-Reductionist aims, let me now return to the first horn of the dilemma and give a reason as to why an independent link from informal proofs to underlying formal ones is going to be particularly hard to establish.

The reason is embodied in the desideratum of **(Content)**. To successfully give the type of account that the Formalist-Reductionist is after, one has to go from the informal, implicit, gappy and often hidden structure of the informal proof to a fully explicit formal proof, which has picked out everything down to the smallest details. But one of the obvious reasons that formal proofs are rarely employed in practice is that these minutiae will get in the way of explanation, comprehension and communication of proofs. The

result is, unsurprisingly, that they are often left out. What follows, then, is that the link is adding extra structure and detail in going from informal to formal proofs.

In section 1.5 we saw the mutilated chess board and a selection of substantially different ways that it can be made formal. Although such multiple realisations of the formal proof's corresponding to informal ones don't pose such a problem to a weaker, agent-dependent notion of formalisation, if we want an independent link this is a serious problem because it compels us to go beyond the link to explain which realisation is the correct one or how they can all be correct, and as we have seen neither option is particularly easy. The difficulty of filling in the leaps made in informal proofs is further compounded by the fact that proofs have a great deal of structure, which means that how we fill in a gap at one point can and does affect the options for later stages of the proof. Believing that the answers to these technicalities is somehow already present in the proof and determined is entirely misguided.

The moral, then, is that satisfying **(Content)** is really quite a challenging problem. Interestingly, the problem is one that extends far beyond Azzouni's particular proposal to Formalist-Reductionist projects generally. Whether one wants to reduce all mathematics to formal derivations, claim that informal proofs reveal a complete logical form, or any other proposal in this direction, the hard problem of **(Content)** is a serious roadblock.

Chapter 2

Saving Proof from Paradox: Gödel's Paradox and the Inconsistency of Informal Mathematics

Every formal system is thus incomplete in two respects: 1 insofar as there are propositions undecidable within it, and 2 insofar as there are notions that cannot be defined within it [...] Thus we are led to conclude that, although everything mathematical is formalizable, it is nonetheless impossible to formalize all of mathematics in a single formal system [...]

— Kurt Gödel (Gödel 1935, p. 389)

2.1 Introduction

Is mathematics consistent? While in practice we generally proceed as if it is, for dialetheists such as Priest in (Priest 1987), mathematics is one of the main battlegrounds on which to establish that inconsistencies do indeed arise and require their dialetheist solutions. In this chapter I shall consider two related avenues of argument that have been used to make the case for the inconsistency of mathematics: firstly, paradoxes which lead to contradictions internal to mathematics and, secondly, the incompatibility of completeness

and consistency established by Gödel's incompleteness theorems. These two strands of argument are closely connected, for the most apparently problematic paradox in the case of mathematics is *Gödel's paradox*, that of the sentence which says of itself that it is unprovable, which is closely related to common constructions of Gödel sentences for formal systems whereby we get to the balancing act between completeness and consistency.

My response to the two lines of dialetheist argument will bring in considerations from the philosophy of mathematical practice on the nature of informal proofs. One thing I will argue for is that we should add to the two axes of completeness and consistency a third axis of *formality* and *informality*. Given this third axis, we can consider the dialetheist arguments in two different ways. At the informal end, the previously problematic paradoxes may be genuine, but I argue that there is no compelling reason to see them as internal to mathematics. Meanwhile, at the formal end of the scale, considerations of the practical role of formalisation in mathematics will allow me to make a positive case for incompleteness over inconsistency without begging the question against the dialetheists. My main conclusion will be that the dialetheist arguments considered do not establish that mathematics is inherently inconsistent.

Answering the ultimate question of whether mathematics is consistent from this perspective which encompasses informal proofs and mathematical practice would, I believe, be a major undertaking, and one which I am not intending to complete here. The intention is rather to take the first step in this direction by demonstrating that the matter is not already settled, since the standard arguments from Gödel's theorems and the paradox of provability do not succeed. In fact, I believe these arguments fall apart through a number of the assumptions they need about informal proofs, the nature of mathematics and the process of formalisation, so I shall proceed to raise these objections in turn.

To begin, section 2.2 will introduce the key distinction between formal and informal proofs that my arguments will focus on. Next, in section 2.3 I will lay out what Gödel's paradox is and why I do not take it to be a concern for mathematics. In section 2.4 I present Priest's longer argument for the inconsistency of informal mathematics based on the application of Gödel's first incompleteness theorem to informal mathematics and the conclusions he draws from this concerning the *inherent* inconsistency of informal math-

ematics. In section 2.5, I argue that the way of understanding formalisation on which Priest's argument succeeds is a bad one, then show that a better understanding means the argument no longer goes through. In sections 2.6 and 2.7, I argue against the thought that we can formalise mathematics as a single theory, proposing that a better thought would be to approach formalisation in a fragmented way. Finally, in section 2.8 I consider formality and informality as a third axis, and a final argument against Priest that he changes the subject in switching between the formal and the informal.

2.2 Formal and Informal Proofs

Before we can begin, we need to be sufficiently clear on the distinction between formal and informal proofs, as this will play a central role in the remainder of this chapter.¹

Formal proofs are those which are studied in logic and proof theory, and may be defined in the usual way. For example, we might define a formal language, give rules for well-formed formulae in that language, specify axioms to be taken as basic and lay down inference rules for stepping between formulae. A *formal proof* (relative to such a specified system) will be a (usually finite) sequence of formulae where each is either an axiom or follows stepwise from previous formulae by an application of one of the inference rules, where the final formula is a statement of what was to be proven and is thus established as a theorem in the system.

However, formal proofs are rarely seen in actual mathematical practice. Instead the type of proofs that are employed by mathematicians in their daily activities, teaching and published work tend to be very different. In most cases no formal language is specified, axioms are rarely given and inferences are not confined to just the basic rules. Steps in these proofs can rather be leaps, and they can invoke the background knowledge of the target audience, the semantic understanding of the terms being employed, visualisation, diagrams and topic-specific styles of reasoning. Let us call proofs in this sense *informal proofs*. Although this would be extremely unsatisfying as a definition, it is certainly not intended as such as one of the

¹A terminological note: while I speak of 'informal proofs' and 'formal proofs', some of the literature on this subject instead speaks of 'proofs' and 'derivations' to get at the same distinction. In (Priest 1987), Priest also uses the term 'naïve proof' to refer to the informal proofs.

main challenges for philosophers of mathematical practice is to pin down exactly what counts as a good, legitimate, correct and rigorous informal proof and filling this out further would take me beyond the scope of this chapter. Nonetheless, there is a good deal of literature that does deal with this issue that elaborates on the distinction I am invoking (see Robinson 1991; Hersh 1997; Rav 1999; Leitgeb 2009; Antonutti Marfori 2010; Larvor 2012).

A number of the differences between these two types of proof will affect the assessment of whether the arguments I am considering successfully establish that mathematics is inconsistent. Gödel's first incompleteness theorem relates to proof as an explicitly defined, formal notion attached to a formal system and one of my main counter-arguments in what is to come is that this will not transpose across to apply to informal proofs. Gödel's proof tells us about the limits of formal systems *which meet certain conditions*, like having a certain degree of expressive power, being able to prove a certain amount of basic mathematics (enough to allow for the required coding etc.) and having an effective procedure for enumerating its theorems. What will be required for the dialetheist line to work, then, will be to show that informal proofs are close enough to formal ones to even begin applying these conditions. I will argue to the contrary that informal proofs are sufficiently different from formal proofs that the argument does not succeed. Some key differences of informal proofs that will play a role later include the social and contextual components of whether such a proof is successful or not; the partially-fragmented nature of modern (informal) mathematics; and the fact that informal mathematics extends to include diagrammatic proofs which have more intuitive inferential rules. Finally, even if the dialetheist arguments manage to establish that informal proofs can be formalised appropriately, there will still be the need to show that the conditions are met.²

Before getting the details of the argument from Gödel's theorems, let me assess whether a simpler argument from paradox outlined by Beall is sufficient to show that mathematics is inconsistent.

²In (Priest 1987), Priest argues that these conditions will be met. I believe that the flawed step in the argument is the earlier one of formalisation (as will be covered in section 2.5), so I will not actively engage in a discussion about whether this formalisation will have an effective calculus etc.

2.3 Gödel's Paradox and Beall's Argument

The first argument I will consider is based on *Gödel's paradox*. Let us begin, therefore, by examining the paradox:

GP: This sentence is (informally) unprovable.

Suppose GP is false; then it is informally provable. Since we take our informal mathematical proofs to establish mathematical truths, it follows that GP is also true. Yet this contradicts the assumption that GP is false, so using proof by contradiction we establish that GP is true. However, since we have just proved GP, it is informally provable. But GP states that it is unprovable, so it must be false. Contradiction.

Now consider how it is that this paradox might show that mathematics is inconsistent. Beall gives the following argument:

There seems to be little hope of denying that [GP] is indeed a sentence of our informal mathematics. Accordingly, the only way to avoid the above result is to revert to formalising away the inconsistency— a response familiar from the histories of naïve set theory, naïve semantic theory, and so on. If one does this, however, then (by familiar results) one loses completeness, which can be regained only by endorsing inconsistency. Either way, then, we seem to be led to inconsistent mathematics. (Beall 1999, p. 324)

Setting aside the option to formalise away the inconsistency until section 2.4, the initial argument is that since GP is part of mathematics and GP leads to an inconsistency, it must therefore be that there is an inconsistency in mathematics. In the rest of this section I will undertake the (purportedly hopeless) task of denying that GP is part of mathematics.

The only sensible suggestion as to why GP should be part of mathematics, it would seem, is that GP concerns the broadly mathematical concept of informal provability. I contend, though, that this is not sufficient to make GP a statement of mathematics. The reason is that I take the concept of informal proof to be used to talk and reason *about* mathematics without it being *a part of* the subject matter of mathematics. While the former is obvious, for the paradox to render mathematics inconsistent we actually need the latter, more contentious claim. Of course, I hold that informal proof

and provability are very important notions in talking about mathematics, but it is crucial to emphasise that these are notions *about* mathematics. To establish that the paradox will render mathematics inconsistent, though, we need the extra claim that it is *a part of* informal mathematics. In general, a statement being about mathematics and a statement being part of mathematics can coincide, but certainly don't always. Consider the following:

- (1) Mathematics is traditionally done on blackboards.
- (2) This square building with $12m$ sides must have an area of $144m^2$.
- (3) $111, 111, 111 \times 111, 111, 111 = 12, 345, 678, 987, 654, 321$.
- (4) Ron likes bacon and eggs.

Here (1) is a statement about mathematics but is not itself a part of mathematics. In contrast, (2) is a mathematical statement which is being applied to a situation, so in a relevant sense is not about mathematics. The third item is both mathematical as a statement and about a mathematical fact, while the fourth sentence is neither. Since these two notions can be pulled apart with minimal effort, that a sentence falls under one of them certainly can't constitute a reason to think that it falls under the other. It can therefore be concluded that the notion of informal provability being about mathematics is not sufficient to establish that GP falls within mathematics.

One can also give positive arguments as to why informal provability should not be considered a concept within mathematics. For example, the lack of a precise mathematical definition we observed in section 2.2 clearly supports the claim that informal provability is not a notion within the subject matter of mathematics. Nor does it interrelate with other mathematical concepts in the way that standard mathematical concepts do (such as, for example, group, integer, derivative, line, etc.). The only notable conceptual link it has is with truth, as exploited by the paradox, but if anything the informal notion of truth in mathematics (before being formalised into some formal theory of truth) will belong to the same category of notions about mathematics that are not within mathematics.

By denying that informal provability is a concept within informal mathematics, it can consequently also be denied that GP is a sentence of our informal mathematics. It is thus reasonable to deny that Beall has showed that informal mathematics is inconsistent by using GP. This certainly does

not provide an ultimate solution to Gödel's paradox, but it does keep the derived inconsistency out of mathematics and allows us to set aside the paradox to be solved in line with whatever one's favourite solution is to paradoxes generally.³

Now, let me note two things about what has gone on here which will be recurrent throughout the chapter. Firstly, although this section does not solve Gödel's paradox, this is not really necessary for the purposes of the current project. Beall, Priest and others have a substantial case for the inconsistency of natural languages, a case which is not the target of this chapter and would have to be addressed separately if one were so inclined. For both of these authors the claim that mathematics is inconsistent is an additional one that is supported by additional argumentation and it is precisely these arguments which I am targeting. Thus, by rejecting that Gödel's paradox is part of mathematics, what has been done is to show that these additional arguments do not cover more ground than the original case for the inconsistency of natural languages and therefore don't provide added support for dialetheism from the realm of mathematics. Secondly, the separation between being part of mathematics and the concepts used about mathematics is not just a way to re-introduce the object language/meta-language distinction for informal mathematics. A separation of languages is not important because the point is not really one about languages, instead it is about the subject-matter of mathematics. While we may use GP to argue that the concept of informal provability is inconsistent, this does no more work than the liar or any other semantic paradox unless it infects the realm of mathematics. As such, showing that informal proof is not the kind of thing to be investigated mathematically blocks the argument considered in this section.

³A final note on Beall: although the argument I am criticising is from an older paper, the response offered here would fit well with Beall's more recent work in (Beall 2009). The suggestion I have made may be appropriated to make the case that informal proof should join truth in the category of useful devices, which when introduced bring 'merely' semantic paradoxes as by-products or 'spandrels' without thereby rendering the base language (in this case, that of mathematics) inconsistent.

2.4 Priest's Argument for the Inconsistency of Informal Mathematics

In Chapter 3 of (Priest 1987), entitled "Gödel's Theorem", Priest makes use of Gödel's paradox in the same way as Beall subsequently went on to do, arguing that it shows that informal mathematics is inconsistent. In Priest's case, however, it is given as the culmination of a longer argument which aims to show that informal proof satisfies the conditions for Gödel's first incompleteness theorem in such a way as to lead to its inconsistency. This section will focus on explaining the details of Priest's argument.

Priest wants to show that informal proof is susceptible to Gödel's first incompleteness theorem. The first hurdle is that the theory of informal proofs is, on the surface at least, not formal and hence not immediately susceptible to Gödel's theorem. Priest addresses this in the following way:

It should be said at once that naive proof, or at least the naive theory it generates, is not a formal theory in the sense of the theorem; but it is accepted by mathematicians that informal mathematics could be formalised if there were ever a point to doing so, and the belief seems quite legitimate. The language of naive proof, a fragment of English, could have its syntax tidied up so that it was a formal language, and the set of naïve theorems expressed in this language would be deductively closed. Hence we may, without injustice, talk about the naive theory as if it were a formal theory. (Priest 1987, p. 41)⁴

In section 2.5, I will claim that Priest's reasoning fails to go through at this point. For now, though, let us complete Priest's argument that informal proof satisfies the conditions of Gödel's theorem. The other pieces that Priest needs are that the formalised theory can express all recursive functions and that the proof relation of the formalised theory is recursive. He rightly takes the first requisite to be obviously satisfied and the second to be the

⁴As the target of his argument, Priest needs to explain what he takes naïve or informal mathematics to be exactly. He says:

Proof, as understood by mathematicians (not logicians), is that process of deductive argumentation by which we establish certain mathematical claims to be true. (Priest 1987, p. 40)

His distinction is, in effect, the same as the distinction between formal and informal mathematics as found in section 2.2 and throughout this thesis.

contentious one, listing a number of possible objections and his replies. A discussion of these would be irrelevant to the purposes of this chapter, so for now we shall grant that the formalised proof relation is recursive.

Given that Priest has now established that informal proof satisfies the conditions of Gödel's theorem, the thrust of his argument is as follows:

For let T be (the formalisation of) our naive proof procedures. Then, since T satisfies the conditions of Gödel's theorem, if T is consistent there is a sentence φ which is not provable in T , but which we can establish as true by a naive proof, and hence is provable in T . The only way out of the problem, other than to accept the contradiction, and thus dialetheism anyway, is to accept the inconsistency of naive proof. So we are forced to admit that our naive proof procedures are inconsistent. But our naive proof procedures just are those methods of deductive argument by which things are established as true. It follows that some contradictions are true; that is, dialetheism is correct. (Priest 1987, p. 44)

Priest soon makes the link between φ and Gödel's paradox. For if we take φ to be the formalisation of GP⁵, the inconsistency of section 2.3 will quickly re-emerge within the formalisation of informal mathematics. A key point is that a standard move towards incompleteness over inconsistency is to separate the object language from the meta-language, but that here we are dealing with informal proof and informal mathematics, for which there is no such distinction, meaning that the orthodox move towards incompleteness is not available. Indeed, this is the entire point of focusing the argument on informal mathematics.

The conclusion that Priest draws is that we are left with true contradictions and dialetheism.⁶ Informal mathematics is seen to be inconsistent, but

⁵The matter is somewhat more complicated than this suggests, of course. Milne discusses in (Milne 2007) the many ways that Gödel sentences can be constructed and what exactly they 'say'.

⁶Not just this, though, since Priest takes it that the theory given by the formalisation of informal mathematics can prove its own soundness and hence must be able to give its own semantics. From here he takes it to follow that it must be able to prove the T-schema for this theory inside the theory, giving him all of the paradoxes he describes as semantic (as opposed to set-theoretic paradoxes). For example, he lists the liar, Grelling's paradox, Berry's paradox, Richard's paradox and Koenig's paradox as falling under the umbrella of semantic paradoxes. In fact, then, Priest argues that "Our naive theory is semantically

even more penetratingly he can claim that there is no escape from this application of the incompleteness theorems to informal mathematics and so “[...] we might say that our naive proof procedures are not just contingently inconsistent, but essentially so [...] dialetheism is inherent in thought.” (Priest 1987, pp. 47–48) That dialetheism is inherent in thought is one of the main claims of *In Contradiction*, supported by several pillars of argument. The argument described here that informal mathematics is *essentially* inconsistent forms one of these pillars, but I shall now argue that this pillar will not hold any weight.

2.5 Formalising Mathematics

The move from the informal version of mathematics to a formalisation thereof is, in my opinion, too quick. By endorsing the claim that mathematicians take it that informal mathematics can be formalised, Priest moves from the informal theory to the formal one without much consideration of what this move entails or how the mathematicians he is invoking conceive of the formalisation process. For one thing, Priest might not want to endorse the naïve claims of mathematicians at all, since they most likely take mathematics to also only be consistent. If such claims were definitive it might thus spell the end of dialetheism.

Nevertheless, it is worth considering how exactly the idea that mathematics should be formalisable will work precisely. In the first half of this section I discuss two options, along with how they interact with Priest’s argument. The first follows a straightforward interpretation of Priest’s claim but is shown to fail as an account of the formalisation of informal mathematics. The second avoids the problems with the first but, I argue, no longer lets Priest’s argument go through.

2.5.1 A First Option

Let us call the first option *many-one formalisation*.⁷ The idea is that one takes the entirety of informal mathematics and tidies up the fragment of

closed and inconsistent. By contrast, any consistent theory cannot be semantically closed.” (Priest 1987, p. 47)

⁷The ‘many’ here is due to the fact that it might end up being case that multiple informal proofs are mapped to the same formal proof.

natural language expressing it to give a formal language. All of the informal theorems will have particular formal counterparts expressed in this one formal language, and the set of these formalised theorems is then deductively closed. For the first option, we consider this as the one single correct formal counterpart for the informal mathematics, a type of super-theory⁸ of mathematics, in which all the current basic assumptions and their consequences are contained. This mirrors a standard idea of formalisation involving a routine procedure of ‘filling in the gaps’ as is discussed, for instance, in the debate between Rav (Rav 1999, 2007) and Azzouni (Azzouni 2004a, 2005a) though ultimately rejected by both. Since the formalisation that occurs is crucial to the application of Gödel’s first incompleteness theorem to informal mathematics, it would be very convenient for Priest’s argument if the picture that is sketched here is the correct one, as this would take formalisation to effectively reduce informal mathematics to something formal, and thereby allow the argument to proceed.

Unfortunately, we have good reason to think that this picture cannot be correct. It is obvious that tidying-up syntax is not going to be a many-one mapping. If we start with the natural-language versions of our mathematical theorems, there will be a whole selection of ways in which we can reproduce these theorems in some particular formal language. Even translating very simple fragments of mathematics into simple formal systems can easily lead to a plurality of results. Scaling this up to include *all* of mathematics exacerbates this problem significantly. Add to that the fact that we don’t start with a particular formal language that we are to be translating the informal into, but instead generate it “on the fly” based on the syntax of our informal mathematics. That there will only be one possible result is not very plausible.⁹

Note also that the conversion of informal mathematics into this super-theory is not really like the standard conversion of informal mathematics into some ‘foundational’ theory such as ZFC set theory, which is potentially

⁸I use the terms ‘super-theory’ and ‘super-system’ throughout this chapter. I do not intend anything of the ‘super-’ prefix besides emphasising that it is all-encompassing of mathematics in the way described.

⁹An anonymous referee suggests that we may be able to distinguish between a plurality of results which are equivalent under translation and those which genuinely disagree. I believe, however, that this will not save the argument. In the critical discussion of Azzouni’s formalist account of proofs in chapter 1, I have argued that such a move is not going to deliver the substantial kind of formalisation required for the argument to proceed.

what Priest has in mind for the mathematicians that he invokes. For if this were the case we would quickly find ourselves with the Benacerrafian problem that there are a large number of different adequate representations for our informal concepts (see Benacerraf 1965). This would lead us out of the first option and its super-theory, into a picture where there are multiple different formalisations of informal mathematics.

I would like to emphasise here that the worry I am raising with the generated super-theory is nothing to do with its inconsistency (for such a theory would undoubtedly be inconsistent) and as such it is not open to the usual charge of begging the question against the dialetheist.

2.5.2 A Second Option

As a second option, Priest could hold that the formalisation process for all mathematics that he is after is actually a case of *many-many formalisation*. As I have already argued, there may be many different formalisations of mathematics, which Priest can accept as the case in order to avoid the problems presented against the many-one formalisation picture. In essence, this approach is embracing the plurality of formalisations as opposed to letting it become a problem.

However, accepting this path immediately adds an extra complication to the argument, in that now Priest's claims about the formalised version of informal proof must implicitly be quantifying over formalisations. In particular, each time he mentions the formalised version of a proof of informal mathematics, there is no one thing this refers to but instead a selection of different formalised versions of the informal proof. The next natural question to follow this up with is how to determine which formalisations fall under this quantification for any given proof. Put another way: which formalisations of informal mathematics will be adequate and acceptable? For example, a formal language which is too expressively weak to even state standard theorems would be inadequate and unacceptable. The question, then, comes down to finding (and defending) criteria of adequacy for these formalisations of informal mathematics.

Formalisation, as it is being conceived of here, is not a process of exposing an underlying logical form already present in the informal proof, or any thought in this direction. I take this to be the case because informal proofs will underdetermine the language, system and structure that such a

proof would adhere to and have. It is instead taken to be a process that is inextricably linked to the context in which it occurs. Relevant factors include the agent performing the formalisation, their purposes in doing so and the formal theory they intend to formalise the given informal proof into. It might be useful here to consider an analogy to Carnap’s notion of explication (as in Carnap 1945) where there is also no definitive fact of the matter as to what the correct explication is for some given concept. Instead the different results are compared and evaluated using pragmatic measures such as usefulness, simplicity, explanatoriness and precision.

In a similar way, there could be a whole range of formalisations that can be of varying degrees of usefulness in making some informal piece of mathematical reasoning fully formal. In Priest’s formalisation of all of informal mathematics we may find a number of different results which are of varying degrees of usefulness, explanation, accuracy, simplicity etc. Of course, amongst these there may be a number of formalisations that we would want to recognise as inadequate, such as that in the above example of an expressively weak language. We want some way of excluding these examples of ‘bad’ formalisations of informal mathematics from being implicitly quantified over in Priest’s argument. However Priest would want to go about this project, we can see that it adds significant philosophical ground that needs to be supplemented to the argument in question before it goes through.¹⁰

2.6 On Mathematical Super-Theories

A new worry that emerges from the consideration of different formalisations concerns the reliance on one (or indeed many) mathematical super-theories. Since we have seen the analogy to Carnap and want to evaluate our formalisations using pragmatic principles, we must consider whether unified mathematical super-theories, in the sense that Priest has proposed, are indeed the best when evaluated in this way. In this subsection I will briefly consider three reasons why this might not be the case.

Before I begin, though, let us just make explicit why for Priest’s argument there is now the need to formalise all of informal mathematics in one

¹⁰An anonymous referee proposes an additional argument against Priest based on this section: that the translation on the many-many case is not effective means that informal proof can therefore not meet the minimum requirements for falling under Gödel’s theorems. Grist to the mill!

go, in its entirety, into a super-theory. If this is not done another key step of the argument cannot go through, namely the step where it is insisted that the Gödel sentence is indeed provable. If we were to replay the argument just in arithmetic, for example, we would code in (the formalisation of) *informal provability in arithmetic* and soon discover the Gödel sentence is not provable in this formalisation. But here we would be free to take the traditional lesson that this is just a limitation on the formalisation, which may well be incomplete.¹¹ It is only by squeezing out all room for this incompleteness by quantifying over all mathematics and informal proof simpliciter that the argument could hope to successfully establish that the answer is actually inconsistency rather than mere incompleteness.¹²

Let us now consider why this super-theory will run into difficulties.

One worry may be that different fields or areas of mathematics might be best served by different formal systems, or even different styles of formal systems. For example, the study of algebra, set theory and geometry all appear very different at first glance, and so it may be that they are best served by being formalised into different formal systems (say, with different proof rules which better track the kinds of inferences made in these fields). Of course, the judgment here must be relative to some purpose of formalisation, but we may take the purpose at hand to be (something like) giving a formal reconstruction of the informal proofs, which tracks the inferential steps that were being used. To justify this, recall that Priest’s treatment of informal mathematics as a formal theory was meant to be “without injustice”.

The first problem I am proposing, then, is that it might be that different formal systems, which are tailored to different sub-areas of mathematics, might allow the more accurate reconstruction of the reasoning present in the informal proofs for those different areas. It also seems that Priest cannot point to the fact that the super-system(s) he is after are those that represent a ‘tidying up’ of the fragment of natural language that mathematics is expressed in, because the point that is being pressed here is that this talk is an over-simplification of a more complex process.

Relatedly, the second concern I have is that diagrammatic proofs may

¹¹And we are well used to theories being incomplete for more reasons than Gödel theorem. For instance, Peano arithmetic also has examples like Goodstein’s theorem and the Paris-Harrington theorem.

¹²Note that this cannot be avoided by insisting that the Gödel sentence must be part of naïve arithmetic without running afoul of the distinction of section 2.3.

lead to a significant worry for Priest. In referring to the “fragment of English” that informal mathematics is expressed in, Priest seems to miss a wide selection of mathematics that is communicated pictorially. Pictures can serve to communicate mathematical facts, but can also function as components of informal proofs or proofs in their entirety (see Nelsen 1993, 2000). How is this to be accommodated in the super-systems which are meant to formalise all of informal mathematics? What will the formalisation process do to diagrammatic proofs? If they are simply to be eliminated, this once again means that informal mathematics is undergoing a drastic change in the formalisation process. Alternatively, there are formal systems for diagrams which may serve to formalise some of the diagrammatic proofs. For example, there is work towards formal systems of mathematical diagrams in (Manders 2008) and more explicitly in (Mumma 2010) and (Avigad, Dean & Mumma 2009). However, we are now engaged in a project of making the super-systems, which originally sounded straightforwardly close to informal mathematics, encompass much broader pieces of mathematical reasoning. At the very least, this is a non-trivial undertaking which involves constructing a mixed-mode formal system which combines traditional syntactic components with formal diagrammatics. The work just cited suggests that this might be possible in certain respects, but it is certainly no mere triviality. A deeper worry, however, is that we are now able to question whether it will even be possible to capture all of the mathematical reasoning that occurs in informal proofs in formal systems, without doing violence to the source material. I shall return to this line of thought in section 2.8.

A third problem we encounter for the mathematical super-theory can draw on Priest’s own considerations of mathematical pluralism in (Priest 2012). Modern mathematical investigation extends to examining which results obtain from adopting different logics to work in. Yet if all the various investigations of different logics are taken to be part of informal mathematics, what happens when we formalise them into the one super-theory? Not only do we face the prospect of systems collapsing into one another, but the more alarming danger of triviality looms. Observe that some of the logics we might want to use will include the principle of explosion, most notably classical logic. As soon as a contradiction arises somewhere in the system (which is exactly what Priest’s argument is attempting to force), immediately it follows that the whole super-system is trivialised. This is regardless

of whether we think that there is something *philosophically* wrong with classical mathematics, and the principle of explosion in particular, since we are just formalising informal mathematics as we found it. This worry also doesn't rely on logical pluralism, instead just the more uncontroversial fact of logical plurality.¹³ In the case of this worry, Priest's argument will still go through but using the fact that a trivial super-system is also inconsistent, which is hardly a desirable result.

2.7 Fragmented Formalisations

The counter-suggestion to formalising all of informal mathematics simultaneously into one super-theory, with which we have seen some serious difficulties, is that the formalisation process may be one that can only be successful when done in a *fragmented* way. The suggestion is that constructing a formal system is achievable when we take smaller "chunks" of mathematics that we want to formalise, just not when we want to take it all at the same time. Such an understanding would provide reasonable solutions to dealing with the problems of previous section, without giving up the possibility of formalising parts of mathematical reasoning.

Let us see why switching from the idea of a super-theory to the fragmented approach is not a good option if we want to maintain Priest's argument that informal mathematics is inconsistent by Gödel's First Incompleteness Theorem. The issue is that the argument relies on capturing informal mathematics fully to insist that the sentence φ , which is unprovable in the formalised version of informal mathematics but is nonetheless established by informal proof, must also be provable in the formalised system. If, however, it fails to obtain that any one theory does successfully formally represent all of informal mathematics as a whole, then it cannot be insisted that the last step holds. The point is that we get to the fact that the sentence must be true in the system because the system includes all informal mathematical reasoning. If we do not guarantee this, then the inconsistency is not guaranteed either.

Undermining this last step is sufficient for giving a criticism of Priest's argument, but what we have seen so far forms a deeper difficulty. Priest's

¹³I take it that, as mathematicians, we don't need to commit ourselves to the truth, in some philosophical sense, of the mathematics that is being carried out.

more general project in *In Contradiction* is to re-examine the balance between completeness and consistency, insisting that it is the latter we jettison in light of Gödel's theorems rather than the former, which is the orthodox choice. Recall that in section 2.3 we set aside Beall's use of the same balancing act, where he suggests that when formalising mathematical reasoning we are returned to the completeness/consistency dichotomy. What has implicitly been done here, then, is to use considerations of the process of formalisation to give an independent motivation for why we might prefer to end up with an incomplete system when formalising informal proofs, without making reference to any concerns about consistency.

2.8 On The Formal and The Informal

For all that has been said, I think there is another more devastating objection to Priest's argument. In part 2.5.1 we saw that the idea that there would only be one formalised counterpart of informal mathematics would not hold any water. However, it was only on this reading that it seemed acceptable to treat informal mathematics as if it were a formal theory, at least superficially, stemming from the fact that there was one 'body' of informal mathematics and one formalisation thereof. Nonetheless, having been discussing the difficulties involved in formalising theories, it should now be clearer that there was something fishy going on in this step of the argument.

The objection is the following: by moving from informal proof to a formalised version thereof, Priest's argument is guilty of changing the subject. The argument intended to show that informal proof was inconsistent, and not just coincidentally but *inherently* so. Yet, almost immediately in the reasoning, to get the application of the incompleteness results off the ground, Priest needs the subject of his argument to be a formal theory. The answer, therefore, is that mathematics is not a formal theory and that transforming it to be one will do an injustice to its source material. The argument speaks as if the multiple representations that informal mathematics can have as formal systems are identical to the informal mathematics itself, but this is just a confusion of distinct things.

While Priest was looking to demonstrate that informal mathematics was inherently inconsistent, an option that is now on the table is that mathematical reasoning is *inherently informal*, a view common in the mathematical

practice literature (e.g. Larvor 2012), or that it may be inherently incomplete, or indeed both. The thought would then, in these cases, be that no formal system would suffice to adequately capture mathematics in its entirety. Indeed, this is the traditional lesson that people take from the incompleteness results, but this standard result relies on the question-begging move from consistency to incompleteness. Now, though, we have seen independent motivations for thinking so and rejecting the argument, motivations stemming from mathematical practice and paying attention to formalisation as a process.

Priest's challenge was looking to adjust the balance between consistency and completeness in favour of the latter over the former. But now, by considering the third axis of formality and informality, we have obtained a way to defend incompleteness over inconsistency in the formal setting without begging the question, and to see incomplete and inconsistent systems as both serving purposes which may be justified by pragmatic principles. For the argument relies on a number of assumptions about the nature of formalisation which allow one to easily and without injustice take informal mathematics into formal mathematics. I have, to the contrary, argued that this distinction runs deep and cannot be bypassed lightly, meaning that arguments that work for formal theories cannot be straightforwardly applied to informal mathematics, and ultimately that Priest's argument does not go through.

Chapter 3

Mathematical Concepts: Open-Texture, Dialectics and Engineering

Proof suggests new mathematics. The novice who studies proofs gets closer to the creation of new mathematics. Proof is mathematical power, the electric voltage of the subject which vitalizes the static assertions of the theorems.

— Philip J. Davis & Reuben Hersh (Davis & Hersh 1981, p. 151)

3.1 Introduction

In this chapter I will explore something deeply connected to the over-arching theme of proofs and their formality: the nature of *mathematical concepts*. For proofs feature and operate on mathematical concepts and as such the degree of formality or informality of a proof is closely related to the exactness of those concepts it deploys. If one believes, like the Formalist-Reductionist, that informal proofs correspond to formalised counterparts, then one should also see mathematical concepts as having exact, formal definitions which can be deployed in those formal proofs. Conversely, rejecting this idea leads to the opportunity to be more historically sensitive in seeing that mathematical concepts develop over time. Furthermore, we can also be more open

to the possibility of mathematical concepts not being fully fixed in their applications.

The philosophy of mathematical practice is heavily influenced by Imre Lakatos on precisely these issues. Lakatos thought that mathematical concepts were not fully fixed, but instead are changed and developed through the proofs they appear in. One way to fill out what this means could be through Waismann's notion of *open texture*, where a concept is open-textured when it is not fully delimited for all potential applications. The first half of this chapter will bring out the connection between Waismann's open-texture and Lakatos's dialectical approach to the philosophy of mathematics. With the door hereby open to ongoing conceptual development in mathematics, in the second half of the chapter I look to connect these ideas to recent work on *conceptual engineering*. I will suggest that deploying particular distinctions and strategies pertaining to conceptual change found in the conceptual engineering literature is a promising route for integrating the formal/informal axis discussed in the previous chapter with change of concepts in mathematics. This will ultimately leave us with three major questions. One that underlies the conceptual engineering literature concerns whether the concepts are to be *revised* or *replaced*. A second question is whether all mathematical concepts need to be changed or just some. Finally, we can wonder if conceptual change is needed for all mathematical contexts or just some restricted range of them. Of course, the answers relate to one another and jointly contribute to a view on mathematical conceptual change. I shall return to these questions at the end of the chapter.

The precise plan for this chapter is as follows. In section 3.2 I will start by explaining what Waismann holds open texture to be, and comparing it to Shapiro's more recent usage of the term, showing that the two are very close but are different in the particular phenomena they intend to pick out. Next, in section 3.3, we shall explore Lakatos's *Proofs and Refutations*, focusing on the roles of concepts and proofs. Following this, in section 3.4 I will briefly cover G. T. Kneebone, a figure who has received little to no attention but pre-empts Lakatos in several crucial respects in proposing a dialectical philosophy of mathematics. In section 3.5, I will bring the three figures together and discuss how Waismann's notion of open texture is a useful way of describing aspects of the Lakatosian and Kneebonian accounts. However, we will also see that open texture is just one tool in

the toolbox for the contemporary undertaking of conceptual engineering, introduced in section 3.6. Indeed, it has been argued in (Corfield 1997) that Lakatos fails to appreciate the amount of formal and axiomatic work that is involved in modern mathematics. In response to this I suggest that we can draw on this new work in conceptual engineering to supplement the Lakatosian picture of concepts. In sections 3.7 and 3.8, I consider two examples of this: Haslanger’s distinction between *manifest* and *operative* concepts and Scharp’s *replacement strategy*, showing how these might apply to the mathematical concepts found in mathematical practice.

3.2 Waismann and Open Texture

Let us begin by investigating what it is for a concept to display open texture, as it is used by Waismann in (Waismann 1968) and Shapiro in (Shapiro 2006).¹ It should be noted that although Shapiro follows Waismann in general, the exact characterisation of open texture does shift somewhat between them. I shall set out the two definitions and see how closely they agree. The main difference between the two seems to come down to the difference between the potential for *sharpening* concepts versus the potential for *extending the domain* of the concepts.

Waismann introduces the term ‘open texture’ through examples such as the following regarding our concept of *cat*:

What, for instance, should I say when that creature later on grew to gigantic size? Or if it showed some queer behaviour usually not to be found with cats, say, if, under certain conditions, it could be revived from death whereas normal cats could not? Shall I, in such a case, say that a new species has come into being? Or that it was a cat with extraordinary properties?
(Waismann 1968, p. 119)

Further examples include people who show the unusual feature of disappearing and re-appearing, someone who is old enough to remember King Darius and a lump of gold which emits a new kind of radiation. The point is that in such surprising situations the concepts we possess do not settle whether

¹Shapiro’s book has an appendix focussed entirely on Waismann’s account of open texture (Shapiro 2006, pp. 210-215).

these are just unusual manifestations of the familiar concept or the appearance of some new thing entirely. Waismann links this up to two crucial points: (1) that our concepts are only delimited in some possible directions and not others, and (2) what he calls “the *essential incompleteness* of an empirical description” (Waismann 1968, p. 121).

Waismann never gives an explicit definition of open texture², but the relevant property found in the examples appears to be that of (1): that the concepts we deploy and use to understand the world around us are not delimited in all possible ways. Indeed, let us take this as the central component of Waismann’s account of open texture to give the following definition:

Open Texture 1 (OT1) A concept or term displays open texture iff there are possible objects falling outside of the standard domain of application for which there is no fact of the matter as to whether they fall under the concept or not.

In other words, Waismann’s idea of open texture concerns the potential to expand concepts to new or larger domains. Concepts generally do have a standard domain of application, a range of practical situations in which we know which objects do and do not fall under that concept. Picking Waismann’s examples, we are usually good at identifying and agreeing on which things are and aren’t cats (in contrast with any other pets we might see), humans (in contrast to shop dummies) and lumps of gold (in contrast to other metals).³ However, these normal domains do not cover all potential applications and his examples provide cases where we are asked to apply the concepts beyond the standard situations, and there is no fact of the matter as to whether the objects in question should fall under the relevant concepts or not. Shapiro puts it as follows:

²It may be that open texture itself is the sort of thing that avoids full specification, such that the concept of open texture is itself open-textured. Such considerations are familiar from the literature on vagueness. The connection between open texture and vagueness will be discussed shortly.

³Note the contrastivist spirit of how I have described this agreement on the standard domains of application and disapplication. Waismann doesn’t have an explicit story of how we pick out these domains, but seems to favour empirical descriptions. These don’t seem incompatible if the empirical descriptions are only fine enough to separate out particular examples from others. The contrastivist pull also seems to be present in the Bartha quote below. On the other hand, we can just read the contrast here as providing cases of application and disapplication which are clear in practice.

We language users introduce terms to apply to certain objects or kinds of objects, and, of course, the terms are supposed to fail to apply to certain objects or kinds of objects. As we introduce the terms, and use them in practice, we cannot be sure that every possible situation is covered. (Shapiro 2006, p. 210)

A good way to think about this is presented by Bartha in terms of hard and easy cases:

Typically, the decision about whether the predicate applies involves a nontrivial comparison to paradigm cases or prototypes. There will be “easy cases” where there is general agreement that the predicate does or does not apply, and “hard cases” where applicability is open to debate. (Bartha 2010, p. 9)

Importantly, open texture does not require that we have already identified such areas of openness, only that it is possible that there are such areas. The mere potential for new situations to arise in which we can question how the concepts are applied is enough. We do not need to know in advance which kinds of questions will push us outside of the standard domain of application, and in fact we usually don’t know this in advance and usually don’t predict the difficult situations that might arise for our application of concepts.

Shapiro’s account of open texture is slightly different. In (Shapiro 2006), Shapiro is arguing for an account of vagueness which incorporates a strong flavour of open texture. However, we can observe that the definition he uses is not quite the same:

Suppose, again, that a is a borderline case of P . I take it as another premise that, in at least some situations, a speaker is free to assert Pa and free to assert $\neg Pa$, without offending against the meanings of the terms, or against any other rule of language use. Unsettled entails open. The rules of language use, as they are fixed by what we say and do, allow someone to go either way. Let us call this the *open-texture* thesis. (Shapiro 2006, p. 10)

We may take the characterisation here to be definitional of open texture in the sense he uses it:

Open Texture 2 (OT2) A concept or term displays open texture iff there are cases for which a competent, rational agent may acceptably assert either that the concept applies or that it disapplies.

While Waismann’s account of open texture is about the extent to which a concept is delimited, in contrast Shapiro’s definition is about competent language users being able to go either way on whether or not a term should be applied or a proposition asserted. What seems to be driving this definition in the text, though, is that the open texture in OT2 will be applicable to borderline cases, thus also includes the potential for competent agents to go about deciding one way or another and sharpening the concepts. The difference becomes less marked, however, if we realise that on Waismann’s account, in going beyond a concept’s standard domain a competent speaker can go either way on the application of the term corresponding to the open-textured concept, as we have on Shapiro’s definition.

The main difference between the two definitions is with respect to vagueness, through which we see that Shapiro’s notion is somewhat broader in a certain sense. On Shapiro’s account borderline cases of vague terms can count as open-textured because we can sometimes go either way on whether the term applies or not, while on Waismann’s definition these will still fall within the standard domain of application (even if they might be hard to decide) so wouldn’t count as cases of open texture. For example, taking a paradigm case of vagueness such as *bald*, there will be some borderline cases in which it isn’t clear whether the head in question is bald or not, satisfying the second definition, nonetheless if the person whose hairline is under discussion is a normal one then this will be within the standard domain of application and disapplication of the term ‘bald’ so the first definition will not be satisfied. However, the term ‘bald’ will still be open-textured on the first definition because of the possibility of other hard cases not previously considered, say a two-headed person—maybe where one head has no hair and the other has some— how do we apply a term like ‘bald’ then? An interesting upshot, then, will be that the two definitions seem to agree on the extension of which concepts count as open-textured, because any concepts which permit the vagueness Shapiro is interested in where competent users are allowed to settle borderline cases either way (thereby satisfying OT2) will happen to have cases outwith their standard domain of application for which there is no fact of the matter (thus also satisfying OT1) by

the ubiquity of the latter phenomenon. Similarly, if there is no fact of the matter regarding these non-standard cases, it would be strange to insist that a rational agent could not be free to settle it either way if we do choose to extend the cases in this direction.⁴

Nonetheless, there is clearly a difference in the kind of phenomenon the two notions are intended to capture. Shapiro's definition is aimed at the broader idea which incorporates both open texture in Waismann's sense and vagueness, by focusing on the openness in the agent's settling of the hard cases. In contrast, Waismann is more concerned with the openness of new cases and potential applications which the concepts might be put to. Indeed, Waismann does compare open texture to vagueness, describing open texture as "something like *possibility of vagueness*" (Waismann 1968, p. 120), which is explained as follows:

...a term like 'gold', though its actual use may not be vague, is non-exhaustive or of an open texture in that we can never fill up all the possible gaps through which a doubt may seep in.
(Waismann 1968, p. 120)

As Shapiro points out (Shapiro 2006, p. 211), Waismann does not say much more concerning vagueness and certainly doesn't offer an account. However, I believe that I have here identified the source of the differences in the two definitions and their intended phenomena: Shapiro is interested in the sharpening of terms, with the openness to decide either way, while Waismann is concerned about how our terms apply to entirely new cases which might arise.

As a final component in setting out the nature of open texture and Waismann's definition of it, let us return to point (2) left aside above, concerning the "the *essential incompleteness* of an empirical description". Indeed, I think this will be useful in elaborating on what is meant by a domain of application. The main idea concerns how we describe and define our concepts or terms⁵, where Waismann says:

⁴Although it might well be that there are other demands on a rational agent which push them one way or the other, changing the picture I am sketching here. For example, one choice might be rationally preferable over another to remain consistent with the way other concepts have been settled.

⁵One thing I have not yet discussed is the difference between a concept and a term, i.e. whether these are meant to be mental, linguistic or something else. Clearly, Shapiro's notion leans far more towards the linguistic side and thus towards ascribing open texture

A term is defined when the sort of situation is described in which it is to be used. (Waismann 1968, p. 122)

Waismann's idea is that the difference between complete and incomplete descriptions and definitions is crucial. A complete description exhausts all the details of its subject and a complete definition "anticipates and settles once for all every possible question of usage" (Waismann 1968, p. 122). Conversely, an incomplete description can be extended with more details that haven't previously been mentioned and an incomplete definition fails to anticipate and settle usage. Waismann's point is that most descriptions of things we find in the world will be incomplete, but even further this incompleteness will not be settled as we add further details:

[H]owever far I go, I shall never reach a point where my description will be completed: logically speaking, it is always possible to extend my description by adding some detail or other. Every description stretches, as it were, into a horizon of open possibilities: however far I go, I shall always carry this horizon with me. (Waismann 1968, p. 122)

Similarly, according to Waismann a definition will not be able to anticipate all eventualities meaning "the process of defining and refining an idea will go on without ever reaching a final stage" (Waismann 1968, p. 123). A standard domain of application and disapplication, then, might just pick out those cases which a definition does anticipate and straightforwardly applies to, which is just to say: Bartha's "easy cases".

Let us briefly return to an issue footnoted earlier: whether the notion of open texture is itself open-textured. The question is very similar to that of higher-order vagueness, where if we take vagueness to be about having borderline cases, there are further borderlines as to where those first borderlines start and finish. That is to say, the conception of what counts as a borderline case is itself vague. There are several problems pertaining to the fact that it seems that we may be able to iterate the vagueness of borderlines of borderline onwards indefinitely, problems which I shan't cover further but are discussed with regards to Shapiro's open-texture account of vagueness by

to terms. Waismann, however, was unsurprisingly less clear on this matter, switching between concept-talk and term-talk. This is not something I intend to resolve here, nor do I think much hangs on it for my discussion.

Greenough in (Greenough 2005). With open texture being defined in both OT1 and OT2 using other terms which are not given full specifications, I believe that these terms are prone to open texture and thereby open texture itself is too. For instance, OT1 refers to domains and, indeed, concepts and it is not hard to suppose that unexpected cases and interpretations may arise for these. Even more so, OT2 is in terms of rational, competent agents, and these will certainly face new cases.⁶ One could even claim that the slight shift in the intended usage of open-texture from Waismann's view to Shapiro's demonstrates an instance of the expansion of the domain to include vagueness as a kind of open texture.

To finish, given that we are going to be investigating *mathematical* concepts, it is interesting to see that Waismann's main comparison for the essentially incomplete empirical terms are the concepts found in mathematics. For example, he takes the description of a triangle by giving the side lengths to be complete and discusses enumerating all possible situations in chess to leave no room for new possibilities to emerge. Put explicitly:

Goldbach's hypothesis [...] may be undecidable [...] But this in no way detracts from the *closed* texture of the mathematical concepts. If there is no such thing as the (always present) possibility of the emergence of something new, there could be nothing like the open texture of concepts. (Waismann 1968, pp. 123-124)

But I believe Waismann is wrong on both the claims that mathematical concepts are closed-textured and that there is no possibility in the mathematical case of new possibilities arising. Let us turn to Lakatos to see how this might be the case.

3.3 Lakatos and *Proofs & Refutations*

Imre Lakatos's *Proofs & Refutations* was published posthumously in 1976 as a book edited by John Worrall and Elie Zahar, including additional work developed from his 1961 PhD Thesis. Previous to this the dialogue sections had also appeared as a series of four articles as (Lakatos 1963a,b,c,d). In the

⁶It is already not uncommon to hear computers discussed as a case outwith the standard domain of application, where we might have reasons to count them as agents in relevant respects.

main parts of the text, Lakatos presents a classroom dialogue of a teacher and students named after letters of the Greek alphabet, discussing various proofs of Euler’s conjecture for polyhedra, which links the vertices (V), edges (E) and faces (F) by the formula $V - E + F = 2$. The brilliance lies in the fact that the dialogue is also a *rational reconstruction* of the historical development of the proofs, counter-examples and concepts involved, where the various students are used as representatives of various historical positions and reactions. Sometimes this is in a very direct sense, with footnotes to the text indicating that the words of the students are quotes from various mathematicians. All of the main ideas that Lakatos presents are important to the themes of this thesis, so it is worth going through them in some detail.⁷ In particular, for the current setting, the account of mathematical concepts and their dialectical development will provide us with the main version of the view which takes mathematical concepts to change over time rather than being timeless and immutable.

3.3.1 A Briefing on the Proof and Responses to Counterexamples

Let us spend a moment on the mathematical case study which Lakatos uses. The dialogue opens with a statement of the problem: that all polyhedra satisfy the formula $V - E + F = 2$, followed by a purported proof⁸ (originally stemming from Cauchy) which proceeds roughly by the following method.

1. Imagine a polyhedron to be made of a thin surface of rubber. Then we can remove one face and *stretch* the remainder flat onto a surface, for which we then must show $V - E + F = 1$ (having removed one face).
2. *Triangulate* the flat network that the previous step delivered. That is, for any face that is not already a triangle, we add diagonals and keep doing so until all faces are triangulated. Since we add an edge and a face for each diagonal, the equation isn’t affected.
3. *Remove* triangles from the triangulated network one by one. For this either we remove an edge and a face, or we remove two edges, a vertex

⁷It should be noted, though, that I am just focusing on *Proofs & Refutations*, so will not spend time on Lakatos’s other papers on mathematics or the philosophy of science.

⁸One of the big points concerns what proofs amount to, so the exact status of the various proofs doesn’t divide neatly into correct proofs and merely purported ones.

and a face. In both cases, the equation $V - E + F = 1$ is unaffected.

4. Ultimately, we are left with a single triangle, for which $V - E + F = 3 - 3 + 1 = 1$ as required.

Despite indicating that many historical figures were fully convinced by the above as a proof of Euler's conjecture (Lakatos 1976, p. 8, fn. 2), the description of the method as a 'proof' does not last long in the dialogue. Immediately the students are suspicious of all three main steps in the method. For example, the third step fails if one removes a triangle from the inside of the triangulated network, and furthermore can fail if we choose the wrong order of removing faces, as pointed out by student Gamma (Lakatos 1976, p. 11).

Prompted by the different kinds of problems that arise, the Teacher distinguishes between *local counterexamples* and *global counterexamples*. Global counterexamples present a counterexample to the theorem, while local counterexamples demonstrate a flaw in a lemma. For instance, the cases just mentioned which are brought up by Gamma are local counterexamples to the third step in the proof. Further examples appear of polyhedra which can't always be stretched flat after removing a face, such as the nested cubes, or can't necessarily be triangulated, such as the crested cube with its ring-shaped face, which therefore are local counterexamples to steps 1 and 2 respectively. For the nested cubes and the crested cube, though, the equations are $V - E + F = 4$ and $V - E + F = 3$ respectively, meaning they are also global counterexamples to Euler's conjecture.

Lakatos reviews numerous possible responses to the problematic cases that arise, again made concrete by the footnotes revealing that these were all positions taken by historical figures. Let us review the responses in turn with some comments on how well they work.

Method 1: The Method of Surrender

A single counterexample refutes the conjecture as effectively as ten. The conjecture and its proof have completely misfired. Hands up! You have to surrender. —*Gamma* (Lakatos 1976, p. 13)

The response offered by Gamma is passed over rather quickly, mainly because it offers very little. The idea is that if a global counterexample appears

to a conjecture, then the conjecture is false and must be abandoned. Such a response might be appropriate in some cases, and might seem natural to modern mathematicians, but the response is too extreme in that it also is clear that we are unlikely to hit on the right answers on our first attempt, and yet the ways in which we go wrong can be informative in improving and refining the key ideas, something which is a main point that Lakatos is making.

Method 2: The Method of Monster-Barring

It is the ‘criticism’ that should retreat. It is a fake criticism. The pair of nested cubes is not a polyhedron at all. It is a *monster*, a pathological case, not a counterexample. —*Delta* (Lakatos 1976, p. 14)

The response espoused by Delta is that of monster-barring, in which global counterexamples are themselves rejected as not being genuine instances of the key concepts, and therefore should be rejected as counterexamples too. For instance, in the text numerous polyhedra are proposed which are counterexamples to the conjecture, but Delta offers new definitions of the concept of ‘polyhedron’ which each add additional necessary conditions which serve to rule out the tricky cases and thereby the counterexamples. This is aptly mocked in the text with the ‘definition’ of a polyhedron as anything which satisfies the conjecture (Lakatos 1976, p. 16), revealing the *ad hoc* nature of ruling out problematic cases. In essence, the notion of a polyhedron “defines the domain of application” (Martin & Pease 2013, p. 102) for the conjecture and thus deciding on what counts as a polyhedron can be employed to decide whether the theorem will remain valid or not. Initially, the stream of new definitions given by Delta are seen as contracting the concept of polyhedron to rule out the many examples that have already been encountered⁹, but monster-barring is given a second hearing later in the book, where the student Pi suggests that it may be that Delta was merely scrambling for definitions to defend a view of what counted as a polyhedron which had seemed natural but turned out to be naïve:

Let us go back to the time of the first explorers of our subject.

⁹Despite protestations from Delta: “I do not *contract* concepts. It is you who *expand* them.” (Lakatos 1976, p. 21)

[...] For the polyhedra they had in mind, the conjecture *was* true as it stood and the proof was flawless. Then came the refutationists. In their critical zeal they stretched the concept of polyhedron, to cover objects that were alien to the intended interpretation. —*Pi* (Lakatos 1976, p. 84)

This discussion of what delimits concepts and their application will be relevant to the notion of open texture when we return to it later.

Method 3: The Method of Exception-Barring

No conjecture is generally valid, but only valid in a certain domain that excludes *exceptions*. —*Beta* (Lakatos 1976, p. 24)

The method here is to endorse the positive aspects of monster-barring, namely to give a precise characterisation of the domain of validity for the conjecture, while avoiding the endless moves just intended to not have to give up the theorem to counterexamples. Of course, the precision with which the domain can be picked out is quickly undermined for a distinctly Waismannian reason: the Teacher points out the possibility of new cases arising which show that the excluded exceptions might not be the only problematic cases. Under pressure, Beta is left to retreat into smaller and smaller domains while still failing to arrive at any confidence that these domains are necessarily safe, i.e. that the theorem will be true for all cases there.

Method 4: The Method of Monster-Adjustment

Monsters don't exist, only monstrous interpretations. One has to purge one's mind from perverted illusions, one has to learn how to see and how to define correctly what one sees. —*Rho* (Lakatos 1976, p. 31)

The method of monster-adjustment works not by rejecting the global counterexamples, but instead by changing the way they are interpreted. In particular, this can work by changing our understanding of sub-concepts. The motivating case in the text is how to interpret the small stellated dodecahedron. In the text this shape is introduced as having twelve star-pentagon faces, twelve vertices only at the tips and thirty edges, whereby $V - E + F = 12 - 30 + 12 = -6$, which means that the stellated dodecahedron

does not satisfy the Euler conjecture. However, Rho’s suggestion is to see the faces instead as the individual triangles, thereby reinterpreting the shape as a “triangular hexacontaeder”. On this reading, there are sixty triangular faces, vertices joining them all to a total of thirty two, and edges on each of the triangular faces adding up to ninety. Now $V - E + F = 32 - 90 + 60 = 2$ satisfying the Euler conjecture.

What is interesting in what the method of monster-adjustment brings out is the deep way in which concepts in mathematics are interrelated. Euler’s conjecture does not merely rely on the concept of a polyhedron, but also on a range of related and interrelated concepts, especially the concepts of face, edge and vertex. Changing the interpretation of these will result in a different outcome for the Euler formula. While the change of interpretation in the one case considered might be a good one or not, the reason that it is not taken up much beyond its initial proposal seems to be that it is severely limited in scope.¹⁰

Method 5: The Method of Lemma-Incorporation

I build the very same lemma which was refuted by counterexample *into* the conjecture, so that I have to spot it and formulate it as precisely as possible, on the basis of a careful analysis of the proof. —*Teacher* (Lakatos 1976, p. 36)

The idea here is that many global counterexamples will also be local counterexamples to certain lemmas in the proof. For example, many of the problematic examples for polyhedron fail on the first step of the proof of flattening out the polyhedron minus one face onto a surface, so the Teacher’s method is to restrict the domain of validity for the theorem to the domain where the first lemma does indeed hold. The difference from exception-barring, then, is that exception-barring is *ad hoc* in ruling out examples as they arise, whereas lemma-incorporation is intended to achieve similar ends but by analysing what went wrong in the *proof* and adding restrictions to improve it. Thus lemma-incorporation has a reason for the exceptions it

¹⁰The method of monster-adjustment is brought up again at (Lakatos 1976, pp. 38-39, fn. 2). Here it concerns the positing of hidden faces and edges to show that counterexamples are Eulerian after all. Lakatos is not impressed and rightly so: if one can posit these willy-nilly then nearly any proposed polyhedron can be adjusted to fit the Euler conjecture.

introduces: these are the restrictions generated by examining why and how local counterexamples undermine the proof steps and how to avoid them.

On the downside, lemma-incorporation is seen as still too tied to exception-barring methods. The lemmas one builds in might be sufficient to mark out a safe domain, but still this is not guaranteed, as is evidenced by the Teacher needing to backtrack on the exact lemmas incorporated into Euler's conjecture. Furthermore, the listing of proof-generated restrictions without a record of why they are given is the main feature of the Deductivist approach which Lakatos criticises, something we will return to later.

Method 6: The Method of Proof(s) and Refutations

[O]ne cannot put proofs and refutations into separate compartments. —*Lambda* (Lakatos 1976, p. 49)

The method of proof and refutations (soon re-dubbed to the titular method of proofs and refutations) is the final methodology emerging from the others, and the main approach advocated by Lakatos. This method, according to Lakatos, encapsulates the dialectical development of mathematical concepts, proofs and theories. The method proceeds by four stages (see Lakatos 1976, p. 127):

1. A naïve conjecture.
2. A 'proof' is offered for the conjecture, where 'proof' is not a success term but instead represents a "rough thought-experiment" where the naïve conjecture is decomposed into a series of lemmas or subconjectures.
3. Global counterexamples are found.
4. 'Proof re-examined': the global counterexample is examined and found to also be a local counterexample to particular lemmas in the original proof. A re-examination of the problematic lemmas leads to incorporating restrictions into the conjecture, or a development of the concepts being deployed (to *proof-generated concepts*) which can then be used to state an improved conjecture-and-proof pair.

Of course, the central example of Euler's conjecture presents a rational reconstruction and Lakatos points out that therefore this is not an exact pattern that will be followed in general. The third and fourth stages might

be bound up more closely, such as if a careful proof-analysis uncovers the counterexamples. The responses in the fourth item show a fork in the road, where several mathematically interesting things might happen, ranging from a simple restriction on the domain of validity to the creation of a new concept through the proof-analysis. In many ways this method is the culmination of the lessons learned from the various other methods considered, taking on their positive attributes as part of the heuristic, while minimising their various problems. While the four stages describe the main idea of this method, further discussion reveals that there are additional important stages that often feature in the analysis of proofs with an end to mathematical discovery and development. These are (see Lakatos 1976, p. 128):

5. Checking related proofs of other theorems to see if any of the new concepts or lemmas occur in them, which might show that these are of additional importance.¹¹
6. Checking accepted consequences of the original conjecture.
7. Converting counterexamples into new examples and thereby revealing new fields of inquiry.

These additional steps broaden out the scope of where the process is taking place, rather than being limited to a particular conjecture, proof and their various developments, instead these additional stages acknowledge that the conjectures and theorems are not isolated from surrounding mathematics.

The link between the method of proofs and refutations and the development of concepts is what we shall turn to next.

3.3.2 Concepts and a Dialectical Philosophy of Mathematics

One of the major themes of *Proofs and Refutations*, is the development and change of mathematical concepts, as well as the close connection this has to practices of proving.

A particularly important notion for conceptual change in Lakatos is that of the *heuristic counterexample*. These serve not to refute the theorem as logically false, but to show that it falls short in its scope. In particular,

¹¹This stage is shown clearly at work in Lakatos's description of the discovery of uniform convergence in (Lakatos 1976, Appx. I).

these may be cases of relevant examples which are not covered by the theorem or concepts as they stand. In the central example, the initial proof only applied to a very limited class of polyhedra and so many of the counterexamples initially appeared to be logical ones. However, in the reconsideration of monster-barring discussed above, it emerges that the initial concept of polyhedron may not have been intended to apply to the broad class of examples that were raised in the dialogue. As such, these were instances of heuristic counterexamples, in that they showed the limited scope of the notions involved in the proof. What we shall return to in section 3.5 is the echoes of Waismann in this: one can see heuristic counterexamples as falling outside of the standard domains of application and disapplication, showing thereby the poverty of the concepts with respect to the larger domain.

Initially, the class begins with poorly delineated, naïve concepts of things such as polyhedron, edge, vertex and face. This comes out particularly clearly, for instance, in the difference between treating polyhedra as hollow surfaces and treating them as solid objects, which both appear early on. As the students examine more problems with the initial proof and counterexamples to the initial conjecture, they start to change their concepts in response to the challenges that appear. Lakatos's term for the concepts that emerge is *proof-generated concepts*. Proof-generated concepts are those concepts that arise from the naïve concepts we begin with through the method of proofs and refutations. Lakatos is careful to point out, though, that it would be inaccurate to see them as simply specifications, generalisations, expansions or the like of the naïve concepts:

The impact of proofs and refutations on naive concepts is much more revolutionary than that: they *erase* the crucial naive concepts completely and *replace* them by proof-generated concepts.
—*Pi* (Lakatos 1976, pp. 89–90)

So what has gone on is that the method of proofs and refutations involves the identification of (1) flawed or limited lemmas and (2) what has gone wrong in particular instances that have been encountered, drawing out the problems in them and incorporating additional conditions to give a new proof which is not subject to the counterexample. However, through repeated iterations of the cycle of proofs and refutations, the fresh conditions can form into new concepts. As Larvor puts it:

If this pattern is repeated sufficiently often, these conditions may accumulate to the point where they collectively define a new concept. (Larvor 1998, p. 13)

In fact, it seems that what has gone on in the dialogue is something more intricate than this. As was noted in describing monster-adjustment and discussed by Larvor immediately after the above quote, mathematical concepts are *deeply intertwined*, such as in the relations between ‘polyhedron’, ‘face’ and ‘edge’. It is not at all clear that you can change one without changing the other related concepts. For example, in the case of the ‘urchin’ (that is, the stellated dodecahedron) the difference between seeing it as having star-pentagram faces as opposed to triangular faces is fully intertwined with how one counts the edges, i.e. whether one sees each face as having five or ten edges. But this does not seem to necessarily fall out of simply putting together new conditions as they appear through proof-analysis, without some associated conceptual analysis.

Of course, that an analysis of the proofs is required to see the place of the concepts within them, and to generate new ones, is just part of Lakatos’s very idea. For one thing, the listed rules of the method of proofs and refutations are not strict rules but instead an attempt to induct the reader into a different methodology for the philosophy of mathematics. Observe the difference between the two summaries given of the method of proofs and refutations. The first description of the method of proofs and refutations (Lakatos 1976, p. 50) is given as a list of general *rules*, putting them in an imperitival form stating how one should go about applying the method. The later summary of the method (which I have used above) is a description of the *stages* of the pattern of mathematical discovery. Clearly these two descriptions are not the same¹², and the reason is that neither an overly-specific set of rules nor a description of the patterns of mathematical discovery are meant to capture the methodology in full. Rather the method of proofs and refutations is part of Lakatos’s broader project of offering a *dialectical philosophy of mathematics*.

Already this provides an answer to a particular criticism offered by Ferman in his response to Lakatos:

¹²Most papers that describe the method of proofs and refutations seem to ignore this fact, simply setting out one of the descriptions. e.g. (Corfield 1997, p. 100). Ernest does better, setting out only one description but observing the different functions that the method plays (Ernest 1997, pp. 117–118).

A related question is whether the method of proofs and refutations is supposed to be *descriptive* or *normative*. It seems at best that it could be descriptive of progress since 1847. But much of the tenor of the discussion leads one to view it as normative. (Feferman 1978, p. 316)¹³

Of course there is the potential for confusion here. The first description of the method of proofs and refutations in terms of rules seems to fit a normative reading, while casting it in the light of stages in a pattern of mathematical discovery seems to fall into a descriptive reading. The answer is, I believe, that this is a false dichotomy. For one thing, the dialogue is a rational reconstruction, so is intended as neither an entirely accurate description¹⁴ nor simply to declare how to go about discovering new mathematical theorems, concepts and proofs. Rather, it is meant partially as a demonstration that mathematical discovery is a rational process predominantly driven by mathematical ends. Certainly, this involves both normative and descriptive elements, but Feferman is wrong to expect it to fall neatly into one or the other.

Dialectical philosophy of mathematics has several major features which distinguish it from the more traditional approaches, as described in (Larvor 2001). Larvor sets out the following key points of what such a dialectical philosophy of mathematics involves on the Lakatosian programme:

Internal Stance: To see that “changes in the body of mathematics normally take place for mathematical reasons” (Larvor 2001, p. 215) While this is certainly defeasible (mathematical changes can take place due to many other factors), the idea is to see that the development of mathematics follows a rational pattern, rather than one marked by arbitrary decision or spontaneous insights.

Human Minds in Mathematics: “Dialectical philosophy [...] typically recognises that human minds, however fallible, are the only available vehicles for the greater rationality of science.” (Larvor 2001, p. 215)

¹³The year 1847 is significant as the year Lakatos mentions the method of proofs and refutations to have been discovered by Seidel, (see Lakatos 1976, p. 131, p. 139). I have changed the emphasis in the quote here from underlining to italics.

¹⁴Lakatosian wit tells us a little about his view on the relationship between history and its reconstruction: “Pi’s statement, although heuristically correct (i.e. true in a rational history of mathematics) is historically false. (This should not worry us: actual history is frequently a caricature of its rational reconstructions.)” (Lakatos 1976, p. 84)

Concepts Over Propositions: To be interested in the nature and development of concepts, ahead of the propositions that involve them. Larvor notes this to be something Lakatos follows Hegel on.

Dialectical and Formal Logic: The previous item involves a split between dialectical and formal logic, where dialectical argument develops concepts while formal logic keeps them fixed to avoid fallacious equivocation.

Different Notions of Rigour: Another idea rooted in Hegel, that the “formal and dialectical logics have different aims and incompatible standards of rigour; so we ought not to mix them up.” (Larvor 2001, p. 217).

Ontological Neutrality: dialectical philosophy of mathematics has nothing to say on the ultimate ontological and metaphysical basis of mathematics, since the methodology makes neither assumptions concerning them nor does it have any way of deciding between rival accounts.

All of these are, of course, evident in Lakatos’s work towards building this dialectical school of philosophy of mathematics. The “fundamental dialectical unity of proof and refutation” (Lakatos 1976, p. 37) is at play precisely because this dialectical approach is to be found in the method of proofs and refutations. After all, the subtitle of the book is that of “The Logic of Mathematical Discovery” and in the work we find that this logic is a dialectical one too, besides the formal logic which dominates traditional philosophy of mathematics. The existence of heuristic counterexamples in mathematics demonstrates the need for a rigorous method for dealing with them and the conceptual change that they bring about.

3.3.3 Fallibilism and Formality

Let us now turn to another aspect of Lakatos’s philosophy: fallibilism about mathematical knowledge. In the introduction, Lakatos presents his work as a new step in the battle for certainty:

In this great debate [between sceptics and dogmatists], in which arguments are time and again brought up to date, mathematics has been the proud fortress of dogmatism. Whenever mathematical dogmatism of the day got into a ‘crisis’, a new version once

again provided genuine rigour and foundations, thereby restoring the image of authoritative, infallible, irrefutable mathematics [...] Most sceptics resigned themselves to the impregnability of this stronghold of dogmatist epistemology. A challenge is now overdue. (Lakatos 1976, p. 5)

The dogmatist position in mathematics which is Lakatos's (indirect) target is that of the *formalist*, or in keeping with our terminology, the *Formalist-Reductionist*. As always, the Formalist-Reductionist takes pieces of knowledge of mathematical propositions to be certain and infallible as a result of formal, logical proofs. The advent of modern logic brings with it the 'fortress' of dogmatism, settling what it is for a proof to be immune to doubt, error and leaps of reasoning. Closely related to the Formalist-Reductionist position is the *deductivist approach*, discussed in the second appendix if *Proofs & Refutations*.

In deductivist style, all propositions are true and all inferences are valid. Mathematics is presented as an ever-increasing set of eternal, immutable truths. Counterexamples, refutations, criticism cannot possibly enter. An authoritarian air is secured for the subject with disguised monster-barring and proof-generated definitions and with the fully-fledged theorem, and by suppressing the primitive conjecture, the refutations, and the criticism of the proof. Deductivist style hides the struggle, hides the adventure. (Lakatos 1976, p. 142)

The deductivist style thus also encompasses a presentational style of defining concepts, setting out theorems and providing proofs, which separates these from the dialectical environment in which they were created. Indeed, Lakatos suggests that mathematics would be greatly improved if we presented mathematics in the heuristic style which makes explicit the growth of the theorem, proof and concepts involved.

The challenges to the Formalist-Reductionist, the deductivist and the dogmatist, then, come from presenting an alternative picture of mathematics and its development. By making the dialectical logic of mathematical discovery clear, it demonstrates that mathematics cannot be equated with its formal shadow, that mathematical concepts change and emerge from mathematical practices, and that the status of mathematical theorems will change

as they and the concepts they feature change. In a way, this deflates some of the radicalism that talk of fallibilism—and indeed scepticism— might suggest, for the sense in which mathematical statements are fallible is merely as the by-product of the idea that concepts and theories are changeable through dialectical logic. Meanwhile, this doesn't cast doubt in a more radical sense, as Larvor puts it:

[I]t has been and will remain the case that an apple taken together with two oranges makes three pieces of fruit. (Larvor 1998, p. 36)

The point is more subtle than the radical sense of fallibilism. Larvor continues:

It is also the case that an apple released in mid-air will fall to earth. Nevertheless, in both cases the theoretical apparatus we use to describe and to account for the phenomenon is highly complex and open to criticism. (Larvor 1998, p. 36)

Ultimately, then, the sense of fallibilism which we find in Lakatos is about the fact that we don't find certainty in having established some mathematical theory because that theory will always be open to potential change, revision and development in light of new counterexamples, new ideas and new mathematics. One might even see the major part of the project as being the attempt to show that this does not collapse into pure subjectivity and that there are good mathematical criteria for these kinds of dialectical changes.

Moving on for the moment, what is of particular interest to consider for our purposes is the distinction between formal and informal proofs in Lakatos's picture. In arguing against 'formalism' and 'deductivism' so widely, there is a point of view which is natural to Lakatos: that informal proofs are the central method of mathematical demonstration, with their own associated notion of rigour operative in differentiating correct and incorrect proofs. While at times controversies do arise, these are beneficial for mathematics in that they drive the development of proofs and concepts, as discussed above. Conversely, formal proofs and the purely deductive picture of mathematics cannot underlie the rigorousness of proofs as they appear in mathematics nor can they account for mathematical discovery. In fact, in the second

chapter of *Proofs & Refutations* Lakatos has the student Epsilon go through a separate proof of the Euler conjecture in the ‘Euclidean’ style, where the theorem and concepts are translated into the language of vector algebra:

I analysed our concepts of polyhedra and showed that they are *really* vector algebraic concepts. I translated the circle of ideas of the Euler-phenomenon into vector algebra, thus displaying their essence. Now I am certainly proving a theorem in vector algebra, which is a clear and distinct theory with perfectly known terms, neat and indubitable axioms, and with neat, indubitable proofs.
—*Epsilon* (Lakatos 1976, p. 118)

The class raise several problems for this line of thinking. Firstly there are the standard worries of whether the translation fully captures the informal concepts they are formalising and whether the ‘certain’ system is guaranteed to be consistent.¹⁵ This, of course, relates closely to the discussions of chapters 1 and 2. Moreover, though, the Lakatosian view seems to be that while such a translation does successfully limit the counterexamples that may appear within the theory, it also loses out on much which we had before the translation:

[Y]ou may push out the original problem into the limbo of the history of thought— which in fact you do not want to do. (*foot-note: This process is very characteristic of twentieth-century formalism.*) —*Alpha* (Lakatos 1976, p. 122)

and

Epsilon wanted, “in virtue of a series of startling definitions to save mathematics from the sceptics”, but what he saved was at best some crumbs. —*Gamma* (Lakatos 1976, p. 123)

The standard idea in this direction, then, is that we can formalise any mathematics we choose, but formality is balanced against meaning and so fully formalising leads to theories devoid of meaning, pushing content into the meta-mathematical interpretation.

¹⁵Indeed, Epsilon admits that they must

[...] forget about the old meaning. I create freely the meaning of my terms while scrapping old vague terms. —*Epsilon* (Lakatos 1976, p. 122)

There is some dispute over where Lakatos stands on this exactly, though. On the one hand, there are Davis & Hersh (Davis & Hersh 1981, pp. 345–359) and Larvor in (Larvor 1998, pp. 33–34), offering a reading much like the above, but on the other hand there are Worrall & Zahar, the editors of *Proofs and Refutations*, who insert several substantial footnoted comments into the book offering an alternative reading, backed up by Corfield in (Corfield 1997) who defends their interpretation of Lakatos. In the inserted footnotes, Worrall & Zahar suggest that Lakatos is mistaken about the need to reject the infallibilist idea that “deductive, inferential intuition is infallible” (Lakatos 1976, p. 138), writing in their footnote that:

This passage seems to us mistaken and we have no doubt that Lakatos, who came to have the highest regard for formal deductive logic, would himself have changed it. First order logic has arrived at a characterisation of the validity of an inference which [...] does make valid inference essentially infallible. (Lakatos 1976, fn. 4, p. 138)

Both Davis & Hersh and Larvor see Worrall & Zahar as making mistakes of the precise sort that Lakatos is arguing strongly against. Firstly, Davis & Hersh argue that the mistake is one of conflating mathematical proof with its formal representation as a derivation in some fixed formal system. In essence, they accuse Worrall & Zahar of Formalist-Reductionist thinking, and reject it for several of the classical reasons. Secondly, Larvor argues that their mistake is to focus on language-statics rather than language-dynamics. That is, the fallibility is not of the logical validity in some given system (say, first-order logic) but rather that the counterexamples will be heuristic in a way that might lead to change or abandonment of the system altogether.

Corfield sets out a defence of Worrall & Zahar against Davis & Hersh’s criticism, trying to show that Lakatos was more conservative than they appreciated. The claim is that while they are correct that Lakatos focuses on informal mathematics, and that most mathematics found in practice is informal, in fact Lakatos made a distinction between different levels of informality. The difference lies between informal proofs proper and ‘quasi-formal’ proofs which really are formal proofs with some of the interim steps left out or suppressed. Corfield then says:

Thus, while Davis and Hersh imagine that Lakatos’s account of

informal mathematics extends to present day proofs in established branches of mathematics, surely the majority of the estimated 200,000 produced each year, Lakatos himself wishes to count them as ‘almost formal’ or ‘quasi-formal’. (Corfield 1997, p. 115)

Furthermore, the point is that formal systems, logic and axiomatics do play a major role in modern mathematics and that even in these cases the development of new concepts does not cease, but instead works alongside the axioms. With respect to axiomatics, their role in mathematical discovery is brought out particularly clearly by Schlimm in his case study of lattice theory in (Schlimm 2011).

In general, I think that Corfield’s reading, along with that of Worrall & Zahar, cannot be right. The claim that most modern proofs are quasi-formal rather than informal, and as such are not subject to the Lakatosian arguments, strikes me as entirely wrong. Indeed, the central argument of my first chapter, that of the over-generation of formalisations relative to some given informal proof, stands strongly against such a view. The Formalist-Reductionist line seems to hold that the modern work in formal mathematics supports their view of proofs, as is made explicit by Corfield:

Given the stabilization that has occurred in the idea of what constitutes a rigorous proof, this gives them [Davis & Hersh] little room for manoeuvre against their formalist adversaries. (Corfield 1997, p. 117)

I think the advantage of my argument against such a view is precisely that it draws so directly on work in formalisation, showing that the idea that modern proofs are no longer informal in a strong sense is false, but also delivering the fact that formal mathematics projects are no help in defending against the criticisms of the informalist camp. In addition to this argument, it seems wholly unlikely that the Formalist-Reductionist interpretation of Lakatos is correct when read in the context of the rest of the text. For instance, the opening passage quoted above on the constant retreat of the dogmatist in the face of sceptical challenges, and it being time to finally storm the stronghold of dogmatism in mathematics, makes for a poor fit with the acceptance that actually most modern mathematics is suitably infallible by virtue of being quasi-formal.

Despite believing that Corfield is totally mistaken in following the Worral & Zahar interpretation, I do think there is a point to his argument that cannot be ignored. That point is that the Lakatosian framework does not fit well with the methodology of modern mathematics. Axiomatisations, formalisms and metamathematical results are used alongside traditional, informal proving; the relationship between syntactic proofs and associated model theory is highly complex and fruitful; reverse mathematics offers genuine insights into the provability strength of mathematical statements; and computational mathematics is developing at a phenomenal rate. The point is that while I have been stressing the importance of taking informal mathematics and proof seriously, this has to be taken alongside formal methods rather than instead of them. I think Corfield is right to stress that the strict dichotomy found in the Lakatosian picture between fruitful and contentful informal mathematics and the sterile and static formal theories does not fully do justice to modern mathematical practice. I will come back to this final point shortly as a motivation for investigating whether other work towards conceptual development, conceptual change and conceptual engineering might be fruitfully applied to the case of mathematics. Before then, let us consider another dialectical proposal for mathematics: that of G. T. Kneebone.

3.4 Kneebone on Mathematics

While Lakatos is the leading figure we turn to in defending the kind of dialectical account of proofs and mathematical concepts described above, such as being credited as the source of the maverick tradition in (Kitcher & Aspray 1988), there is another figure who espouses views similar in certain key respects but has received seemingly no attention: G. T. Kneebone. This despite the fact that Kneebone pre-dates Lakatos, writing in the 1950s already.¹⁶ In his papers (Kneebone 1955, 1957), he discusses the relationship between intuitive and rigorous mathematics, the fixity of mathematical concepts, and dialectical versus deductive mathematics.¹⁷ In this section I will briefly run through Kneebone's main views and arguments, bringing out the

¹⁶The plot thickens: in the LSE archives there exists some Kneebone-Lakatos correspondence. Future work ahoy!

¹⁷Despite being two papers, these do have the advantage of being explicit in his views rather than exposing them through the dialogue form that Lakatos uses.

similarity to the Lakatosian position.

The starting point for Kneebone is the consideration of the role of logic in mathematics. Tracing the emphasis on symbolic logic and the formalisation of mathematical reasoning back through mathematical history¹⁸, he places particular emphasis on Peano as seeing the formalisation of mathematics as a move away from intuition towards rigour:

Peano, on the other hand, realized that if complete rigour is to be achieved intuition must be banished completely from mathematical argument. (Kneebone 1957, p. 206)

Peano's project was of major influence on two of the big schools of thought in the philosophy of mathematics: Logicism and Formalism. The former encapsulated by the *Principia Mathematica* project by Russell and Whitehead¹⁹, intending to define all mathematical concepts in terms of explicitly logical ones, reduced mathematical rigour to formal, logical rigour. Formalism, headed up by Hilbert, aimed to formalise all of mathematics in order to allow us to prove that it is consistent. Of course, the failures of both are well-known: the axiom of infinity Russell needed was not purely logical, the various class-theoretic paradoxes stood in the way of a naïve picture of classes and membership, and Gödel's theorems seemed to undermine the possibility of consistency proofs which were more reliable than that which one is proving consistent. From this Kneebone concludes:

The failure of both undertakings suggests that the relationship between mathematics and logic may perhaps have been wrongly understood, and prompts reconsideration of the nature of this relationship. (Kneebone 1957, p. 210)

The answer, according to Kneebone, is to acknowledge the divide between formal logic with its notion of logical validity and *dialectical* reasoning:

¹⁸Although it is worth remarking that Kneebone's picture of history is shaky at best, for example describing Logicism and Formalism as sequential rather than the actual relationship between the two. In fact, he says "Up to this point in the history of the philosophy of mathematics the idea of rigour had developed naturally and smoothly" (Kneebone 1957, p. 207) which seems even further from the truth than Lakatos's rational reconstruction of history. The bumpy history of the development of rigour is covered in (Kleiner 1991).

¹⁹Kneebone sets Frege aside as having had little influence on philosophers besides Russell, which might be true historically of Frege's time, but is another amusing reminder of the age of the paper.

Mathematics, on the face of it, is completely undialectical; but this is only appearance, and it is my thesis in this paper that the basically dialectical character of mathematics is precisely the feature that has been neglected hitherto. (Kneebone 1957, p. 212)

Taking the earlier list of features of dialectical philosophy of mathematics in the vein of Lakatos, as set out by Larvor, we can now run through and see some of these same features at play in Kneebone's work.

To begin, the 'internal stance' of seeing mathematical developments come about for rational, mathematical reasons is certainly not as fully developed as Lakatos's detailed case studies but is implicitly present in Kneebone's consideration of historical examples which cannot be fully explained on the purely deductivist model, such as in his discussion of Kummer's ideal factors. More explicitly, Kneebone acknowledges that mathematics must not be separated from those that practice it, the human agents, but that this gives rise to one of the puzzling aspects of a dialectical logic for mathematics: that it is both personal and impersonal. He says:

[I]n so far as mathematical thinking is dialectical it appears to be both personal and impersonal at the same time. It is personal because thinking is a process that can only take place in the minds of individual mathematicians, and impersonal because it produces a body of mathematical knowledge that is accessible to every individual mathematician, and valid for all alike. (Kneebone 1957, pp. 220–221)

The solution to this puzzle is that the mathematicians who are carrying out the mathematics are embedded in particular cultural traditions, and that concepts within these manage to span the personal and impersonal divide:

The concepts of the cultural tradition are thus part of the mind's equipment and at the same time part of the structure of the known world, and they are able therefore to be both personal and impersonal together. (Kneebone 1957, p. 221)

Next, the emphasis on concepts and conceptual development instead of seeing mathematics as a body of propositions is also undoubtedly central to Kneebone's dialectical philosophy of mathematics:

[W]hereas deductive reasoning operates with fixed concepts, which might be represented by the symbols of a logical calculus, dialectical reasoning always brings about development of concepts. (Kneebone 1957, p. 212)

Mathematics is usually thought of, by philosophers of the subject no less than by those who take it more for granted, as a body of propositions. But although it is true that the mathematics of the textbooks is such a propositional edifice, propositions are by no means all that the creative mathematician is concerned with. (Kneebone 1957, p. 215)

As we have already seen, Kneebone separates out deductive and dialectical aspects of mathematics, but additionally he is clear in the fact that these must come with separate notions of *rigour*:

There is rigour of demonstration and also rigour of dialectical development, and the two are by no means the same. (Kneebone 1957, p. 222)

On Larvor's last point, ontological neutrality, we don't find anything explicit, although the desire to see beyond the limitations of logicism, formalism and intuitionism (which take mathematical ontology to reside in the world, in language and in the mind respectively) to a dialectical form of justification beyond them seems to indicate sentiments in this direction.

While Kneebone's position is covered briefly in just the two papers, from looking at the various aspects of his introduction of a dialectical philosophy of mathematics, we can see that it largely aligns with the aims and positions Lakatos has in engaging in such a project in general. Indeed, often the language the two use is incredibly close, such as concerning the difference between dynamics and statics:

In other words for the purposes of philosophy we have to conceive of [dialectical] rigour in dynamical not in statical terms—as the rigour of a *process* which yields knowledge, not of a system of propositions which summarize a particular *state* of knowledge. (Kneebone 1957, p. 223)

Heuristic is concerned with language-dynamics, while logic is concerned with language-statics. —*Pi* (Lakatos 1976, p. 93)

In actual fact, the importance of the dynamics underlying mathematical activities will be a major point we shall return to in chapter 4.

While the two are close in their positions on many things, there are some notable differences between them, most markedly the role of proofs and refutations. Lakatos's study is deeper precisely because it begins the process of delivering on the investigation of how a dialectical development of mathematical concepts might play out (which is strongly connected to proving practices and the discovery of counterexamples) something which for Kneebone is only the ultimate aspiration. Even then, there is no suggestion that proofs and counterexamples have any particularly important role in the dialectical story for Kneebone. Another point of difference between the two is their attitudes towards formal logic: Kneebone seems content for formal logic to stand supreme as the canon for correct, rigorous deduction (something which I have argued against in chapter 1) while Lakatos seems to reject this.

In summary, it seems that Kneebone managed to slip through the cracks of history. While Lakatos is seen as one of the guiding figures of the philosophy of mathematical practice, Kneebone appears to have missed out on any widespread acknowledgement for his work. Nonetheless, Kneebone's ideas about the need to step towards a dialectical philosophy of mathematics are worth taking seriously alongside Lakatos's.

3.5 The Open Texture of Mathematical Concepts

Having discussed Lakatos at length and then the similar direction found in Kneebone, let us now connect this to Waismann's notion of open texture. Recall that Waismann claimed that mathematical concepts display closed texture, in contrast with empirical concepts which are open-textured. The aim now is to reject Waismann's position here by applying Lakatos and Kneebone's arguments concerning conceptual development in mathematics to show that mathematical concepts do display open texture after all. This has been noted as a central point to be taken from Lakatos by several authors, such as Shapiro in (Shapiro 2006, 2013), Schlimm in (Schlimm 2012) and mentioned by Bartha at (Bartha 2010, p. 10). Let us go through the argument first, then consider what the other authors have to say in turn.

The big point of putting together Waismann and Lakatos is that math-

ematical concepts can display open texture. Heuristic counterexamples fall outside of the domain of application and disapplication as it stands, showing that these are too narrow in that there are interesting and pertinent cases which have not been covered by the accepted definitions. The first definition of open texture in **OT1** is precisely this, so the existence of heuristic counterexamples with respect to some concept in mathematics is enough to show that that concept displays open texture. The second definition in **OT2**, concerns the freedom rational agents have to go either way on the application of the term. With respect to the heuristic counterexamples there does seem to be a rational freedom to choose different responses. While Lakatos does poke fun at the more simplistic implementations of monster-barring, monster-adjustment etc. the ultimate upshot of the dialectical approach is that there is the mathematical freedom to respond and develop the concepts in a way that a purely deductivist approach cannot accommodate.

Waismann takes open texture of concepts to be the result of a lack of full delimitation and definition of when a concept or term applies. The reason he takes this to hold for empirical concepts and not mathematical concepts boils down to this. On the one side, Waismann sees empirical descriptions as never being able to anticipate all potential cases and possibilities that might arise concerning their applications. Recall:

Every description stretches, as it were, into a horizon of open possibilities: however far I go, I shall always carry this horizon with me. (Waismann 1968, p. 122)

On the other side, Waismann suggests that this kind of endless possibility of new cases arising is not applicable to mathematics, with its strict, explicit definitions. But this is where he is mistaken, as demonstrated by Lakatos and Kneebone. Indeed, Kneebone cites Waismann's notion of open texture in (Kneebone 1955, p. 37) as one way in which conceptual development occurs.²⁰ If one restricts oneself to deductive logic and ignores the development

²⁰The exact quote is:

One of the ways in which conceptual evolution is actually brought about has been described by Waismann, who has drawn attention to the open texture of empirical concepts. Such concepts are never finally and completely defined, as concepts can be in pure mathematics. (Kneebone 1955, p. 37)

While this might seem to go against the open texture of mathematical concepts and agree with Waismann on their closed texture instead, the wider context of the quote shows that more is going on. Since this is the earlier paper by Kneebone, he has not fully articulated

of mathematical concepts, ideas and theories then this may sound convincing. However, it seems clear that Lakatos has put this myth to rest; that there are ongoing developments and discoveries which are driven by reasons internal to mathematics and require a dialectical philosophy of mathematics. In light of this, there is little question that mathematical concepts can display open texture just as much as empirical concepts do.

One issue that this ties into is that of the formality of mathematics. Above, I discussed the controversy over Lakatos's attitude to modern mathematics, according to Worrall & Zahar, Davis & Hersh, Larvor and Corfield, where the issue is over whether something like first-order logic captures some infallible inference patterns and thus whether modern mathematics is still subject to the kind of dialectical developments Lakatos and Kneebone describe. Implicit in Waismann's suggestion that mathematical concepts are closed-textured is the perspective on which modern logic has succeeded in pinning down mathematical concepts exactly for all applications.²¹ But on the more natural reading of the core idea of a dialectical philosophy of mathematics, such a perspective presents a mistakenly narrow view of mathematics. We should agree with Corfield that formal systems and axiomatisations both play important roles in mathematics, but as I have argued in the previous chapters of this thesis, even in modern mathematics there are good reasons to want to keep both formal and informal proofs as essential parts of mathematics which play dual roles.²² One important role for informal

the difference between deductive and dialectical theories of mathematics, instead focusing on *concrete* and *abstract* modes of thinking. After mentioning Waismann, he discusses how conceptual development comes about, where he presents an example of the way in which conceptual development occurs particularly through the use of analogy, which takes as its object mathematical concepts:

An example of a different kind, in which the analogy is not merely heuristic, may be seen in the extended use of such geometrical concepts as "point" and "space" in modern mathematics. Here the analogy goes very deep, for what has happened is that a logical structure has been isolated which is exemplified in classical geometry and which can now be seen to pervade the greater part of pure mathematics. It is to this development of the analogical mode of thinking that mathematics owes much of its recent spectacular progress. (Kneebone 1955, p. 40)

Certainly this appears to be more in the direction of his later picture and in agreement with Lakatos.

²¹In the quote at the end of section 3.2, Waismann mentions undecidability, which suggests to me that he held the corresponding Formalist-Reductionist type view of proofs as fixed relative to formal systems.

²²We can now also reframe the problem with Priest's argument for the inconsistency of mathematics. The arguments I offered against this in terms of formalisations being

mathematics more generally is its place in a dialectical process, wherein concepts display open texture and can thus be extended and developed as new ideas arise.

Let us turn now to briefly consider the other authors who have also connected Lakatos's views with the notion of open texture.

In the articles cited, Shapiro is discussing the notion of computability and whether or not we can hope for a mathematical proof of the Church-Turing thesis. We needn't go into the details of this, but the aim Shapiro has is of showing that the mathematical notion of computability has developed in a broadly Lakatosian sense, and that the informal notion of computability that they began with displayed open texture. It is in this context that Shapiro discusses both open texture and the Lakatosian framework, making it explicit that he takes one of Lakatos's main ideas to be that mathematical concepts display open texture. Referring to Lakatos's case study of the Euler conjecture, he says:

[T]he notion of polyhedron exhibited what Waismann calls open-texture. This open-texture did not prevent mathematicians from working with the notion, and proving things about polyhedra. Still, at the time, it simply was not determinate whether a picture frame counts as a polyhedron. (Shapiro 2006, p. 432)

Shapiro also discusses the translation of the main notions such as polyhedron into set-theoretic terminology, such as in Epsilon's vector algebra version of the proof:²³

The student [Epsilon] then gives a fully formal (or at least easily formalizable) proof of a generalization of Euler's theorem from these definitions. The only residual question left, it seems to me, is the extent to which the set-theoretic definition captures the essence of the original, pre-theoretic (or at least pre-formal)

fragmented into different systems which are useful for different purposes, stem from this direction of thought, of formal and informal mathematics playing useful interdependent roles. Meanwhile, the view that all of mathematics can be straightforwardly formalised is Formalist-Reductionist in its approach and misses out on the broader perspective that the dialectical model can offer.

²³Shapiro treats the vector algebra proof as a set-theoretic one. We may suppose for the sake of argument that there is nothing substantial in the differences between these, without conceding that this is ultimately true.

concept of polyhedron. Lakatos had that exactly right. (Shapiro 2013, p. 167)

So Shapiro suggests that the formalised versions of the concepts do not display open texture, and are properly fixed up to certain considerations:

One can perhaps claim, now, that the final, austere and rigorous set-theoretic definition of “polyhedron”—as a set of “vertices,” “edges,” and “faces” under certain conditions—is not subject to open texture. Its boundaries are as determinate as one could wish—assuming that there is no flexibility concerning the logic or the underlying set-theoretic model theory. (Shapiro 2013, p. 168)

There is frequently flexibility over logic and set-theoretic model theory, as Shapiro well knows, so the ultimate security here is hardly unchangeable. However, the point is just that formal and axiomatic systems are powerful tools for fixing meanings within a particular domain and that mathematicians are very good on agreeing on large stretches of mathematics in this way. What is interesting is that Shapiro might be right that even if we encounter new mathematical possibilities that had not previously been considered concerning set theory or its related logic, the open texture seems to be located in the meta-language now and no longer applying to the formalised definitions of the concepts discussed by Lakatos directly.

Shapiro also brings out nicely the connection between Waismann and Lakatos concerning the questions surrounding the identity of concepts. Recall Lakatos quoted above saying:

The impact of proofs and refutations on naive concepts is much more revolutionary than that: they *erase* the crucial naive concepts completely and *replace* them by proof-generated concepts.
—*Pi* (Lakatos 1976, pp. 89–90)

But there is the obvious problem that there are two possibilities (which I have thus far kept in play alongside one another) which are either that the concepts are *replaced* with entirely different concepts, or else that they are *revised*, developed and extended in a way that is consistent with them being fundamentally the same concept. On the latter take, Kneebone says:

There is frequently enough affinity between the concepts involved in a succession of accounts of the same body of experience, or of related bodies, for it to be legitimate to speak of changes undergone by a single enduring concept; and so we have the idea of conceptual evolution. (Kneebone 1955, pp. 38–39)

Shapiro quotes Waismann on the response to this dichotomy which is to mostly not worry about the difference:

[T]here are no precise rules governing the use of words like ‘time,’ ‘pain,’ etc., and that consequently to speak of the ‘meaning’ of a word, and to ask whether it has, or has not changed in meaning, is to operate with too blurred an expression (Waismann 1951, p. 53)

While concepts and expressions are developed in the dialectical accounts they offer of mathematics, whether this ultimately is about replacing concepts or revising them will come down to how coarse or fine the identity conditions are.

Moving on to Schlimm (Schlimm 2012) and Bartha (Bartha 2010), these authors both discuss Lakatos in terms connected to Waismann’s open texture. Schlimm’s central thesis is that we should be pluralists about mathematical concepts, accepting both *Fregean* and *Lakatosian* concepts in mathematics. These are summarised as follows:

According to the first, concepts are definite and fixed; in contrast, according to the second notion they are open and subject to modifications. (Schlimm 2012, p. 128)

The difference here clearly matches the difference as I have been deploying it between Lakatos’s dynamic, dialectical approach and the more traditional, static approach. The characterisation of Lakatosian concepts has them as “subject to modifications” and “fluid” (Schlimm 2012, p. 43), which is not quite open texture by either of the definitions we have seen above. However, Schlimm does discuss the term ‘open-textured’ with reference to Bartha’s work and Bartha, in turn, applies the framework explicitly to Lakatos:

Very similar observations apply to mathematics. There are open-textured mathematical concepts. Lakatos (1976) famously reconstructs the reasoning that leads us to include or exclude

certain objects from the category of regular polyhedra. Here, too, open-textured concepts are the ones under investigation. (Bartha 2010, p. 10)

Bartha's ultimate point is one about the role of analogy in reasoning and argumentation, which we may set aside. Nonetheless, this is another recognition of the close proximity between Waismann's ideas on open texture and Lakatos's ideas on mathematical concepts.

3.6 Conceptual Engineering in Mathematics

Taking stock for a moment, we have now seen that one of the ideas underlying a dialectical philosophy of mathematics, be it Lakatos's or Kneebone's, is that mathematical concepts are not fully fixed and immune to change but instead display open texture, developing in response to new examples, ideas and methods that arise through mathematical practice.

For all that, though, there are two remaining issues for this claim concerning the open texture of mathematical concepts.

The first of these outstanding issues is that there is more going on in the dialectical philosophy of mathematics than the simplistic picture on which we have open-textured mathematical concepts, find new examples which are not covered by them, expand the concept to apply or disapply, repeat. While at times this pattern might be operative, the reason for dedicating so much space to the Lakatosian and Kneebonian frameworks is that these are significantly more broad and more subtle. A better conclusion, in light of this, would be that open texture is one feature of mathematical concepts which participates in the more complex patterns of conceptual change and evolution.

This leads us to the second remaining issue: that even this broader framework does not seem to be enough to give an account of modern mathematics. As discussed above, Corfield has presented the case that a lot of mathematics now is bound up with axiomatic systems and formal provability, in a way not captured by the Lakatosian story of proofs and refutations. In fact, axiomatics, formal derivations, metamathematics and computational mathematics are involved in large parts of the modern mathematical landscape, including discovery, and I fully agree that if we are to adequately account for mathematical concepts we need a picture which also sees the interaction be-

tween informal and formal mathematics, and proofs in particular, as playing an important, rational role in the dialectical development of concepts.

What I propose to do in the remainder of this chapter is to explore some more recent work on the nature of concepts and apply this specifically to mathematical concepts. In particular, I shall draw on some recent work which falls under the general heading of *conceptual engineering*. Conceptual engineering extends beyond the well-known philosophical method of conceptual analysis to include a more active participation in knocking down bad or defective concepts and building new ones as we want them to be. The term seems to come from Blackburn where he says:

I would prefer to introduce myself as doing conceptual engineering. For just as the engineer studies the structure of material things, so the philosopher studies the structure of thought. Understanding the structure involves seeing how parts function and how they interconnect. It means knowing what would happen for better or worse if changes were made. This is what we aim at when we investigate the structures that shape our view of the world. Our concepts or ideas form the mental housing in which we live. We may end up proud of the structures we have built. Or we may believe that they need dismantling and starting afresh. But first, we have to know what they are. (Blackburn 1999, p. 1)

The idea is not new to Blackburn, though. A more traditional modern starting point which leads in a similar direction would be Carnap's method of *explication*. According to Carnap explication is

[...] the transformation of an inexact, prescientific concept, the explicandum, into a new exact concept, the explicatum. (Carnap 1950, p. 3)

The purpose of explication is to gain new understanding and insight into this explicandum, but importantly it aims to revise an informal or natural language concept, to end up with a more scientifically acceptable concept which is useful in academic research, for instance.²⁴ In other words, the

²⁴This does lead to some amusing mismatches in practice. See <http://www.mrlovenstein.com/comic/643>

purpose is to serve as a more fruitful and precise concept for the more exacting demands of scientific realms. Carnap held that explication should be guided by the following principles: similarity of the explicatum to the explicandum, exactness, fruitfulness and simplicity. These allow us to assess how good an explication is, but ultimately there may be several different explications which do well by the different criteria. For Carnap, then, there is no single correct explication, rather it is guided by our specific purposes and motivations for going about the process in the first place.

While Carnapian explication may be a kind of conceptual engineering, it does not cover the whole spectrum of possible approaches. Indeed, in their (Burgess & Plunkett 2013a,b), Burgess & Plunkett emphasise that conceptual engineering has a large normative component.²⁵ They identify four areas where conceptual engineering is already taking place, calling on philosophers to investigate this methodology more widely. These four areas are about personal identity through time; inconsistency, truth and logic; fundamental metaphysics; and race and gender.

The two examples of conceptual engineering I draw on fit into the second and last of these. The idea will be to see how the treatment of concepts in these separate domains might be transferred to mathematical concepts and, in particular, we are looking for a way to see how the open-textured, informal concepts usually deployed in proofs can work with and alongside formal definitions and derivations to spur the development of concepts and the growth of mathematics more generally.

3.7 Haslanger's Manifest and Operative Concepts

Haslanger's ameliorative project is ambitious, wide-ranging and has important political and social ramifications. I shan't go into depth on the project, but the central idea is the following. Beginning from the philosophy of race and gender, while we might want to investigate and describe how we deploy racial or gender terms, or explicate these to get to more robust categories, Haslanger argues that there is another project: to investigate the pragmatics of how this language is used and whether it serves our purposes. Haslanger says:

²⁵Their term is actually 'conceptual ethics', precisely because of the normative aspects of such an approach to concepts.

What is the point of having these concepts? What cognitive or practical task do they (or should they) enable us to accomplish? Are they effective tools to accomplish our (legitimate) purposes; if not, what concepts would serve these purposes better? (Haslanger 2000, p. 33)

In response to these questions, within feminist and anti-racist theorising one purpose we have is to move towards and achieve social justice. But the language of race and gender terms are part of sustaining the oppression of the groups it picks out. The ameliorative approach then, looks to come up with concepts of ‘gender’, ‘race’, ‘man’, ‘woman’, ‘black’, ‘white’ etc. which best suit our purposes in exposing and fighting injustice, oppression and inequality. To this end, Haslanger gives definitions of these terms which build in the subordination and power dynamics which they are associated with. For instance, to be a woman in part involves being subordinated along some dimension in virtue of the social role ascribed to women (see Haslanger 2000, pp. 42–43).

The ameliorative project itself falls under the general heading of conceptual engineering, in that constructing new, ideal concepts is part of the process of bringing about equality. For now, though, I want to pick out just one distinction which Haslanger makes frequent use of, such as in (Haslanger & Saul 2006) and (Haslanger 2012), between manifest and operative concepts. The insight here is that the concepts that we deploy in our practices do not always align with the concepts we take ourselves to be deploying. The *manifest* concept is that which we take ourselves to be using or working with, while the *operative* concept is that which we are actually putting into practice.

Many examples are used to give substance to it actually being rather common in practice for manifest and operative concepts to part ways. In (Haslanger 2005) the example is of what counts as being late to school (or ‘tardy’), with the difference between the official school policy and its practical implementation providing two different concepts. In (Haslanger & Saul 2006, pp. 99–100), Haslanger & Saul use the example of what counts as a parent, with the concept splitting between a manifest concept being the biological parents, and the operative concept of primary care-giver. In (Haslanger 2012, p. 92), the example is that of the concept of *being cool*. While cool dudes and the ‘in-group’ take themselves to be applying some

objective standards of what is and isn't cool (allowing a distinction between being cool and merely acting cool), in the end coolness is just the product of the interests and standards of the in-group, according to Haslanger. These simple examples lead up to the big one:

As in the case of “cool,” we debunk the idea of Woman's Nature and find two concepts at work: The manifest concept of Woman's Nature—understood as defining what women are by nature in traditional terms—is an illusion; the operative concept being masked by it is constitutively constructed in terms of men's (socially conditioned) sexual responses. (Haslanger 2012, pp. 93–94)

Further to the distinction in place, Haslanger adds the third notion of a *target* concept, which is the concept which

[...] all things considered (my purposes, the facts, etc.), I should be employing. In the ideal case, I adjust my practice and my self-understanding to conform to the target concept. (Haslanger & Saul 2006, p. 16)

This adds the normative component and the conceptual engineering, where we are able to reflect on, construct, and choose which concepts suit us and our purposes best.

Moving back to the realm of mathematical concepts, there are two levels at which we can identify a similar distinction at work. The first is in the deployment of mathematical concepts as we have been discussing so far in this chapter, while the second is at the level of concepts *about* mathematics (as discussed in chapter 2) such as the concept of proof itself. Let us consider these both in turn.

We might think that the relationship between informal mathematical concepts and formalisations thereof follows the manifest versus operative distinction in the following way. Formal definitions and definitions in axiomatic systems can be seen as useful for a number of reasons already described, such as their relative exactness, investigating foundations for mathematics, unifying systems, computational mathematics, metamathematics etc. The work of figures such as Frege, Russell, Whitehead and Hilbert (plus many others besides) demonstrated that large-scale formalisation was possible. A

favourite point of Azzouni's is to emphasise the widespread success of Frege's project and that of Russell & Whitehead's *Principia Mathematica* in their translational endeavours, e.g. (Azzouni 2005b, p. 19) and (Azzouni 2009, p. 10). The point being that the failure of naïve comprehension does not take away from much of the successful formalisation that went on, as is taken advantage of by the neo-Fregeans for example. As such, it is not surprising to find that there is often the view that these are ultimately what mathematical concepts pick out. On the other hand, the way we learn mathematics is very much about being inducted into particular mathematical practices, which lead more naturally to the informal concepts.²⁶ As such, the operative concepts which we learn through mathematics education do not necessarily coincide with the manifest concepts which mathematicians might take themselves to be using.

To make this more concrete, let us look at an example: set-theoretic foundations. Declarations to the effect that mathematics is at its core just set theory are ubiquitous in philosophy of mathematics, especially as espoused by mathematicians. Here are some samples (though opening up a random selection of set theory textbooks will be enough to furnish you with further examples):

[T]he mathematicians identify the natural numbers with the finite von Neumann ordinals. So, contrary to received wisdom, I suggest that philosophers follow mathematical practice and identify the natural numbers with the finite von Neumann ordinals. [...] Numbers are sets. (Steinhart 2002, p. 356)

and

All branches of mathematics are developed, consciously or unconsciously, in set theory or some part of it. (Levy 1979, p. 3)

and

Set theory is the foundation of mathematics. All mathematical concepts are defined in terms of the primitive notions of set and membership. In axiomatic set theory we formulate a few simple

²⁶Being inducted into a practice is an important way of obtaining *knowledge how*. This will be the focus of the next chapter.

axioms about these primitive notions [...] From such axioms, all known mathematics may be derived. (Kunen 1980, p. xi)

and

All this is in stark contrast to what we now regard as the answer to the question what mathematical entities exist? The working practitioner of classical mathematics can answer the question with one word — *sets*. Every mathematical entity is a set and all sets are objects, some of them being infinite objects. (Clark 2009, p. 347)

and

[M]athematical objects (such as numbers and differentiable functions) can be defined to be certain sets. And the theorems of mathematics (such as the fundamental theorem of calculus) then can be viewed as statements about sets. (Enderton 1977, p. 10)

Or slightly more cautiously:

[A]xiomatic set theory is often viewed as a *foundation of mathematics*: it is alleged that all mathematical objects are sets, and their properties can be derived from the relatively few and elegant axioms about sets. Nothing so simple-minded can be quite true, but there is little doubt that in standard, current mathematical practice, “making a notion precise” is essentially synonymous with “defining it in set theory”. Set theory is the official language of mathematics, just as mathematics is the official language of science. (Moschovakis 2006, p. vii)²⁷

You get the picture. Frequently these claims are in the opening pages of introductory texts for students, setting out from the start the manifest concepts for all of mathematics.

However, in practice mathematicians mostly are not doing set theory. In the Levy quote above he even says as much with the admission that the set theory supposedly underlying all mathematics might be unconscious.²⁸

²⁷Mathematics as the language of science is an allusion to a well-known quote by Galileo.

²⁸This is also reminiscent of Azzouni’s claims that mathematicians don’t need to be aware of underlying derivations, as discussed in chapter 1.

The operative concepts will vary with the situation and who is deploying them, of course, but I contend that frequently these will come apart from the manifest, set-theoretic concepts. There is a traditional exemplification of the manifest and operative concepts I have just described coming apart, that found in (Benacerraf 1965). With the excellent story of two children educated by “militant logicians” to truly believe that all mathematics boils down to set theory, Benacerraf shows how quickly problems can arise. For the children have been taught different set-theoretic representations of the natural numbers: one knows the von Neumann ordinals, the other the Zermelo ordinals. There are then examples of theorems for one of the children which are false for the other, such as $3 \in 17$ or whether any n -membered set can be put into one-one correspondence with the set n itself. Importantly, for my purposes, these issues can’t be resolved by asking other people:

Attempts to settle this by asking ordinary folk (who had been dealing with numbers *as* numbers for a long time) understandably brought only blank stares. (Benacerraf 1965, p. 54)

Why is this so understandable? Because it is common knowledge that, while the orthodox foundationalism declares that all mathematics reduces to set theory, this is simply not what is found in practice. Not just that, but the manifest concepts of the foundational picture come apart from the operative concepts in direct mathematical ways, as claims about the membership relation holding between numbers don’t even make sense unless you are working directly with sets. Now, it could be argued that the two representations are not on equal footing. Steinhart, for instance, in (Steinhart 2002) argues for the unique correctness of the von Neumann representation by what is essentially monster-barring, inventing additional demands and conditions on the concept of number which would give Lakatos a field day. Even still, this does not undermine the fact that those not separately trained in set theory would be baffled about statements of the form $3 \in 17$.

To extend the example somewhat further, observe that representing numbers in set theory is a straightforward and well-known construction. But many of the claims about set-theoretic reductions talk of *almost all of mathematics*, which includes a huge deal more. If we consider as an example low-dimensional topology, explored excellently in (de Toffoli & Giardino 2014), it seems likely that the mathematicians working in this area do not

know any actual set-theoretic construction corresponding to their concepts, so even if they were to agree to the claim that mathematics is founded in set theory, this would not be reflected in their operative concepts. However, even if we were to find those who did happen to know how to translate the main concepts into set theory, de Toffoli & Giardino demonstrate that the kind of operations licenced by the topological concepts are often visual, diagrammatic and manipulative (in a tactile sense)— far removed from those which one would find in the austere setting of set theory.²⁹ So the operative concepts, i.e. the concepts as employed in practice, are not going to be set-theoretical.

The foregoing does not stand against formalisation and representation of mathematics more generally in set theory. Indeed, as in the first chapter, the thought does not rely on arguing that formalisation is not possible for some class of cases. Rather, the point is that if one insists on identifying the mathematical concepts with some particular formal representations, then the resulting manifest concepts will still come apart from those which are operative in mathematics. As such, while one can point out that the readings of set theory as foundational I have quoted above are pretty strong ones, and even called ‘simple-minded’ in the quote by Moschovakis, we only need a fairly light version of the position for the distinction between the manifest and operative concepts to come out.

Let us move on to another kind of example. This echoes a lot that has already been covered and doesn’t require full repeating again, but the distinction between manifest and operative concepts may be a useful way to think about the formal and informal proof distinction. As has been discussed already, it seems to be fairly common for mathematicians, when pressed, to offer formal derivations as the manifest concept in play when it comes to the epistemological justification of mathematical propositions. This in turn leads to the need for the Formalist-Reductionist to explain how this can be, given how rarely these actually show up in practice. The idea raised by the proponents of the philosophy of mathematical practice is to observe that the operative concept, that of informal proofs, is doing all of the epistemic work and that this is the idea we should be taking more seriously. I certainly agree, and believe this terminology helps bring the idea out.

²⁹This is only austere, of course, when seen in the foundationist way. Actual set theorists will in practice draw plenty of diagrams to illustrate their ideas.

To finish, let me return to the starting questions which lead us into discussing Haslanger’s distinction in the first place. We wanted to find a way of seeing how formal and informal mathematics and their respective conceptual contents could operate side by side, both contributing to the understanding we have of the phenomena they describe and the discovery of new concepts, ideas, techniques and theorems. Applying Haslanger’s distinction, we have made progress because it is now clear that even in the realm of mathematical concepts, the (operative) concepts we put into practice do not always coincide with the (manifest) concepts we take ourselves to be using. Therefore, we do have the two notions working side by side and, furthermore, I have been displaying that in the mathematical case the difference between the manifest and operative concepts does span the formal and informal divide.

Nonetheless, it seems like the work is not yet complete. The discussion has showed that manifest and operative concepts coming apart in mathematics should be somewhat worrying to us. In the case of formal definitions not coinciding with the operative concepts they are meant to be identifying, there is an underlying philosophical difficulty that comes out in the fact that this seems like mathematicians are confused about what they are doing and the philosophical status thereof, with the danger that this might have mathematical repercussions (as in the Benacerraf and low-dimensional topology cases). Concerning the concept of proof, again we find that most mathematics does not fit to the standards that mathematicians frequently proclaim are to be demanded of correct proofs, with the potential difficulties that entails. So although we have achieved the ‘side-by-side’ aspect of our goal, the duality here so far seems problematic rather than virtuous.

What is needed is a story about how the operative concepts and the manifest concepts can work alongside one another in mathematics, rather than being in tension. I will in the next section investigate how we might make progress by looking at how Scharp manages this in the case of *truth*.

3.8 Scharp on Replacing Defective Concepts

Scharp holds the view that the concept of truth, as well as many other concepts, is inconsistent. The project of his book *Replacing Truth* (Scharp 2013) is to engage in conceptual engineering to design a replacement concept (or, as it turns out, two replacement concepts) which can succeed at the various

roles truth plays in our conceptual toolbox while also being consistent.

What does it mean for a concept to be inconsistent? Scharp discusses this at some length to rule out an opposing view that concepts cannot be inconsistent. On Scharp's view, a concept is inconsistent if and only if its constitutive principles are inconsistent. The toy example he uses is that of the 'rable', with the following constitutive principles:

(1a) 'rable' applies to x if x is a table.

(1b) 'rable' disapplies to x if x is red.

Importantly, the concept of rable is not logically inconsistent, but rather leads to a contradiction give the empirical fact of the existence of red tables, to which the concept both applies and disapplies. Given the environment we wish to deploy it in then, the concept of a rable is defective. Being a toy example, the defectiveness is not so important. However, the important idea is that the concept of truth is also defective because it also leads to inconsistency. The central constitutive principles for truth are the two directions of the T-schema:

(T-In) If φ then $\langle\varphi\rangle$ is true.

(T-Out) If $\langle\varphi\rangle$ is true then φ .

Then by the well-know reasoning of the liar paradox or one of its ilk, we can arrive at an inconsistency. There are good reasons not to give up either principle, as Scharp argues that these encode two important uses of truth: as a device for endorsement and rejection. To say that φ is true can function as a way of endorsing φ , thus T-Out encodes the endorsement function. Similarly, to say that φ is not true is a way to reject φ , so T-In encodes the rejection function (by its contrapositive).

The problem is that in a classical setting, no single concept will satisfy T-In and T-Out without leading to inconsistency. The conceptual engineering Scharp engages in solves the difficulty by replacing truth with *two* concepts which take over the two functions of endorsement and rejection separately:

[I]f we replace truth with two concepts, we can split the workload, allowing one to serve as a device of endorsement and the other to serve as a device of rejection. (Scharp 2013, p. 147)

Scharp replaces the concept of truth by the two concepts he calls *ascending truth* and *descending truth*, which each only satisfy one of the principles. Descending truth obeys T-Out, so functions as a device for endorsement; ascending truth obeys T-In, so functions as a device for rejection. Scharp furthermore defines the predicates expressing the concepts to be duals, i.e. $D(\langle\varphi\rangle) \leftrightarrow \neg A(\langle\neg\varphi\rangle)$ and $\neg D(\langle\neg\varphi\rangle) \leftrightarrow A(\langle\varphi\rangle)$.

Now we can begin to see how Scharp’s project might be a useful parallel to dealing with the concepts of mathematics. For Scharp proposes a formal, axiomatic theory to go along with the new concepts, a theory which exhibits a large number of principles which we would want truth to obey, but splits these between ascending truth, descending truth and several hybrid principles given in terms of both new predicates.³⁰ The crucial result is that this formal theory is consistent (or, to be clear, consistent if set theory is), which Scharp establishes by building a model with what he calls Xeno semantics.

We don’t need to go into a full assessment of Scharp’s theory of ascending and descending truth.³¹ Rather, what is of interest is the way in which the conceptual engineering Scharp engages in uses formal, technical machinery to develop the new twin concepts of ascending and descending truth, as well as the relationship this has to the defective concept of truth they are replacing. There is a particular objection based on practicality which Scharp can set aside, which is that it does not seem realistically achievable to convert the general public to the replacement concepts which are based on subtle and complex philosophical motivations.³² The reason Scharp can set this aside is that he does not believe that such widespread adoption is necessary:

It is essential to remember that, on the proposal defended here, it is legitimate to continue using ‘true’ for most purposes. Only where the difference between ascending truth and descending truth is not negligible does one need to use ‘descending true’ or ‘ascending true’ instead of ‘true’. (Scharp 2013, p. 174)

³⁰To be exact, the theory of ADT is meant to be a kind of minimal theory for ascending and descending truth, in that it will be a subtheory of any adequate theory for these two concepts. For instance, Scharp does not demand that a theory of ascending and descending truth should be fully axiomatisable, so this might only be a part of such a non-axiomatised account. (See Scharp 2013, p. 153).

³¹For a start towards this, see reviews of the book such as (Ripley 2014; Read 2014).

³²This point is frequently raised against Haslanger’s project too, but is even more pressing because of the centrality of trying to achieve political and social aims with the new concepts.

That is, a lot of the time the new framework is not necessary for the purposes at hand. Comparing the concept of truth to the concept of mass, Scharp elaborates:

Just as in casual conversation, people use ‘mass’ with the understanding that what they are saying might not be, strictly speaking, correct, but it is good enough for the purposes at hand. That is, those involved would have reached the same conclusions even if they had used the more complicated replacement concepts instead (with more effort). If a conversational participant wants to insist that the questions under consideration warrant a more precise conceptual framework, then those in the conversation can switch to the more precise terminology of relativistic mass and proper mass. Likewise, if necessary, conversational participants can switch from talk of truth to talk of ascending truth and descending truth. (Scharp 2013, p. 275)

So the replacement concepts, those which are consistent, are there in case the common but inconsistent concept of truth is likely to get us into trouble. Nonetheless, the everyday uses of truth are often harmless and actively useful—the well-known existence of paradoxes of truth has not hindered its functioning in endorsement and rejection, for instance. The defective and the engineered concepts both have roles in our theorising about and interaction with the world, and this perspective allows for them to operate side-by-side, each prevailing in suitable settings.

What about the case of mathematics and mathematical concepts? I have been investigating how the literature surrounding conceptual engineering might be applied to this question. While little previous work seems to have been done explicitly on this approach, there is an exception to be found in (Scharp & Shapiro forthcoming). In it, Scharp & Shapiro are still primarily concerned with truth, but run through three other examples of inconsistent concepts to illustrate their approach, two of which are mathematical: the naïve concept of set and naïve infinitesimals.³³ These two examples are

³³The final example is a practical situation in which the barber’s paradox appears, concerning a club called the “Secretary Liberation Club” precisely for secretaries of clubs they are not eligible to join, with trouble arising when they hire a secretary who wants to know if they can join the Secretary Liberation Club. The example is attributed to (Chihara 1979).

common in philosophy of mathematics, at least partly because analysis and then set theory were major stopping points in the development of the modern era of mathematics. Indeed, the process of rigourisation in analysis in turn led to the drive towards the now-dominant foundationalist picture. So it is an important observation that major concepts in the history of mathematics have turned out to be inconsistent, but Scharp & Shapiro equally stress that this does not mean that the practices they found themselves in were wholly broken:

The lesson of this episode (at least as we have characterized it) is that inconsistent terms need not undermine an otherwise successful and productive intellectual project. The project can go on so long as the practitioners have a good feel for what they can and cannot do with the potentially troublesome terms. And this particular project went on splendidly for some 200 or 300 years, engaging some of the finest mathematical minds ever. Until the trouble arose, internally, the project was not regarded as broken, and was in no need of fixing. (Scharp & Shapiro forthcoming)

As it turned out, infinitesimals were never really replaced, but rather banished. This did lead to the conceptual development of new definitions for all of the central notions of analysis via the $\epsilon - \delta$ definitions, which are replacements of previous concepts. For instance, the concept of convergence was replaced by the two notions of uniform and pointwise convergence, and observing that these two come apart in certain cases was an important mathematical step in the development of analysis. Set theory, on the other hand, has been a huge success *mathematically*, replacing the naïve concept of set with whatever is defined within the system of ZFC. Whether this is one concept of set or many is still an ongoing question underlying the debate between proponents of the universe view of sets (Woodin 2011) and the multiverse approach favoured by Hamkins (Hamkins 2012). I certainly think that debate could do with some input along the lines of this chapter on the nature of mathematical concepts.

I shall not attempt to cover the full range of mathematical concepts just in virtue of the sheer number and diversity of them in the history of mathematics. Still, it does seem that the strategy being championed by Scharp

is a useful one for beginning to untangle the many threads interweaving in the realm of mathematical concepts. However much I am suspicious of foundationalist projects, and their link to the Formalist-Reductionist philosophy, it cannot be denied that formalisation projects are enlightening to us in coming to understand mathematical concepts, theories and methods, as well as deeply important to our confidence in the edifice of mathematics. But the formal and axiomatised aspects of mathematics and their positive contribution to mathematical discovery also does not mean that we have to accept that these are the ultimate foundations for mathematics and that all that came before was broken. Rather, we can deploy Scharp's replacement strategy in recognising that certain concepts in mathematics are defective—and, occasionally, inconsistent—so we replace them and perform conceptual engineering to develop new ones.

Yet we should not be so simplistic as to see the old concepts and the practices they are embedded in as always flawed or misguided, nor that the concepts we end up with are deep or rich enough to embody all the wealth of mathematical thinking. Practices which make use of open-textured, informal concepts do contribute to the development of new mathematics and allow for a great deal of mathematical reasoning. The Scharpian move, involving detailed conceptual engineering of formal models and axiomatic theories of truth, should be mirrored in dealing with mathematical concepts, where the formal models and axiomatic systems do provide us with new insights into the informal and open-textured concepts, but where we only need to replace them in certain settings and contexts, such as those where we are investigating meta-mathematics, or how mathematical theories can be represented within one another. Instead of being in tension, the two ends along the axis of formality work in parallel and are useful to mathematics in different ways, or useful to different mathematical projects.

3.9 The Concepts of Mathematics

To finish, let us return to the three general questions about mathematical concepts change that we began with. These were:

1. Do mathematical concepts need to be replaced or merely revised?
2. Do *all* mathematical concepts need to be changed or merely some of

them?

3. Does the change of concepts need to apply to all contexts or just some?

The answers to these questions are going to be tied to some set of purposes we have for changing mathematical concepts. For instance, in Haslanger and Scharp their answers were linked to achieving social justice and providing consistent replacements for truth, respectively. So we can ask what the purposes for mathematical concepts and their development are? Well, from what we have seen there are multiple answers to this question. As we have seen, the big traditional answer is that we replace defective concepts in mathematics to ensure rigour. Indeed, the conceptual development underlying modern analysis was brought about as the “rigorisation of analysis” by Cauchy, Weierstass, Dedekind etc. Similarly, the follow-up project of the formalisation of mathematics by Frege, Russell & Whitehead, and Hilbert was guided in a large part by the desire to secure mathematics on a fully rigorous foundation. Besides rigour, though, we have seen from the dialectical approach to the philosophy of mathematics that conceptual development is also part of the broader methodology of mathematics. Hereby, the purpose of conceptual change will be connected to broader mathematical aims, such as problem-solving, structural understanding and discovering proofs. A full discussion of what mathematical conceptual development is aimed at when we are treating it with conceptual engineering would be part of larger project, one which I certainly don’t have the answers to, so let me set it aside for now.

Instead, let us consider where some of those writers discussed above land on the three main questions.

In section 3.5 we already compared Lakatos, Waismann and Kneebone on the identity of concepts. Identity conditions are clearly hugely important to the first question, as the weaker the identity conditions, the easier it is to call concepts “the same” and the more straightforward it is to claim that concepts are just being revised rather than being replaced. Lakatos, first of all, was above quoted as holding that proof-generated concepts *erase* and *replace* the naïve concepts. This offers a strong answer on the first question, but I suspect that it is slightly too strong to accurately capture Lakatos’s views. More rightly, the context of the line seems to refer to larger jumps in the understanding of our concepts, where these involve “characteristic”

proof-generated concepts, i.e. those central to a particular proof idea. This does not seem to rule out the thought that for smaller alterations to concepts it would be acceptable to call these “revisions” rather than “replacements”, as is somewhat suggested by the term ‘concept-stretching’, which brings to mind a single concept being deformed. On the other two questions, Lakatos believes that interesting mathematical concepts will all change through the proofs they occur in, barring perhaps formal proofs and the concepts deployed therein. On the other hand, Kneebone appears to place much more emphasis on the slow development of concepts, calling it ‘evolution’, ‘enlarging’ and ‘transformation’ of concepts, which is suggestive of revision over replacement. Like Lakatos, he seems to think that all mathematical concepts are subject to development and change, but unlike Lakatos he seems to think that context manages to fix concepts well enough in the moment to avoid equivocation and to put formal logic to work, pointing to a slightly more conservative answer to the third question. Meanwhile, we saw that Waismann thinks that concepts are too loose in their definitions for there to be a meaningful discussion about whether they are or are not the same before and after they have been confronted with new cases.

Turning to the figures from conceptual engineering, Carnap is somewhat ambiguous on the revise/replace question. For instance:

The task of explication consists in transforming a given more or less inexact concept into an exact one or, rather, in replacing the first by the second. (Carnap 1950, p. 3)

This seems to give both readings. On the other hand, Carnap is clearer that explication is important mainly for moving from ordinary language terms to scientifically precise ones, so on the last two questions the concepts are restricted to the ones we want to treat this way, and the context is specifically the scientific or academic one. Scharp is quite explicit about the use of replacing concepts rather than just revising them. This matches Scharp’s fine-grained view of the identity of concepts via constitutive principles. On the other hand, Haslanger often sounds like she wants to replace the unjust operative concepts with better target concepts, but also describes the ameliorative project as being akin to Carnapian explication in its method of improving concepts in (Haslanger & Saul 2006, fn. 5), ultimately seeing the question of what is going on in such conceptual analyses as inquiry-

dependent:

It should be understood, however, that on my view, whether or not an analysis is an improvement on existing meanings will depend on the purposes of the inquiry. (Haslanger & Saul 2006, fn. 5)

As such, there seems to be an openness similar to Carnap's position.

So what should we say to these questions for the specifically mathematical cases of conceptual engineering? Let us draw some morals from the preceding discussion. Having argued that mathematical concepts are open-textured like empirical concepts, it results that we should be open to the possibility of concepts being revised in the light of new cases. However, beyond this there are clear cases of replacements in the well-known examples of analysis and set theory. For instance, the previous concept of convergence was replaced by two concepts of uniform and pointwise convergence. Therefore we should conclude that both revision and replacement are to be found in mathematics and should be utilised. Taking a lesson from the dialectical philosophy of mathematics, the answer to the second question will be that interesting mathematics will develop all concepts we use, although pragmatically this might be restricted to less than that. For instance, it might be that different concepts are more or less fixed in our practices at different times. We might also limit replacement processes to concepts which are defective for our purposes, whatever those happen to be, which again might be a far more restricted class of mathematical concepts, one which varies over time with our purposes and our opinions on what defectiveness amounts to. The answer to the third question is that the contexts in which revision and replacement of mathematical concepts occur are not all of them but just some. As discussed, I believe this is important for the resolution of puzzles about the relationship between formal and informal mathematics. In different contexts our purposes will be best served by different levels of formality and rigour, and it is part of mathematical thinking to be able to move between these as the broader context demands. As such, we can look to conceptual engineering for strategies on how to connect the different concepts, to allow them to work in parallel.

3.10 Conclusion

In this chapter I have been looking at mathematical concepts and how they should often be seen in a dynamic rather than static way, in the sense that they are developed over time via mathematical practices. In the first half of the chapter I looked at how we can make it clear that not all mathematical concepts are fixed and ‘Fregean’, by examining Waismann’s notion of open texture and how it applies to mathematical concepts; Lakatos’s dialectical theory of mathematical conceptual change and the connection this has to proving; and Kneebone’s earlier work on changing mathematical concepts. However, there was the central problem that if concepts are tracking mathematical practices, then we need an account of how this extends to the modern side-by-side usage of informal proofs and formal results. I turned to conceptual engineering as a place to look for a strategy for dealing with the axis of informality and formality in modern mathematics, finding two approaches that might provide answers from Haslanger and Scharp. Finally, I examined what morals should be drawn from the discussions of this chapter towards how to give a full account of modern mathematical conceptual development.

The place we have reached leaves a great deal open and further work to be done. Nonetheless, I hope to have set out a path which leads from the work of Waismann, Lakatos and Kneebone, to a consideration of the relationship between proof, formality and concepts in modern mathematics.

Chapter 4

Proof and Mathematical Know-How

A proof is like the mathematician's travelogue. Fermat gazed out of his mathematical window and spotted this mathematical peak in the distance, the statement that his equations do not have whole number solutions. The challenge for subsequent generations of mathematicians was to find a pathway leading from the familiar territory that mathematicians had already navigated to this foreign new land. Like the story of Frodo's adventures in Tolkien's Lord of the Rings, a proof is a description of the journey from the Shire to Mordor.

— Marcus du Sautoy

4.1 Introduction

In *Why do we Prove Theorems?* (Rav 1999) Rav argues that mathematical epistemology should focus on proofs rather than truths or, as he puts it:

[P]roofs rather than the statement-form of theorems are the bearers of mathematical knowledge. (Rav 1999, p. 20)

The argument proceeds from here to various conclusions about mathematics such as rejecting both foundationalism and the Formalist-Reductionist approach to the justification of proofs. I will not focus on these so much here, rather I will investigate more deeply the epistemological picture that

we should adopt in response to Rav's arguments. The Ravian idea is to move away from the shallow picture of mathematical knowledge as knowing the propositional statement-forms of mathematical theorems, and towards one more focused on knowledge of techniques, methods and general proof-constructing tools. Crucially, what Rav taps into in the re-conception he provides of mathematical knowledge is a distinction familiar from epistemology between *knowing-how* and *knowing-that*. However, Rav does not explicitly relate his move to an emphasis on knowing-how in mathematics over knowing-that, nor is this drawn out in subsequent literature.¹ On the other hand, the literature in epistemology has a great deal to say on the relationship between knowing-how and knowing-that, so drawing out how these insights apply in the case of mathematics will provide a number of links, ideas and proposals that will in turn lead to a more subtle and worked-out theory of the practical side of mathematical knowledge.

The main idea of this chapter is the following. While accepting the point Rav makes that we should focus on mathematical methods, skills, techniques, concepts and connections, I will argue that the actual conclusion to be drawn from Rav's main argument is that mathematical epistemology should be investigating both the traditional propositional knowledge and knowledge-how, as well as the way in which the two connect. The last part is crucial: throughout I will argue that propositional and practical aspects of knowledge in mathematics are thoroughly entwined, and the job of mathematical epistemology should in part be to look at the relationship between them. I will argue later that this has a large impact on the epistemological function of proofs, in that activities of proving are epistemologically primary, while proofs are to be seen as a guide to action which are used to communicate directions on how to carry out the activities.

In section 4.2, I will weigh up Rav's argument for the switch of emphasis from mathematical truths to mathematical proofs to find that while it is successful at rejecting the opposing view, the outcome of his thought-experiment is not exactly the one he envisages, with interesting philosophical results. Next, in section 4.3, I will set out the distinction between knowledge-how and knowledge-that as it is found in the epistemology literature. I go

¹With the exception of (Löwe & Müller 2010) who argue that knowing-how in mathematics should be analysed in terms of mathematical skills. I will return to this later in section 4.5.

on, in section 4.4, to apply the epistemological picture to the case of mathematics, discussing benefits of the cross-over between the literatures and several examples. In section 4.5, I describe and critique the arguments from two papers by Löwe & Müller which argue that mathematical knowledge is context-dependent and should be filled out in terms of mathematical skills. From here I will be lead into a discussion of the epistemological importance of the activities of proving in section 4.6. Finally, I will return to the relationship between knowledge-how and proofs in section 4.7.

4.2 Rav and *Pythiagora* the Oracular Computer

One of Rav's main arguments for the claim that proofs are more important to mathematics than truths comes from the thought-experiment about an oracular machine called *Pythiagora*. *Pythiagora* is a universal decision machine, where if you enter any mathematical statement then the machine will respond immediately with a declaration of the statement's truth or falsity. Rav imagines a world in which such a machine sits on every mathematician's desktop. The result is that all open problems can be resolved simply by entering them into *Pythiagora*, while refereeing papers becomes far more straightforward since *Pythiagora* can immediately check any claimed theorems. All the hours of trying to establish the truth or falsity of some proposed claim are done away with, replaced by submitting the statement to the computer.

The argument is that if the role of mathematical proof is merely to deliver unto us knowledge of mathematical truths, then such a machine would be a welcome triumph, replacing the cumbersome and hard process of proving with an instant resolution to all mathematical problems that take our interest. However, Rav argues that *Pythiagora* would not be welcome at all:

A universal decision method would have dealt a death blow to mathematics, for we would cease having ideas and candidates for conjectures. [...] But conceptual and methodological innovations are inextricably bound to the search for and the discovery of proofs, thereby establishing links between theories, systematising knowledge, and spurring further developments. Proofs, I

maintain, are the heart of mathematics, the royal road to creating *analytic tools* and catalysing growth. (Rav 1999, p. 6)

Expanding on this point, Rav claims that the above benefits are to do with the epistemological role of proofs compared to that of truths. The claim is that knowing particular mathematical facts is by itself not very mathematically interesting, a point Rav makes through a case study of Fermat's Last Theorem and the following related quote from Gauss:

[I]f I succeed in taking some of the principal steps in [algebraic number] theory, then *Fermat's Theorem will appear as only one of the least interesting corollaries*. (Quoted in Bell 1937, p. 261)

On the other hand, the knowledge found in proofs is of a more interesting type, for proofs contain the problem-solving techniques which mathematics is most concerned with. Thus Rav can conclude with the primacy of the knowledge contained in proofs over that of mere truths, facts and theorems:

[P]roofs rather than the statement-form of theorems are the bearers of mathematical knowledge. Theorems are in a sense just tags, labels for proofs, summaries of information, headlines of news, editorial devices. The whole arsenal of mathematical methodologies, concepts, strategies and techniques for solving problems, the establishment of interconnections between theories, the systematisation of results—the entire mathematical know-how is embedded in proofs. (Rav 1999, p. 20)

Now, Rav makes a compelling case that the kind of knowledge embedded in proofs is both interesting and importantly different from knowing the theorems proved, as well as for the general point that mathematics without proofs would be significantly depleted. Nonetheless, the response to the *Pythagora* example that Rav uses does not seem entirely plausible.

In stark contrast to Rav's claim that mathematicians' constant access to *Pythagora* would deal a death blow to mathematics, a better interpretation of the proposed scenario proceeding is the very opposite: not a dark age but a golden age of mathematics. Of course, on first getting access to the oracular powers of *Pythagora* the reaction of many, most or maybe all mathematicians might be to check the truth of both famous open problems

and the problems they personally are working on. It is only natural for people to be curious whether the Riemann hypothesis is true or false, whether $P = NP$ etc. But unlike the description Rav gives, this doesn't seem to be the end of the mathematicians' curiosity. Now certainly Rav is right to reject the simple picture on which we are only interested in truths, and this is intuitively clear from the repulsion of doing mathematics with these alone, and no proofs to reveal deeper structural connections, but the focus on proofs with truths relegated to mere tags and labels suggested by Rav's rhetoric strikes me as equally implausible.

The golden age of mathematics that would come about in the thought-experiment of *Pythiagora* would be due to the fact that having access to an oracular machine would not do away with proving in mathematics, but would on the contrary facilitate it. For one thing, knowing whether a claim or its opposite is true lets us concentrate on the proof of that. Take, for example, the question of $P=NP$. While this is an open question, the general suspicion amongst computer scientists and mathematicians appears to be that it is false.² Nonetheless, if it turns out that it is true after all then we may be directing the majority of our efforts at proving the wrong thing. Thus, knowing the truth of $P = NP$ is immediately useful for mathematics.

But knowing the final truth of the ultimate theorem is just the tip of the iceberg in terms of usefulness to proving. We would be able to use *Pythiagora* to not just verify theorems, but also check lemmas and intermediary steps on the way. Such an application is key to the point I am making: proofs do not just use skills, techniques and methods to suddenly arrive at the truth of the theorem, but rather they interweave the application of these skills, techniques and methods with a whole series of propositional facts. Stated like this it may sound rather obvious, but the point is that downplaying the importance of theorems and facts relative to proofs in the epistemology of mathematics misses the deeper interconnection between different types of mathematical knowledge.

Let us examine a little further the usefulness of having access to *Pythiagora* for broader mathematical endeavours. For one thing, the time wasted on blind alleys and poorly chosen lemmas in mathematics must be sub-

²I say suspicion rather than consensus because it isn't hugely unusual to think the opposite. A poll, now somewhat out of date, found out of 100 well-respected respondents, 61 expect that $P \neq NP$, to only 9 believing $P=NP$. See (Gasarch 2002) for details.

stantial, and being able to check these in an instant would be informative, time-saving and helpful. Of course, one might be concerned that there is some benefit to the honest toil of going in wrong directions, such as that seeing why they are wrong directions will be helpful in formulating improvements and better choices. However, having access to *Pythiagora* would also be very useful in providing insights into where incorrect lemmas and theorems go astray, since it would make it significantly easier to discover actual counter-examples, from which one may well be able to glean more than a failed proof alone. The use of an oracle can thus augment and direct our proving efforts rather than destroy them, because epistemic access to truths is a vital component of successful proving.

The world in which we have access to *Pythiagora* is a fantasy land, of course, for it is well-known that such a universal decision engine is mathematically impossible. It might therefore seem indulgent to imagine this scenario beyond the playful use Rav himself puts it to. However, we don't need anything so wild as a *universal* decision machine to see that Rav's conclusion is too extreme by far. All that is needed is something which is found increasingly extensively in the real world: computational mathematics. As an example, consider the system *GAP* (standing for 'Groups, Algorithms and Programming'), which is a programme used in computational group theory, which includes a programming language, an interactive environment, huge data libraries of examples from group theory and related areas, and implemented algorithms.³ In (Martin 2015), *GAP* is examined from the point of view of mathematical practice, revealing the following:

Research users of *GAP* typically use it to experiment with conjectures and theories. Whereas pencil and paper calculation restricts investigations to small and atypical groups, the ready availability in *GAP* of a plethora of examples, and the ease of computing with groups of large size, makes it possible to develop, explore and refine hypotheses, examples and possible counter examples, before proceeding to decide exactly what theorems to prove, and developing the proofs in a conventional journal paper. (Martin 2015, p. 42)

In essence, this is exactly what I am describing: ready access to confir-

³Its current homepage can be found at: <http://www.gap-system.org/>

mation and computational results, which serve a similar function as the oracular *Pythiagora* outputs in Rav's example, doesn't hinder mathematics but instead function as a useful and productive part of mathematical discovery. A similar example is that of *SnapPy*, which focuses on the geometry and topology of hyperbolic structures.⁴ Again, it provides computational results, such as being able to compute whether an inputted planar projection diagram⁵ has a canonical decomposition. On the other hand, a major focus of their project is also for the program to provide mathematicians with quick, computer-generated visualisations of the various objects it deals with, which in turn provide deeper understanding of the objects involved. In both cases, the computational approach to mathematics certainly provides the sort of propositional feedback that Rav's thought experiment portrays as a danger to mathematics. Yet, a more careful analysis of the effects on mathematics that computational work has demonstrates that this does not lead to abandoning proofs and proving in the mathematical methodology, rather it augments it and opens up new ways for mathematicians to do mathematics. Indeed, this was part of lesson to be learned from the previous chapter.

A potential response that could be offered by Rav is to point to examples of computer-assisted proving which are less easy to frame in such a positive light. For example, the proof of the Four Colour Theorem involves reducing the problem to some large number of cases and then using a computer to verify the truth of the theorem for those cases. Now there is a case to be made here for how the computational work fails to deliver the knowledge we would usually expect from a proof, or maybe more accurately, it doesn't give us a proper understanding of the truth of the theorem. Certainly a great deal of philosophical literature is concerned with the particular case of the Four Colour Theorem and the shift in proving practices it represents, e.g. (Tymoczko 1979), (Detlefsen 2008) and (Davis & Hersh 1981, pp. 380-387). And it may well be true that the essential use of a computer to produce a proof too long to be checked by a human mathematician represents a radical shift in mathematical practices and proof in particular. However, this is precisely the point as I see it: mathematical proof might be augmented and changed, but this is not to the detriment of mathematical methods, ideas

⁴Its current homepage can be found at <http://www.math.uic.edu/t3m/SnapPy/>. Thanks to Adam Epstein for bringing this system to my attention.

⁵The programme comes with software for anyone to draw their own planar diagrams, which are a way of representing objects in low-dimensional topology.

and progress. As an important observation, the Four Colour Theorem itself was not the end of mathematics on this topic, not the “death blow” that might be expected, but has been the source of another milestone of progress when it was *formalised* by Gonthier, as described in (Gonthier 2008). The controversy had been about the computational enumeration of the huge number of cases of reductions for the Four Colour Theorem, with no doubt cast on the traditional but previously informal parts of the proof leading up to that point. However, Gonthier’s report describes the significant and interesting mathematics that came out of formalising the proof fully. Two points result from this. Number one: formal proofs and formalisation do not hinder mathematical discovery (including of methods and techniques) but instead provide new ways of examining, generalising and coming to understand the mathematics involved. This is in line with what I argued in the previous chapter. Number two: the Formalist-Reductionist once again doesn’t benefit from this point. For if one believes that informal proofs are justified by their formalisations, then the fact that an accepted informal proof might require substantial mathematical discoveries to be formalised casts serious doubt on this position. The informal proof was held to be justified by all of the usual standards, but the formalisation was inaccessible without the substantial extra mathematical work.

As a further response to the *Pythiagora* argument by Rav, it should be noted that the “death blow for mathematics” rather exaggerates the impact of no longer needing proofs in another direction too, in that it unnecessarily restricts the realm of mathematics. In fact, mathematics extends far beyond the pure mathematics which Rav focuses on throughout his article, to include the vast swathes of applied mathematics, statistics, computational mathematics, etc. The point is that the majority of these other areas of mathematics are bound up with modelling in some form or other rather than establishing ultimate truths, and as such are better assessed holistically rather than individually. The point is very similar to Elgin’s point about scientific knowledge in (Elgin 2006). Indeed, Elgin’s argument that the factivity of knowledge does not fit well with practices such as scientific theory-building which are not aiming at truth per se transfers directly to the broader mathematical situation. In the same way that scientific theories involve purpose-dependent simplifications, abstractions and idealisations, meaning that the full theory does not purport to be literally true,

many mathematical theories outside of pure mathematics are focused on modelling phenomena.

For instance, in an undergraduate applied mathematics course one might encounter all kinds of models, from mechanics to population dynamics. Without the theory they are embedded in many of the propositions the theory is built on and generates will not strictly be true, so is hard to describe as ‘knowledge’ given that this is standardly taken to be factive. Elgin’s conclusion is that a more important epistemic goal in these cases is *understanding*, meaning the understanding of the phenomena afforded to us by the theorising or mathematical techniques employed. Understanding, which can come in degrees and different forms, clearly relates to both propositional facts and practical knowledge but is not obviously reducible to either of them. The thought-experiment of *Pythiagora* is all about delivering truths, but if large swathes of mathematics are more focused on understanding phenomena mathematically than on finding the ultimate truth of things, then the thought-experiment cannot be expected to see the same devastating effect on mathematics. For these other areas of mathematics, the same honest toil will still be required to get their results. Again, this point is in the same direction as Rav’s intended broadening of the class of interesting mathematics, but actually goes beyond it. As part of the philosophy of mathematical practice, it is worth remembering that the narrow view of mathematics which is treated in the standard philosophical accounts is not exactly faithful to the diversity one finds in mathematics departments.

Thus it would appear that contrary to the simple view Rav attacks, mathematical knowledge must be more than propositional knowledge of mathematical truths, but contrary to Rav’s proposed picture, the practical knowledge embedded in proofs is not by itself an adequate replacement. The two responses I have given to Rav’s *Pythiagora* thought experiment suggest that mathematical epistemology should be concerned with at least three key concepts: propositional mathematical knowledge, mathematical know-how and mathematical understanding.

4.3 Knowing-How and Knowing-That

Rav’s discussion goes into detail on why it is that we are more interested in proofs than in the truths of mathematics. The ultimate point, though, rests

on the kind of knowledge of mathematics we're after: that knowledge of isolated truths (such as those offered by *Pythiagora*) is not in itself of much interest, whereas knowledge of proofs contains the structural connections, the methods, the ideas and the relations of mathematics. Recall from above:

[...] the entire mathematical know-how is embedded in proofs.
(Rav 1999, p. 20)

We find this line in and amongst the many other advantages knowing proofs has over knowing truths. However, there is a large literature in epistemology working on precisely the relationship between knowledge-how and knowledge-that, one which has hardly been brought to bear on mathematical knowledge at all. If we take Rav's point seriously, as I believe we should, then it will be a worthwhile project to see what the status of knowing-how is in mathematics and how it relates to the standard propositional knowledge which has been the common subject of mathematical epistemology in the past. To enable such an investigation, I will in this section set out the state of play on this topic as it stands in epistemology. I will proceed from Ryle to Stanley & Williamson, ultimately settling on the picture by Wiggins that knowledge-how and knowledge-that are tightly connected and difficult to pull apart in practice.

To begin, the first important observation is that the intended distinction is between the standard objects of knowledge, which are propositional, and some distinctly practical kind of knowledge, but that this distinction is not necessarily tracked exactly by the linguistic constructions. For instance, consider the following two sentences:

- (1) Caroline knows how to play Doppelkopf.⁶
- (2) Ryo knows how Alper cheated at Doppelkopf.

While the first sentence typically ascribes a practical sort of knowledge to Caroline, the second might amount to no more than:

- (3) Ryo knows that Alper cheated using secret hand-signals.

where this is straightforwardly seen to be essentially propositional. The point is made by Rumfitt in (Rumfitt 2003) by emphasising the difference

⁶This is a German card game similar, in certain respects, to Whist.

between the semantics and metaphysics of knowledge-how. Similar points are also seen in (Moore 1997, ch. 8) and (Glick 2011, p. 403) and likely originate in (Hornsby 1980, p. 84). While the linguistic structure can be used to pick out both forms of knowledge, as in (1) and (2), the metaphysical interest lies in investigating just the practical kind of knowledge. That is to say, although the usual way of talking about the distinction is between knowing-how and knowing-that, the interest is not in just any knowledge identified by “knows how” constructions. In the literature, it is fairly common to talk about *knowledge-how-to* instead, but the emphasis on the kind of knowledge rather than the semantics of sentences is an important issue we shall return to shortly.

The main source for the philosophical distinction in the literature is Ryle in (Ryle 1946) and (Ryle 1949, ch. 2) who took himself to be refuting the ‘intellectualist’ position:

Mathematics and the established sciences are the model accomplishments of human intellects. [...] They thus bequeathed the idea that the capacity to attain knowledge of truths was the defining property of a mind. Other human powers could be classed as mental only if they could be shown to be somehow piloted by the intellectual grasp of true propositions. [...] the intellectualist doctrine [...] seeks to define intelligence in terms of the apprehension of truths, instead of the apprehension of truths in terms of intelligence. (Ryle 1949, p. 27)

Hereby we come to a debate between the intellectualists on one hand and Ryle’s anti-intellectualism on the other. Firstly, *intellectualism* is the view that holds that all knowledge-how is just propositional knowledge, or that knowledge-how is just a form of knowledge that. On the other hand, *anti-intellectualism* is the position that knowledge-how and knowledge-that are distinct, or that knowledge-how is practical in a way that knowledge-that is not or cannot be. As it stands, Ryle’s anti-intellectualist view is generally acknowledged to have become the philosophical orthodoxy. However, the modern debate was re-ignited by Stanley & Williamson in their paper (Stanley & Williamson 2001), who took on the intellectualist position against Ryle. Let us go through both in turn.

Ryle's form of anti-intellectualism sees making a distinction between the two types of knowledge as hugely important for rejecting a narrow approach to philosophy which is largely based on our knowledge and discovery of facts:

In [philosophers'] theories of knowledge they concentrate on the discovery of truths or facts, and they either ignore the discovery of ways and methods of doing things or else they try to reduce it to the discovery of facts. They assume that intelligence equates with the contemplation of propositions and is exhausted in this contemplation. (Ryle 1946, p. 4)

He thus sees the exclusion of the ways and methods as failing to explain intelligent action. A useful example for our purposes is his discussion of logical inference, for which he turns to the Lewis Carroll's story of "What the Tortoise said to Achilles" (Carroll 1895). The idea is to consider a scenario in which a student understands two premises A and B and a conclusion Z in a valid argument, but fails to appreciate that the conclusion follows from the premises. For the intellectualist, Ryle suggests, what is needed is to add the proposition C : "if A and B are true, then so is Z " to their consideration. But Carroll's Tortoise observes that the problem is not resolved: the student might still fail to make the inference. Again, we can add a further instructive premise D : "if A , B and C are true, then so is Z " and again the student might fail to infer Z from A , B , C and D , and so on ad infinitum. Ryle argues that what has gone wrong here is the assumption that intelligent action can be reduced to propositional theorising, for the problem is that the student can accept all of the new premises *in theory* without any of them forcing the student to accept the conclusion. The upshot is:

Knowing a rule of inference is not possessing a bit of extra information but being able to perform an intelligent operation. (Ryle 1946, p. 7)

Similarly, Ryle's main objection against intellectualism takes the form of a regress:

If a deed, to be intelligent, has to be guided by the consideration of a regulative proposition, the gap between that consideration and the practical application of the regulation has to be bridged by some go-between process which cannot by the pre-supposed

definition itself be an exercise of intelligence and cannot, by definition, be the resultant deed. (Ryle 1946, p. 2)

If the intellectualists are right, and knowledge-how is just a kind of knowledge-that, then Ryle's regress puts the following problem to them. Exercising knowledge-how in intelligent action will involve two distinct parts: the mental consideration of the relevant proposition and the resultant action. But then the mental consideration of the proposition needs to be an intelligent action too (so that the consideration is done in the right way, at the right time, in the correct circumstances etc.)⁷ If so, then we need a further act of mental consideration of a proposition to underlie that action, which in turn requires another etc. Ultimately then there is the challenge of finding something which "reconciles these irreconcilables" (Ryle 1946, p. 3), a challenge which Ryle believes the intellectualist will be hard-pushed to meet.

As for a positive view, Ryle's position links knowledge-how closely to the actions and activities that the intellectualist struggles to accommodate. For instance:

When a person knows how to do things of a certain sort (e.g., make good jokes, conduct battles or behave at funerals), his knowledge is actualised or exercised in what he does. (Ryle 1946, p. 8)

And while in all of these we can find rules, maxims, canons, principles etc. (especially in the case of logic), for Ryle these are separate from their judicious application:

In short the propositional acknowledgement of rules, reasons or principles is not the parent of the intelligent application of them; it is a step-child of that application. (Ryle 1946, p. 9)

which is to say that rules are abstracted from practice, rather than knowledge of them being necessary for knowledge-how. Such principles are helpful for many things, such as pedagogy, but the metaphysical nature of knowledge-how for Ryle is more than will be captured by any such rules or maxims. In particular, knowledge-how will surpass these rules precisely

⁷In their discussion of the regress, Stanley & Williamson leave out the 'intelligent' part of this, focusing simply on considering propositions as necessary for action. For a further discussion of the form of the regress argument see (Fantl 2012).

in that knowledge-how also involves regulating the activity as it goes along, being prepared for eventualities that might arise and adapting as necessary to them. Furthermore, Ryle's position also involves the possibility of a kind of internalisation of the activities in one's know-how:

But very soon he comes to observe the rules without thinking of them. He makes the permitted moves and avoids the forbidden ones; he notices and protests when his opponent breaks the rules. But he no longer cites to himself or to the room the formulae in which the bans and permissions are declared. (Ryle 1949, p. 41)

As a result of the ongoing adaptability and the possibility of internalisation, the picture of knowledge-how that Ryle proposes allows for knowledge-how to be potentially quite complex. This must be right, given the large range of knowledge that it picks out. For instance, consider:

(4) Joe knows how to spell 'rhododendron'.

(5) Carley knows how to fly a helicopter.

While the former picks out a very particular piece of know-how, the latter is complex and broad, requiring the co-ordination of many other skills and competences across a range of situations.

There is a caricature of Ryle as holding an *ability* account of knowledge-how, that is:

Ability: S knows how to V iff S is able to V .

This picture is actually a fairly popular one, found in (Rosefeldt 2004) and (Noë 2005) for example, and does well on a large range of cases. It would be very odd to accept (5), that Carley knows how to fly a helicopter, but then to also hold that she is unable to fly a helicopter (at least in some salient range of situations). However, there are a number of common objections. Mainly, the ability account seems to fail on both necessary and sufficient conditions.⁸ As a counterexample, take the example of a top chef who loses his sense of taste in a dire cheese accident. It is plausible that the chef loses the ability to cook certain complex dishes, without having lost any knowledge. Conversely, at times we are able to do things by luck alone, without knowledge of how to

⁸This is brought out particularly clearly in (Bengson & Moffett 2011).

do them, such as if I were to try to fly a helicopter (I definitely don't know how) and just happen to do everything right. Hawley has a particularly nice example in (Hawley 2011) of a climber getting caught in an avalanche, but mistaking the snow for water and therefore making swimming motions to get out. As it happens, this is the correct way to escape avalanches. While she is able to escape the avalanche by making swimming motions, she is furthermore reliably successful at doing so. Nonetheless, it would be highly implausible to say that the climber has knowledge of how to escape avalanches.

To be fair to Ryle, in (Hornsby 2012), Hornsby makes a convincing case for the idea that Ryle did not hold the abilities view and was aware of the kind of counterexamples that such a view would be subject to. She says:

In connection with knowing-how, he spoke of all of “abilities,” “skills,” “competences,” and “capacities,” and one might assume that he used these various terms in part because he recognized that ‘knowing how’ could not be understood in terms simply of ability. (Hornsby 2012, p. 82)

Hornsby's contention is that Ryle's main aim in discussing know-how is not to give an account of necessary and sufficient conditions, but rather to establish the existence of another category of knowledge which is not directly about propositions. The reason for this is that it is required for Ryle's broader project of rejecting the Cartesian dualist position, and in particular the idea that intelligence and practice come apart:

[T]here is no gap between intelligence and practice corresponding to the familiar gap between theory and practice. (Ryle 1946, p. 2)

So Hornsby's point seems right in reading Ryle in the following way:

The Cartesian thinks that the mental is separate from the physical. Ryle wanted it to be clear that the states of mind implicated in intelligent bodily action are inseparable from bodily action itself. (Hornsby 2012, p. 87)

Hornsby is talking about the way in which the Cartesian “myth” which Ryle stands in opposition to separates thinking from doing, or mental thought

from bodily action, with knowledge situated firmly in the mental realm. Ryle's purpose in exploring knowledge-how, therefore, is not to give a reductive account, but to show that there is another class of knowledge which is not to be relegated to the mysterious mental domain. For this reason, rejecting the gap between the 'inner' mental life and external behaviours was a more central aim than an explicit account of knowledge-how.

Let us move on for now to the modern intellectualist revival brought about in (Stanley & Williamson 2001). Stanley & Williamson take the Rylean regress argument against intellectualism to be invalid and argue for the thesis that knowledge-how is just a species of knowing-that. Their argument against the regress is that there is an equivocation on the kind of action involved in knowledge-how. On the one hand, the sort of actions we employ knowledge-how to perform are *intentional* actions, while the contemplation of a proposition involved in exercising that knowledge how is not necessarily intentional, so the move to form a regress of needing deeper and deeper knowledge-how is rejected. They cite this move as coming from Ginet, who says the following:

I exercise (or manifest) my knowledge *that* one can get the door open by turning the knob and pushing it (as well as my knowledge *that* there is a door there) by performing that operation quite automatically as I leave the room; and I may do this, of course, without formulating (in my mind or out loud) that proposition or any other relevant proposition. (Ginet 1975, p. 7)

As such, the regress argument doesn't succeed. Responses to this are expanded in (Hornsby 2012, pp. 95–95) and (Fantl 2012).

Stanley & Williamson's positive case for intellectualism mainly rests on one central argument based on the language and semantics of knowledge ascriptions. They present the 'standard' semantics for the meanings of sentences of the form "P knows how to V" and demonstrate at length that this analyses knowledge-how as a relation to a proposition, through an embedded-question construction which fits the same structure as knowledge-what, -where, -who, -when etc. constructions which are normally taken to be relations to propositional answers. The key question then is what exactly the propositions known are when someone possesses knowledge-how.

Distilling the argument somewhat, here are two of their examples and the central thesis expressed about them:

(19) Hannah knows how PRO to ride a bicycle. (Stanley & Williamson 2001, p. 424)

(20c) Hannah knows how she could ride a bicycle. (Stanley & Williamson 2001, p. 425)

Relative to a context in which (19) is interpreted as (20c), (19) is true if and only if, for some contextually relevant way w which is a way for Hannah to ride a bicycle, Hannah knows that w is a way for her to ride a bicycle. Thus, to say that someone knows how to F is always to ascribe to them knowledge-that. (Stanley & Williamson 2001, p. 426)

The ‘PRO’ that Stanley & Williamson use is part of the syntactic theory, which functions as “a phonologically null pronoun that occurs [...] in the subject position of untensed clauses.” (Stanley & Williamson 2001, p. 419) but this is not so important for our purposes. The point is that the relevant constructions which ascribe knowledge-how to some agent can be interpreted in terms of propositional knowledge, in particular of knowing that some way is the way that the activity can be done.

In this way Stanley & Williamson use the standard semantic theory to reduce knowledge-how to propositional knowledge of ways. An important additional feature of this account is that the propositional knowledge ascribed in knowledge-how must frequently be held under a *practical mode of presentation*. For there is a potential difficulty in the propositional answer to the embedded ‘how’-question, which is that knowledge that w is a way to V is possible to have in plenty of cases where we don’t seem to know how to V in the substantial way we were interested in. For example, if I point to an expert juggler and observe that the way they do it is a way to juggle nine balls, I might know that THAT (while pointing) is a way to juggle nine balls, without myself possessing any knowledge on how to actually go about doing it. The practical mode of presentation is then the connection to the complex dispositions involved in knowing-how to do something in a way that makes it plausible to link it to ability, skills etc. They can then claim:

It is for this reason that there are intricate connections between knowing-how and dispositional states. But acknowledging such connections in no way undermines the thesis that knowing-how is a species of knowing-that. [...] It is simply a feature of certain kinds of propositional knowledge that possession of it is related in complex ways to dispositional states. Recognizing this fact eliminates the need to postulate a distinctive kind of nonpropositional knowledge. (Stanley & Williamson 2001, pp. 429–430)

The investigation I am interested in of knowledge-how in the mathematical context is premised on there being an interesting kind of mathematical knowledge which has not been adequately appreciated in mathematical epistemology previously, so in turn presupposes the intellectualist position of Stanley & Williamson to be flawed. A full discussion of the merits and failings of the theory could lead us far astray, so let me just set out succinctly four ways in which I believe it is incorrect, though all are closely related. Firstly, the heavy use of semantic theory is significantly less convincing once we see that the account is not uniform across different languages, as aptly displayed in (Rumfitt 2003) and developed in (Ditter 2016). Rumfitt shows us that the best semantic theories for other languages range from largely unhelpful for Stanley & Williamson’s point (in the case of French) to fully antithetical to it (in the case of Russian). Secondly, there is the underlying move of making the jump from the linguistic and semantic facts to the metaphysical realm of what knowledge really is, which is not well-justified. Alva Noë (Noë 2005) and Jessica Brown (Brown 2013), for instance, both argue that the language-first methodology fails to consider empirical science, especially cognitive science, as a potentially divergent source of evidence on the nature of knowledge-how. Thirdly, linguistics and semantics aim to study language and meanings, but in building theories there are significant idealising assumptions made. For example, Noam Chomsky starts his (Chomsky 1965) with the following:

Linguistics is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in

applying his knowledge of the language in actual performance.
(Chomsky 1965, p. 1)

The best understanding of these practices, then, is under the heading of *modelling*, broadly construed. However, with this in place, if one wants to make the kind of argument Stanley & Williamson do, then it must be demonstrated that the conclusions one is drawing are not artefacts of the modelling process. No such argument is given, but furthermore there is a clear story to be told about why modelling knowledge-how ascriptions in terms of propositions does have the propositions merely as a result of the modelling: because these are the standard, well-known building blocks of semantic theories more generally. After all, why expand the theoretical apparatus when instead we can just simplify the complexity involved in practical knowledge. Indeed, this strikes me as the reason that Stanley & Williamson need to posit both ‘ways’ and ‘practical modes of presentation’ as a bridge back to the actual phenomenon we are interested in. This takes us to the fourth point: both ‘ways’ and ‘practical modes of presentation’ act as a philosophical black box, doing essentially all of the work in getting to knowledge-how but without any indication of what is going on inside. Jennifer Hornsby demonstrates that ‘ways’ in particular are not up to the task required of them on this picture (Hornsby 2012, pp. 90-92).

Having seen Stanley & Williamson’s intellectualism and the Rylean anti-intellectualism before that, I want to finish this section with a consideration of positions which are located between the two extremes. One way to find such a position is to follow (Glick 2011) in separating out *weak intellectualism* from *strong intellectualism*, where the former holds that knowledge-how has propositions as a relatum, while the latter equates knowledge-how with full theoretical knowledge, i.e. possessing standard features of knowledge-that like “(some subset of) belief, justification and Gettierizability, linguistic accessibility, availability of content for use in inference, and concept-possession.” (Glick 2011, p. 411). Importantly, this allows one to find positions which satisfy the weak but not the strong intellectualist position, with scope for making the relation between the knowing subject and the propositions be of some sort which would be acceptable to anti-intellectualist thinking.

The middle-way I will look further at, though, comes more from the Rylean direction in Wiggins’s (Wiggins 2012). While Wiggins explicitly

accepts that Ryle's general point is correct against the intellectualists, that there are distinct kinds of knowledge, the interest lies in the way in which he brings out the close connection between the distinct pieces of practical and propositional knowledge.⁹ Wiggins expands further on Ryle's metaphor of the 'step-child' we saw already above:

In short the propositional acknowledgement of rules, reasons or principles is not the parent of the intelligent application of them; it is a step-child of that application. (Ryle 1946, p. 9)

The idea being that a great deal of propositional knowledge, even of rules or principles, rests on prior practical knowledge. Wiggins gives us the following example:

A ship's pilot who is retained by the maritime authorities to bring large ships safely to anchor in an awkward or difficult harbour can tell us, on the basis of his competence and experience, that when the wind is from the north and the tide is running out, the best thing to do is to steer straight for such-and-such a church tower until one is well past a certain bend in the channel. Almost anyone can come to possess that propositional knowledge but the information they get in this way will probably rest indispensably upon the experience and practical knowledge of a handful of people with a different kind of knowledge, namely practical or [...] agential knowledge. (Wiggins 2012, p. 109)

There is a clear flow in one direction, then: that propositional knowledge often comes about through practical experience. But I take it that the thought extends in the other direction too. Once the rule of thumb is in place, the ship's pilot can quickly explain how to anchor in the difficult harbour to other ships' pilots, that they can thereby combine this propositional item of knowledge with other know-how of sailing they already possess to be more knowledgeable generally of how to make it into the harbour safely.

I think it is safe to go even further, in fact, in that a great deal of our learning comes about through a combination of practical and propositional knowledge. It would be strange to expect us to be able to trace back our

⁹Recalling my response to Rav's *Pythagora* thought-experiment above, the idea that practical and propositional knowledge are interdependent will be crucial in what is to come.

knowledge to find it ‘bottoming out’ at either pure knowledge-how or pure knowledge-that. It is much more likely that we would find the two to be fully intertwined and mutually supporting. The Wigginsian picture, then, is one which there is an interdependence between knowledge-how and knowledge-that. Rather than having one being the step-child of the other, we find that the different kinds of knowledge are a close-knit family group.

4.4 Mathematical Know-How

In this section I will return to the mathematical realm and explore how knowledge-how can arise here too. Agreeing with the anti-intellectualist claim that there is a substantial difference between knowledge-how and knowledge-that, and the broader picture from Wiggins on which the two kinds of knowledge are nonetheless strongly interrelated, we now have some ideas to bring back to apply to mathematics.¹⁰

To begin, let us connect Rav’s ideas to the epistemological literature. Rav claimed that the interesting mathematical knowledge is knowledge of proofs, since this is where the methods, techniques, concepts and interesting ideas, which are the essence of mathematics, actually reside. The thought then is that Rav’s big insight, put in the epistemological framework, is that the interesting knowledge in mathematics is *knowledge-how*. On this view, the knowledge that mathematicians are after is knowledge of *how* to solve problems, *how* to prove theorems, *how* to analyse data etc.

However, while I thought that Rav’s argument was sufficient to show that knowledge-how is of serious interest in mathematics in its own right, this is not to the exclusion of knowledge of mathematical truths and propositions. Instead the two are closely connected, in that mathematical knowledge-that of truths and knowledge-how of methods, techniques and strategies, are not easily pulled apart in practice. What do I mean by this exactly? Well, the thought is that each would be severely diminished without the other,

¹⁰Patrick Greenough has commented on this section that it should be noted that there is a coherent position for a kind of Ravian intellectualist, where the focus is on practical modes of presentation for mathematical knowledge-how as a kind of propositional knowledge. This doesn’t seem to be compatible with Rav’s own views due to the other material he has against Formalist-Reductionist claims, standing against there being underlying propositions for proofs or knowledge of them. However, there is room for this position in the debate and it does appear to be open for proponents of intellectualism to explore, though I shall not be doing so here.

both in the history and practice of mathematics. I want to now spend some time on various examples of how the distinction between knowledge-how and knowledge-that will be reflected in mathematics, as well as where we can see the interaction between the two taking place.

For one thing, a focus on knowledge-how in mathematics will contribute to a better picture of mathematics education. Work in maths education has shown a significant awareness of the need to teach students practical knowledge and skills. A prime example is found in the work of Gila Hanna who has over 30 years examined the importance of proof and proving in mathematics education, such as in (Hanna 1989), (Hanna & Jahnke 1996), (Hanna & Barbeau 2008) and (Hanna 2014).¹¹ In fact, in (Hanna & Barbeau 2008), they engage with the Ravian view of proof and weigh up its importance for mathematics education:

We argue that what is true of mathematics itself may well be true of mathematics education: in other words, that proofs could be accorded a major role in the secondary-school classroom precisely because of their potential to convey to students important elements of mathematical elements such as strategies and methods. (Hanna & Barbeau 2008, p. 352)

They make the argument for this around two case studies of the benefits of teaching particular mathematical proofs, one of the quadratic formula (which I will discuss separately shortly) and one concerning angles inscribed in circles. The point is that learning proofs is an important way of also learning strategies that take us beyond merely learning the truth of the theorem, as well as allowing us to come to a more rounded understanding of mathematics.

There is a major Rylean point here about the process of learning: that learning a subject is often about being *inducted into the practice* of that subject rather than merely learning the truths associated with it. Ryle even includes mathematics as an example of this phenomenon:

The fact that mathematics, philosophy, tactics, scientific method and literary style cannot be imparted but only inculcated reveals that these too are not bodies of information but branches of knowledge-how. They are not sciences but (in the old sense)

¹¹This is only a representative sample; Hanna has over forty relevant papers.

disciplines. The experts in them cannot tell us what they know, they can only show what they know by operating with cleverness, skill, elegance or taste. The advance of knowledge does not consist only in the accumulation of discovered truths, but also and chiefly in the cumulative mastery of methods. (Ryle 1946, p. 15)

Actually, Ryle goes on further to suggest that knowledge-that rests on prior knowledge-how, as both discovery and deployment of our propositional knowledge requires practical knowledge of how to discover and where the knowledge fits into the wider framework.

Effective possession of a piece of knowledge-that involves knowing how to use that knowledge, when required, for the solution of other theoretical or practical problems. (Ryle 1946, p. 16)

The case for this in mathematics is particularly strong, as even understanding the language of mathematics is about knowing what can be done with the various concepts deployed. For instance, as one of the first things children learn, the ‘+’ symbol is directly associated with learning the process of adding numbers together and it is hard to imagine understanding what it means independently. Of course, the inculcation into mathematical practice is not a one-off event, but a continuing development of knowledge of mathematics, both practical and propositional.¹² There are some skills, abilities and pieces of know-how which are more general and others which are topic-specific, but this does not affect the point that mathematics involves the cumulative mastery of methods as well as knowledge of theorems and propositional statements.

That learning mathematics involves being inducted into practices, practices which involve both knowledge-how and shared items of propositional knowledge, has clear impact on the claims from the previous chapter concerning mathematical concepts. Something which I quietly avoided flagging up earlier was that Rav’s idea of what proofs give us knowledge of (besides methods, skills, interactions and systematisations) included knowledge of

¹²As well as much more besides: being inducted into the practices of mathematics involves all kinds of additional learning, such as how to behave at conferences; how to present proofs on a blackboard; which journals to send which papers to; which mathematicians are helpful, rigorous, friendly, quick at responding to emails; which funding bodies to apply for grants from etc.

mathematical concepts. There is room here for a strong stance on the nature of concepts, where concept possession is about the ability to make use of them in various ways, such as in distinguishing whether some object belongs in the extension or anti-extension, deploying them in inferential moves, describing them correctly in appropriate linguistic settings, recognising the relations between them and other concepts etc. Such a move would bring concepts and knowledge-how close together, via the tight link between abilities and knowledge-how. I refrain from leaping into this discussion fully, but it certainly does not strike me as implausible in the mathematics case. In particular, the prominent place of informal concepts, open-texture and domain-specific reasoning in mathematics is suggestive of the idea that coming to know how to do mathematics involves coming to a tacit understanding of the sort of activities which are acceptable to carry out. I shall return to this point later in section 4.6.

Thus far, we have seen that on the picture I am presenting, mathematical knowledge-how is frequently prior to propositional mathematical knowledge. However, I also want to demonstrate that the Wigginsian observation is in full effect and the relationship goes the other way too, such that knowledge-that and knowledge-how are interdependent. Following Hanna & Barbeau, consider the example of the quadratic formula, i.e. that the solutions to an equation of the form $ax^2 + bx + c = 0$ are given by $x = (-b \pm \sqrt{b^2 - 4ac})/2a$. The point Hanna & Barbeau make is that learning the proof that the quadratic formula will always deliver the roots of a quadratic equation can teach a student the skills involved for several related techniques, such as the “completing the square” method and applications to examples beyond quadratics, such as quartic equations of particular forms. But, as a parallel point, the propositional knowledge of the truth of the theorem does deliver the knowledge of an easy way to solve a whole class of problems, one which a struggling student can perform almost mechanically even if they don’t understand the reason that it works. That student can now know how to solve more quadratic problems than they did before. Indeed, if their difficulty is localised to just quadratics, that student might even be able to solve much more complex problems that require solving quadratics as a part. Obtaining the propositional knowledge of the quadratic formula acts as a key to unlock further knowledge-how.

To take a more advanced example, consider *equivalence* and *duality of*

categories. These are relations between categories that show them to be “essentially the same” in the case of equivalence, or equivalent to the “opposite” in the case of duality. The power of such results is immense in their ability to bring out connections between seemingly disparate areas of mathematics and to transfer theorems easily from one to the other without a fresh proof. This holds in a very strong sense, as Mac Lane puts it:

For more complicated theorems, the duality principle is a handy way to have (at once) the dual theorem. No proof of the dual theorem need be given. We usually even leave the formulation of the dual theorem to the reader. (Mac Lane 1998, p. 32)

In general, category theory thrives on these kind of links, and there are a large number of theorems about duality between categories. For example, *Stone’s representation theorem* gives an isomorphism between Boolean algebras and certain topologies on sets (in particular: a topology on the set of ultrafilters of the Boolean algebra) and *Birkhoff’s representation theorem* does the same for distributive lattices and partial orders. Generalising, Stone duality refers to the broader class of categorical dualities holding between topologies and partially-ordered sets, which allows us to move between different disciplines while straightforwardly transferring theorems. The philosophical significance here is that there is once again the lock-and-key phenomenon going on of knowledge—that providing the means to open a whole new range of methods and puzzle-solving techniques. While certainly it requires some background to establish dualities, the interesting mathematics lies not necessarily in the proofs or the methods used in the proof, but rather in the fact that the establishing of the representation theorems allows us to think about certain structures in two distinct but equivalent ways. The usefulness of this is emphasised by Abramsky as a ‘creative ambiguity’:

Mathematically, this distinction can be related to the duality between points and properties, in the sense of Stone-type dualities: the duality between the points of a topological space, and its basic “observable properties”—the open sets. The particular feature of domains which allows this creative ambiguity between points and properties to be used so freely without incurring any significant conceptual confusions or overheads is that

basic points and basic properties (or observations) are essentially the same things. (Abramsky 2008, p. 494)

The Ravian picture, on which the theorem is the ‘headline’ to go with the interesting parts of mathematics which are embedded in the proofs, falls short on the example of representation theorems and Stone duality, in that the interesting mathematics does not reside in the relatively mundane proofs of the theorems, but instead in the new connections one can draw once the theorem is in place and the Gestalt-shifting in viewing well-known structures in entirely different ways. Knowledge of how to prove the theorems is an important discovery that establishes the truth of the duality and gives us knowledge thereof, but it is the latter knowledge of the truth of the theorems which is primary in opening up the new connections which can subsequently be drawn. The knowledge that is discovered about the vast network of connections between different mathematical structures is interesting and might well be entirely propositional. The propositional knowledge of these connections then opens up the scope for a whole range of additional methods, techniques and results, once again supporting the idea that the mathematical knowledge is best understood in terms of *interconnected* propositional and practical knowledge.

4.5 Löwe & Müller on Mathematical Skills

I am not the first to pick out the fact that Rav’s claims about the interesting knowledge of mathematics fits directly into the framework of knowing-how as found in epistemology. Löwe & Müller in two papers (Löwe & Müller 2008, 2010) propose a picture of mathematical knowledge as context-dependent which draws on mathematical knowledge-how as defined in terms of mathematical skills. In this section I will set out the arguments they give and critically assess them.

The papers begin with Löwe & Müller setting out their main claim as follows:

We argue that mathematical knowledge is context dependent. Our main argument is that on pain of distorting mathematical practice, one must analyse the notion of having available a proof, which supplies justification in mathematics, in a context dependent way. (Löwe & Müller 2008, p. 91)

The argument to establish this claim is especially relevant to us in that it draws heavily on mathematical practice, and in particular the distinction between formal and informal proofs, placing these in the context of their relationship to mathematical knowledge. The idea is to look at the “standard” account of mathematical knowledge from proofs, given by:

K1 *S* knows that *P* iff *S* has available a proof of *P*.

To then arrive at a mathematical contextualism, the thought they have is to show that both the notions of ‘availability’ and ‘proof’ in this account must vary with context. In order to demonstrate this, Löwe & Müller run through a series of explications of K1, to show that any invariantist sharpening of the two notions is doomed to fail to properly match ascriptions of mathematical knowledge as found in practice.

The first observation is that K1 cannot be read as merely requiring physical access to a proof, else standing in the Mathematics Departmental Library would turn even the slowest dimwit into a mathematical genius. Instead, the idea behind K1 must be spelled out with “a modalised reading in which the epistemic subject *S* plays an active role” (Löwe & Müller 2008, p. 92) such as the following:

K2 *S* knows that *P* iff *S* could in principle generate a proof of *P*.

Again, this leaves the invariantist about mathematical knowledge needing to fix readings for ‘in principle’, ‘generate’ and, of course, ‘proof’. Following the Formalist-Reductionist, one could fix ‘proof’ as ‘formal proof’. In that case we can look to Formal Mathematics to see how long formalisations have taken and use that as an benchmark for how long one should be given to generate a formal proof for K2. Löwe & Müller cite certain Coq formalisations as having taken ten years, suggesting that therefore ‘could in principle generate’ might be best be set within such a timeframe:

K3 *S* knows that *P* iff , given ten years, she could write a formal derivation of *P* in the language Coq.

But they point out that this is far too generous: given ten years, the mathematician could *learn* Coq from scratch as well as a large number of theorems previously unknown to her (Löwe & Müller 2008, p. 99). Clearly K3 will be inadequate. It should also be clear that fixing the length of time some other

way will always be too restrictive (ruling out the hard cases, like early Coq formalisations) or too permissive (ascribing knowledge to those who clearly don't possess it).¹³

In a different direction, Löwe & Müller suppose that 'proof' could be taken along the lines of 'informal proof on a blackboard':

K4 S knows that P iff , given a blackboard and a piece of chalk, she is able to produce an acceptable blackboard proof within an hour.

But once again, setting an exact timeframe—such as the one hour above—will be problematic. The time cannot be too short, because there are normal cases of mathematical knowledge where the mathematician needs to refresh the details if they are to write up a proof on the board:

They need to try one or two standard approaches to tackle the problem, remember the important details, and only after that are they able to provide an acceptable proof. (Löwe & Müller 2008, p. 99)

Too long, however, and someone can in theory have time to figure out something new; they might even believe $\neg P$ and by working for the one hour arrive at the opposite belief. It would certainly be undesirable for us to claim they knew P despite believing $\neg P$.

The upshot is that we should take mathematical knowledge to be linked to the salient context. Any attempt by the invariantist to fix the key concepts in a rigid way is open to refutation by pointing either to a context in which possessing mathematical knowledge is very demanding which the definition is too strict for, or one in which mathematical knowledge is made too easy by it. The argument is that there is no way in which we can fill out the details of the link between proof-possession and mathematical knowledge that isn't inextricably linked to the context.

Löwe & Müller next suggest that mathematical knowledge in their contextualist picture should be explicated in terms of *mathematical skill*:

K5 S knows that P iff S 's current mathematical skills are sufficient to produce the form of proof or justification for P required by the actual context. (Löwe & Müller 2008, p. 104)

¹³They also consider another version based on Steiner's 'midwife logician' idea (Steiner 1975). In this case 'could in principle generate' is set as 'could with aid of a midwife logician produce'. Problematically, as indicated by Löwe & Müller, this blurs the line between some agent's knowledge and the knowledge of the midwife logician.

In their follow-up paper (Löwe & Müller 2010), Löwe & Müller expand on what they mean by mathematical skills, with reference to the Rylean picture and the idea that Rav's claim is about the primacy of knowledge-how in mathematics. However, they quickly set this aside to focus purely on skills as professional skills, suggesting that Ryle took knowledge-how to be synonymous with skill and following him in this usage.¹⁴ They have a lot of useful things to say concerning the nature of professional skills in mathematics, drawing explicit parallels with the case of nursing as a profession:

The notion of a skilled nurse is related, ultimately, to a nurse's job description, which has developed historically. We are not concerned here with a natural kind of human beings, nurses, of which there are more and less skilled ones. Rather, we are assessing human beings who have chosen a specific profession, as more or less skilled as required by the (historically and sociologically contingent and changing) requirements of that profession. Nursing skills are *professional skills*. (Löwe & Müller 2010, p. 270)

They point out that in mathematics there are a number of complex and context-dependent issues surrounding the nature of mathematical skills. Firstly, there are a whole range of mathematical skills needed in being a professional mathematician, ranging from almost essential ones used in doing mathematics and mathematical reasoning, to fairly relevant ones such as giving talks and engaging in informal chats, to mostly peripheral skills such as filling in expense forms or adjusting to jet-lag. Secondly, there is a question of the granularity of mathematical skills and how one individuates them, to which they answer:

Mathematics is one subject, and for most purposes, it makes sense to view *general* mathematical skills as the pertinent level of granularity. For purposes of assessing knowledge claims, local dimensions of skill may however also play a role, depending on context. (Löwe & Müller 2010, p. 274)

Due to the contextualist picture, context (as well as the particulars of the theorem P) plays an important role in picking out the relevant skills in the

¹⁴I don't think this is a correct reading of Ryle, something I shall return to momentarily.

picture of mathematical knowledge, as well as the extent to which the generality of the skills holds up. It might well be that there are theorems which require very particular topic-specific skills, for example. Finally, there is an interesting set of questions concerning the measurement and assessment of skill. In particular, the thought is that skills must go beyond particular performances, as even the most highly skilled individual can go awry occasionally.¹⁵

So let us assess the picture presented by Löwe & Müller. In general, I find the direction appealing and believe that it is in the same vein on many issues as the picture of mathematical knowledge in terms of virtue epistemology to be given in the next chapter, especially with respect to making the mathematician play an active role in their epistemic state. I also share the general naturalistic methodology of beginning with mathematical practice, and think that context and skill must play a role in mathematical knowledge.

However, there are points of contention too. First of all, I do not think it is correct to equate knowledge-how with skill, and have argued that this is not the Rylean picture either. As we saw above, Hornsby makes a convincing case that Ryle was not interested in any sort of reductive analysis such as this, be it to abilities, capacities, skills etc. Such exegesis doesn't hinder the philosophical claims, of course, but does cast doubt on how well the work follows through on the Rylean project as it applies to mathematics.

Secondly, one obvious place to look for criticism of the Löwe & Müller position is in the usual invariantist responses to contextualism, which would seek to explain away the cases of shifting knowledge ascriptions varying with context. One way to do this is to argue that mathematicians might well commonly engage in "loose talk", where they use the term 'knowledge' but actually mean something weaker. They might go so far as to reserve the term 'knowledge' for the idealised, Formalist-Reductionist view, and argue that anything less than this does not meet the high standards we have for mathematical knowledge. Mathematical knowledge from proof, after all, is singled-out as special precisely because of its rigid, deductive form of jus-

¹⁵This point will be extremely important in the next chapter, where skills will be connected to intellectual virtues. For some virtue epistemologists, possessing the right virtues is necessary for knowledge, but we see from the current point that this can't be the whole picture, as neither skills nor virtues generally suffice to make someone immune to performance errors.

tification and, one might argue, if this is lacking then what results does not deserve the title of ‘knowledge’. However, I think Löwe & Müller are right not to accept this, as they are clear in their naturalistic methodology which gives mathematical practice a primary role. The strict account which the invariantist might offer would simply rule out so much of our claimed knowledge of mathematics that it would be a Pyrrhic victory: the concept of knowledge arrived at would be neither the one actually in use nor particularly useful for any mathematical purposes.

A more pressing worry, in my opinion, is about the details of modalising mathematical knowledge through the notion of skill. With skills, and knowledge-how more generally, it makes a great deal of sense to modalise in a way that means we know how to do things that we have not previously done or even considered doing, because the very notions of knowledge-how, skill and ability are bound up with success or reliability across a range of counterfactual scenarios. Furthermore, considerations of granularity mean that an application of knowledge-how can almost always be to a new situation if we describe it finely enough. However, the worry is that Löwe & Müller transfer this modalisation to mathematical knowledge generally, and in particular to propositional pieces of knowledge, which is their main target for the contextualist approach they are advocating. To emphasise, their final account is based around the following thesis:

K5 *S* knows that *P* iff *S*’s current mathematical skills are sufficient to produce the form of proof or justification for *P* required by the actual context. (Löwe & Müller 2008, p. 104)

My concern is that this principle assigns mathematicians too much knowledge, even with the contextual restrictions that they build in. For the context here only supplies the form of the proof or justification needed to know that *P*, ranging from a rough proof-sketch to a full formal derivation (Löwe & Müller 2010, pp. 274–275). The difficulty, however, is that mathematicians may possess many skills which don’t then get applied to some given proof for a theorem *P*. The fact that *S* possesses all the relevant skills is sufficient to satisfy K5, even if *S* never goes through the motions of actually proving *P*.

Indeed, we saw above that performance errors are a reason for talking of skills more generally rather than the instantiations of them, but similarly

they are a reason that having skills won't suffice for possessing the deliverances of successful performances of those skills. On K5, it would still easily be possible to possess all of the relevant skills to produce the level of proof of P required by the context, and thereby to know P on K5, despite believing $\neg P$. For example, adopting one of Löwe & Müller's own cases, imagine a student S going into a maths exam who has studied all of the techniques and methods of the course, but just before entering is told by a usually trustworthy fellow student "psst don't forget that $\neg P$ ". While S does have all of the relevant skills to produce a proof of P up to the standards of the exam, she might well believe $\neg P$ on the basis of the testimony. Simply put, K5 allows for someone to know something they simultaneously believe to be false— a most undesirable result.

The answer to this problem, as I see it, is to go beyond skills to require successful performances or manifestations of knowledge-how to obtain propositional mathematical knowledge. We shall explore the epistemology of this in the next chapter. However, what I would now like to emphasise in the coming section is that this requires observing the importance of *proving as an activity*, and it is going through the actions built into a proof which is what secures knowledge of the truth of the theorem proved.

4.6 Proving in Action

In this section I will consider the difference between proofs themselves and the activity of proving, specifically with respect to their contribution to mathematical knowledge. The focus on actions in proofs builds on previous work by Larvor in (Larvor 2012), so let us extract some key points he makes first.

4.6.1 Larvor's Inferential Actions

In (Larvor 2012), Larvor sets out the case for the existence of *essentially informal arguments* (from which we have essentially informal proofs too, as those proofs which involve essentially informal arguments). Hereby we can identify a substantial area which is not covered by the traditional approaches to the philosophy of mathematics, which also requires investigating mathematical practice. The philosophy of mathematical practice can then gain traction and make the perspectival shift being advocated clear. Larvor's

main proposals are that essentially informal arguments are distinguished from formal arguments by their *content-dependence* and, furthermore, that a proper account of content-dependence will include a broadening of the picture of inference beyond merely linguistic argument to a full class of *inferential actions*. This is because the content of mathematics, and that which we act on inferentially, is not limited to language alone.

In what way are essentially informal proofs content-dependent? Larvor argues that the content of a proof connects to some domain the proof is located in, and that this domain has some class of acceptable inferences that can be employed in proofs in this domain. A proof is then valid if all of the inferences used in the proof are acceptable in the domain, and have been applied properly etc. Importantly, formal rules are acceptable across all domains (such as modus ponens)¹⁶, but Larvor's picture allows us to also have domains with more contentful moves which are not generally applicable across all domains. While this might mean some inferences are specific to a very restricted domain, many are in fact acceptable across a broad range of domains without this being so broad as to include all domains. A content-dependent proof will then make use of these content-dependent inferences.

Further clarification is needed, though. For instance, the above has not yet told us about what content is or what counts as a domain. We would be in big trouble if the domains were so fine-grained to have it that each purported proof is located in its own domain, with the acceptable inferences being precisely those employed, as this could trivialise mathematics and the whole notion of proof. The right answer seems to be that domains are particular areas of mathematics, with particular frameworks for acceptable inferences established through mathematical practice, though it would be wise to follow Löwe & Müller in including some context-dependence in this. The content is merely the subject-matter of the proof, which connects to the domain in the straightforward sense that, for example, a geometrical proof reveals that we are working in the domain of geometry and thus authorises the use of geometrical moves in the proof.¹⁷

¹⁶There is, here, the obvious question of which logic determines the rules that are acceptable across all domains, or whether there even is such a thing. I don't think it would be a bad thing for the position being advocated here if there is indeed no universal background logic. The logical pluralist in me certainly thinks so. Nonetheless, I will set aside these issues as outside the scope of the current discussion.

¹⁷We should be careful, of course, to avoid the subject-matter merely being comprised of some set of acceptable inferences, as this threatens circularity.

The account of informal proofs as depending on content as well as their form does lose us some nice features that a Formalist-Reductionist account would have. For example, if all of our proofs could be reduced to formal proofs alone, then the required logic could remain topic-neutral, which now will be explicitly given up. Likewise, Larvor points out that “we have to abandon the hope of establishing a general test for validity” (Larvor 2012, p. 723). We shouldn’t be unhappy to see these go, however, given that they are so closely linked to formality which we have good reason to reject as the right account of informal proofs (see the first half of this thesis). Rather, it will become clear that these features can have no general place in the new, more dynamic approach to mathematical proof conceived of in terms of content and action.

The second main proposal by Larvor is a switch to emphasising the activities involved in inferring, arguing and proving. Above, we saw a number of mentions of the acceptable moves, steps and inferences in some given domain. Regarding these, Larvor says:

If we think of an argument as a sequence of propositions connected by logical relations, it is hard to see how the content of the argument can play a role in the step from one proposition to the next. This is in part because a classically trained philosophical imagination is dominated by general logic, but also because orthodox philosophical education urges us to forget that the movement from one line of a proof to the next is an *action*. (Larvor 2012, p. 721)

Larvor argues that we should recognise the purely propositional framework as being too limited to properly account for actual arguments found in mathematical practice and mathematical proofs. The point is not merely that we should recognise the actions involved in moving between propositions, but rather that adopting such a focus reveals that the objects of our actions actually form a much broader class than just propositions.

The liberating insight is to notice that in making arguments, we act on all sorts of items in addition to propositions and well-formed formulae. Sometimes, we act inferentially on non-propositional representations of the subject-matter such as diagrams, notational expressions, physical models, mental models

and computer models. (Larvor 2012, p. 721)

More specifically to informal proofs, the kinds of steps found in mathematics are not limited to actions on propositional contents, something Larvor illustrates with a series of examples from diverse areas of mathematics. Indeed, it may even be that these actions do not have objects at all, such as if the subject-matter is the manifestation of the action. Larvor's example of the last point is the demonstration by a gymnast that some complex gymnastics move is possible by performing the move.

The framework being proposed by Larvor, then, is to see proofs as *systems of inferential actions*. This is far removed from the alternative, traditional view of proofs as abstract objects made up of sequences of propositions.¹⁸ Inferential actions are just those actions which can be used in arguments and, in the mathematical case, proofs. Of course, as described above, the inferential actions acceptable for some particular proof depends on the domain the proof is in.

There is a strong dose of Lakatos and Kneebone in this conception of inferences as found in actual proofs. We can view the key point as being that we should switch from a static conception of proofs to a dynamic one. While the static conception is primarily concerned with the stops along the way and the stepping-stones through the proof, in contrast, the dynamic view is concerned with the movement through the proof and the actual steps being made, as it is these which ultimately take us through the proof to establish a theorem. That isn't to say that the places we stop aren't important. The full picture that should emerge of informal proof will be one which takes account of how the non-propositional actions found in the proof relate to the propositional content of that proof as in the Wigginsian anti-intellectualist view of Ryle. This aligns with the approach to practical and propositional knowledge being argued for in this chapter.

Again, the move to the action-oriented perspective gives up on certain desirable features, especially when combined with validity as content-dependent. Primarily, unlike in formal logic, there is now no general test

¹⁸Proofs and arguments as abstract objects is not restricted to the formal proofs and arguments. For example, Leitgeb takes informal proofs to be abstract objects:

[...] we regard mathematical proofs *per se* as abstract entities which are independent of any material instantiation. (Leitgeb 2009, p. 266)

Similarly, (Simard Smith & Moldovan 2011) treats arguments as abstract objects (although abstract objects which can come into existence and disappear again).

for validity because there is no full and final list of inferential actions. Even when limited to some domain, if that domain is anything beyond the simplest cases then it will simply not be possible to fully settle all of the actions that might be permissible within that domain. Although this is certainly a cost of the view, as Larvor points out, it is a rather mild one. This is so because it is something that we have been lead to believe by the Formalist-Reductionist tradition that we will get, yet isn't really something we should be expecting once we pay attention to proof in practice. Indeed the open-ended nature of mathematics and mathematical methods is vitally important to its growth and development, as we saw in the previous chapter.

Let us briefly mention the place of *rigour* as it is sketched in Larvor's view. Larvor says that

[F]or every kind of inferential action, there must be a corresponding means of control, to ensure rigour. Sometimes these controls are simple rules like 'do not divide by zero'. In other cases, these controls may be the fruit of mathematical research [...] Demonstrating rigour involves making the controls on inferential acts explicit, which is why some diagrams disappear from the final published version of a mathematical argument. The problem is not with diagrams as such, but rather that the actions performed on these diagrams in this piece of work do not have established, agreed controls. (Larvor 2012, p. 728)

Such controls are important— mathematicians should be careful not to divide by zero or abuse diagrams and infinite series. Larvor is right that it is often a fruitful project to make these explicit and that this is connected to mathematical rigour. However, we might be sceptical that rigour is fully accounted for by such corresponding controls, for I take correct and rigorous proving to be connected to practical knowledge and it has been argued that this is not fully enumerable in terms of explicit rules or principles in any reasonable sense. Just as in Ryle's point that there may be regulative propositions, rules and maxims which apply to practical knowledge, but these cannot be the whole of what knowledge-how amounts to, nor should we expect there to be a particular list of controls which ensure rigour. As such, we should not expect to be able to *demonstrate* rigour in the way described either, unless it reduces to the Formalist-Reductionist position,

which is not what is intended. If not, though, we need to answer the further question of how we will ever manage to be confident that there is not some further ‘hidden’ rule that we violate in a given proof?

We can also wonder if such controls will form a unified class at all, or whether there is a range of different principles, between those necessary for rigorous proving and those which are merely good proof etiquette. For example, if I were to switch languages (from English to Japanese, to Afrikaans etc.) between each line of a proof, is this a lack of rigour or just poor style? Just like the open texture of mathematical concepts, for any given mathematical domain it might well be that there is a never-ending horizon of ways to mathematically misbehave. The way these are avoided is not about implicit rules, but about learning how to behave well. This is not to say that there are no such rules; in line with the arguments from Ryle we can extract them from practice and describe them in exactly the same way that we can identify logics which our practices cohere with. We certainly have a great deal of rules and heuristics taught in classrooms and lecture halls, for example. The point is just that there is something more than this to rigour. In the next chapter I will argue that the correct account of mathematical rigour should connect it to intellectual, *mathematical virtues*.

4.6.2 Proving as an Activity

Proofs play numerous roles in our mathematical practices and serve many different functions, but one I am here primarily interested in is their role in mathematical epistemology. Previous considerations of mathematical knowledge seem to have paid little attention to the idea that the knowledge we get from proofs is arrived at by the *activities of proving*. Even the Löwe & Müller papers discussed above, which seem to be going in the right direction, start from an idea of merely having access to a proof and finish on a modalised notion which grants us knowledge of everything we are skilled enough to do with respect to a context, without requiring us to actually do the work of obtaining the knowledge we have. Larvor is correct to emphasise that the movement through a proof involves inferential actions, but his focus lies elsewhere and the paper does not make explicit the impacts this has on mathematical epistemology.

The important idea is that it is the *activity of proving* which is of primary epistemological significance in mathematics, with *proofs* themselves of

secondary importance. Outside of philosophy this has been picked up on in more popular reflections on mathematics, such as by Marcus du Sautoy in (du Sautoy 2015), as quoted at the start of this chapter. He sees proofs as narratives, describing the journey across the mathematical terrain from familiar and well-trodden starting points to far-off realms. He continues as follows:

Within the boundaries of the familiar land of the Shire are the axioms of mathematics, the self-evident truths about numbers, together with those propositions that have already been proved. This is the setting for the beginning of the quest. The journey from this home territory is bound by the rules of mathematical deduction, like the legitimate moves of a chess piece, prescribing the steps you are permitted to take through this world. At times you arrive at what looks like an impasse and need to take that characteristic lateral step, moving sideways or even backwards to find a way around. Sometimes you need to wait for new mathematical characters like imaginary numbers or the calculus to be created so you can continue your journey. (du Sautoy 2015)¹⁹

Reading this, there is a touch of formalistic thinking in the further analogy to moves in chess²⁰ which we wouldn't want to take too seriously, but the notion of a journey fits very well with the thought that we should emphasise the activity of proving. The quote is, of course, very reminiscent of the well-known picture from G. H. Hardy:

I have myself always thought of a mathematician as in the first instance an *observer*, a man who gazes at a distant range of mountains and notes down his observations. His object is simply to distinguish clearly and notify to others as many different peaks as he can. There are some peaks which he can distinguish easily, while others are less clear. He sees A sharply, while of B he can obtain only transitory glimpses. At last he makes out a ridge which leads from A, and following it to its end he discovers that it culminates in B. B is now fixed in his vision, and from

¹⁹I am grateful to Ursula Martin for pointing me to this article.

²⁰That old chess-nut.

this point he can proceed to further discoveries. In other cases perhaps he can distinguish a ridge which vanishes in the distance, and conjectures that it leads to a peak in the clouds or below the horizon. But when he sees a peak he believes that it is there simply because he sees it. If he wishes someone else to see it, he *points* to it, either directly or through the chain of summits which led him to recognise it himself. When his pupil also sees it, the research, the argument, the *proof* is finished. (Hardy 1929, p. 18)

The point of these rather long quotes is that the metaphor of mathematics as a huge landscape has been drawn on before. However, du Sautoy's way of speaking is preferable to Hardy's for now, as it brings out the active nature of proving rather than the more passive language of the 'observer'. The claim I am making is that if we want to attain knowledge, this requires finding and following the path to get there. We can think of proofs, via a similar metaphor, as maps or directions providing us with a guide as to how to get from one place to another, from A to B. While the activity of proving is about traversing the mathematical landscape, a proof provides a record of the series of actions required to reach a new mathematical location and is used to communicate what the discoverer of the proof went through to others who wish to follow the same road and gain the same mathematical knowledge.²¹

Indeed, there is clear linguistic evidence for such an idea in the fact that much of the standard terminology in informal proofs is imperatival.

²¹I have recently discovered, thanks to Josh Habgood-Coote, that Ryle actually uses the same metaphor for mathematical discovery:

[...] the pioneering path-finder, Pythagoras say, has no tracks to follow; and any particular sequence of paces that he tentatively takes through the jungle may soon have to be marked by him as leading only into swamps or thickets. All the same, it may be, though it need not be, that in a day's time or a year's time he will have made a track along which he can now guide docile companions safely and easily right through the jungle. How does he achieve this? Not by following tracks, since there are none to follow. Not by sitting down and wringing his hands. But by walking over ground where tracks certainly do not exist, but where, with luck, assiduity and judgement, tracks might and so perhaps *will* exist. All his walkings are experimental walkings on hypothetical tracks or candidate-tracks or could-be tracks, or tracks on appro; and it is by so walking that, in the end, while of course he finds lots and lots of impasses, he also finds (if he does find), a viable track. (Ryle 1971, p. 224)

Common terms are ‘let’, ‘assume’, ‘suppose’, ‘define’, ‘construct’, ‘observe’, ‘consider’, ‘reduce’, ‘rearrange’, ‘note’ and many more.²² Once again, we find ourselves with a point familiar from Ryle:

We certainly can, in respect of many practices, like fishing, cooking and reasoning, extract principles from their applications by people who know how to fish, cook and reason. Hence Izaak Walton, Mrs. Beeton and Aristotle. But when we try to express these principles we find that they cannot easily be put in the indicative mood. They fall automatically into the imperative mood. (Ryle 1946, pp. 11–12)

A proof thus tells the reader what to do in order to prove some theorem, and thereby makes one important role of proofs to guide us through the inferential actions needed to get to a certain place. An equally good analogy, then, would have been recipes in cookbooks: the recipe itself is only important in so far as it directs you how to make the cake in question. While we can talk about better or worse recipes, this is derivative on how well it guides us through our baking activities.²³ In (Robinson 1991), Robinson describes proofs in a way similar to du Sautoy’s narrative idea:

[A] kind of meaningful narrative [...] more like a story, or even a drama, conveyed to us in language calling on our semantic and intuitive understanding. [...] To follow an informal proof as it unfolds in time is to understand the story as it develops. (Robinson 1991, p. 269)

Now certainly this seems right in certain respects, but it appears to suggest that we are passive observers to the unfolding drama with our understanding just being used to follow the action from afar. On the contrary, I take understanding a proof to involve being part of the action. The proof tells us which actions to take; the mathematician acts them out. Proofs thus don’t operate in a vacuum, securing their targets in the abstract, but rather they are secondary to the mathematical activities they guide us through,

²²This should be familiar to anyone who has looked at mathematical proofs, but I invite anyone sceptical of this to open up a few recent pure mathematics articles on the ArXiv and check for themselves. An interesting study to carry out in the future would be to do a proper analysis of some body of real proofs.

²³Though the baking analogy works slightly less well because we also judge recipes for the tastiness of the baked goods they produce.

activities which are themselves embedded in a practical context and carried out by agents.

To conclude, by seeing proofs from this *action-centred perspective* we can thus say something new about the relationship between proofs and mathematical knowledge. Gaining knowledge of mathematics from proofs is actually done through the activities of proving: blazing a new trail through the mathematical landscape or following the paths that others have set, by following the instructions they have given us in their proofs. The epistemic importance thus lies primarily with the activities and actions, not the proof itself or its mere existence. The Formalist-Reductionist project misses out on this crucial idea, according the primary importance to proofs themselves, or worse still to unaccessed or inaccessible formal proofs underlying them, thereby failing to correctly explain the epistemic role of proofs. We, on the contrary, are in a place to explain how mathematical knowledge actually connects to proofs and proving activities. This will be the topic of the next chapter.

4.7 Conclusion: Knowing How to Prove It

To finish this chapter, let us briefly return to the relationship between knowledge-how, knowledge-that and proofs. Following Rav, I have argued that besides knowledge-that of mathematical facts, theorems and propositions, there is also knowledge-how of methods, tricks, techniques, interrelations and more besides, something which has not received proper attention in mathematical epistemology. Rav's idea, put in the epistemological terminology, was that knowledge-that of theorems in their statement-forms is of less interest in mathematics than knowledge-how as embedded in proofs. I have been arguing, though, that in mathematics knowledge-how and knowledge-that are actually very closely linked, with each delivering the other in a range of cases, in a way that does not entirely track the theorem/proof divide. Additionally, I have been arguing that while knowledge-how might be modal, in that it relates to how one acts and behaves across a range of possible scenarios, the propositional knowledge of a theorem we get from having proved it is not, since proving is an activity which delivers knowledge only when it has been successfully carried out.

An interesting upshot of the positions I have taken is that there are, for

any given mathematician, theorems which they know how to prove without having knowledge of the theorem or its truth. This falls out of the fact that the knowledge of how to prove some theorem T is modalised in such a way that one might possess all of the relevant knowledge-how to produce a proof of it without ever carrying out that proof. Meanwhile, the propositional knowledge of the truth of T demands the stronger condition of actually having proved it, which I have argued involved going through the inferential actions the proof is made up of.

I don't take this to be problematic, though, as this certainly matches up with what we should want from an account of mathematical knowledge. For there are plenty of cases where we might possess the know-how, skill and ability to carry out a proof straight-off, without that meaning that we have any knowledge of the particulars before actually proving the relevant theorem. For example, one might be fully competent with quadratics and know how to solve any given example, without already knowing the roots of $1124723477234x^2 - 3419824x + 1 = 0$. This is exactly as it should be, and even a desirable feature of the position, since it demonstrates that there is a clear reason to want our epistemological theory to include both kinds of knowledge of mathematics.

So what is it to know how to prove some theorem? Just as in knowledge-how more generally, an ability account might seem initially appealing:

Ability 2 S knows how to prove theorem T iff P is able to prove theorem T .

One can even argue that the problematic cases don't apply here. Firstly, unlike knowledge-how more generally where counterexamples arise in the general pattern of pianists who break their fingers etc. which impinges their ability without affecting their knowledge, in mathematical cases there seem to be far fewer physical requisites on being able to prove some theorem. Secondly, the cases of lucky success are somewhat harder to generate for mathematics, as there can be a distinction made between actually proving something and producing a proof of it. The thought is that by emphasising the activity involved in traversing the mathematical landscape, we can exclude cases where a proof is put on paper without the mathematical activity being carried out, such as if Jackson Pollock accidentally flicked paint into a proof of the Riemann hypothesis. Hereby, we can also rule out lucky success

cases.

However, once again this will not fly. With a bit more imagination we can think of fresh counterexamples, such as going blind initially hindering someone's ability to do diagrammatic proofs. As to lucky successes, there are more cases than those of the monkey-and-typewriter variety. For instance, we have probably all witnessed students guess at steps they think the teacher wants to hear but subsequently not be confident that what they have done is correct. This may be a case of knowing how to prove the theorem, but it is a worryingly low bar. Furthermore, the Löwe & Müller discussion provides us with other problems for an ability view, such as filling out how long it should be allowed to take for them to be said to be able to. Or another case: what if a student knows thirty techniques, only one of which will work for the proof, but where the student has no idea which to use or even how to go about deciding between them. In the weak sense they do know how to prove the theorem— they might be able to just by trying out all thirty approaches— but their ability might require far longer to enact than if they knew how to select the right tool for the job, which is itself an important piece of mathematical know-how.

There is also reason to be hesitant regarding a mathematical know-how as mathematical skills position. Löwe & Müller do well by linking the appropriate skills to the context relative to which a knowledge ascription is being made. Nonetheless, there are difficulties concerning what kinds of skills and exactly which they are: even a moderately straightforward proof might require a whole range of skills, but more than this, they need to be combined in the right way to form a complete path of inferential actions. Indeed, it might well be that for plenty of open problems in mathematics, we do already possess all of the relevant mathematical skills, techniques and methods, but simply haven't combined them in the right way yet. Presumably, we don't know how to prove these theorems, despite possessing the relevant skills. Of course, we could avoid this problem by positing that there needs to be some kind of "putting it all together" skill which is also in play. But this move is less than ideal for two reasons. Firstly, such a skill seems to be of a different kind to the general-level skills which Löwe & Müller want, being tied to the specific proof. Secondly, this means that the reduction of knowledge-how to skills has not been particularly informative, as once again it seems like the "putting it all together" skill basically amounts to knowing

how to prove the theorem.

Regarding mathematical know-how, I am broadly inclined towards the Rylean position as expounded by Hornsby, namely to avoid such reductive analyses. Nonetheless, there is a great deal more to be said about the interaction between proving and mathematical knowledge, which I will do in the next chapter where I explore the application of virtue epistemology to mathematics.

Chapter 5

A Virtue Approach to Mathematical Epistemology

I was unable to find flaws in my “proof” for quite a while, even though the error is very obvious. It was a psychological problem, a blindness, an excitement, an inhibition of reasoning by an underlying fear of being wrong. Techniques leading to the abandonment of such inhibitions should be cultivated by every honest mathematician.

— John Stallings (1965) “How Not To Prove The Poincaré Conjecture”

5.1 Introduction

In the paper quoted above, Stallings describes the mathematics behind a failed attempt he made at proving the Poincaré Conjecture. Prior to Perelman’s 2003 proof, this was one of the best known open problems in mathematics, so we can certainly sympathise with the excitement and fear expressed by Stallings at the prospect of having solved it. The fatal error in the proof stems from proving a key auxiliary theorem, named ‘Theorem 0’, for all cases of $n > 2$ but then subsequently making essential use of it for a case where $n = 2$, a case for which it is demonstrably false, shown by a counterexample Stallings provides.

My topic in this chapter will be *mathematical virtues and vices* and how these play a crucial role in mathematical knowledge. To introduce this

theme, I want to draw out two key aspects of Stallings's words. Firstly, Stallings blames the mistake in the proof on psychological failings, where these are what lead to him not spotting the misapplication of Theorem 0. He blames himself for the error, seeing it as caused by inhibitions of his reasoning, or we might say performance error in his mathematical skills and competences. Putting this explicitly in virtue-theoretic terminology, there are two ways that we might blame *epistemic vices* here: either in his making the error in the first place or in failing to spot and correct it later on. Secondly, the emphasis Stallings puts on the *honest* mathematician wanting to develop techniques to prevent such errors and failures of rigour in future mathematical proofs or, to put it another way, to develop the *mathematical virtues*.

A central idea in this chapter will be that the investigation of mathematical virtues and vices will be able to develop and draw on some of the major lessons learned in virtue epistemology and virtue ethics in recent decades. One key difference which emerges in virtue-theoretic approaches in epistemology, in contrast to other epistemological theories, is that the agents and communities play a central and irreducible role in their own epistemic positions and states. For the current project I will thus be looking at the place of the mathematicians in mathematics.¹ If we take virtue epistemology seriously, mathematical knowledge may be deeply connected to virtuous mathematical behaviours such as those of Stallings's honest mathematician, and conversely we may find that failings in mathematics may at times stem from mathematical vices. Placing such emphasis on the role of the mathematician will have wide-ranging consequences for the philosophy of mathematics, a number of which I will set out and explore in this chapter. What we will see is that virtue theory will be crucial to resolving difficult questions concerning proof and rigour in mathematics, and in particular how they relate to mathematical knowledge.

One might immediately take a dismissive attitude towards the virtue turn for mathematics by arguing that while we do have terms for virtues and vices in mathematics, the explanatory value of these is limited to just a few

¹We can, for simplicity, call everyone engaging in explicitly mathematical activities a mathematician. Nothing hangs on this in the broad overview of the project here, but it may well turn out that contrasting mathematical virtues and vices as possessed by mathematicians and non-mathematicians reveals interesting aspects of how we learn and behave in different mathematical contexts.

aspects of mathematical practice which are already involved with crediting and evaluating mathematicians. Conversely, one might see in the virtue approach the germ of a radical shift in philosophy of mathematics generally. To draw out the large range of views that we may adopt under the general heading of *virtue-theoretic philosophy of mathematics*, I will separate three levels of claims that one can propose, develop and defend:

Moderate Proposal Virtues and vices of mathematicians will be relevant to mathematical knowledge.

Strong Proposal Virtues and vices of mathematicians will be explanatory of mathematical knowledge. In other words, virtue epistemology should be adopted to give the correct epistemology for mathematics.

Radical Proposal Virtues and vices of mathematicians explain mathematical knowledge and extend to provide alternative answers to other kinds of questions in the philosophy of mathematics, e.g. those concerning ontology, access, metaphysics etc.

As stated, these three proposals are only meant to serve as a general guide to different levels of views one might take rather than exact statements or theories. Work needs to be done to establish the meaning behind the three proposals and their plausibility. In this chapter I will focus on the Strong Proposal, beginning with a brief discussion of the Moderate Proposal and returning only briefly to the Radical Proposal in the conclusion.

Incorporating a theory of mathematical virtues into our theoretical landscape of mathematics and our knowledge of it will in many ways be a grand departure from a lot of traditional thinking in the philosophy of mathematics. In section 5.2, I will begin with the Moderate Proposal, showing that the evaluation of skills, competences and character traits of mathematicians is already commonly taken to be relevant to mathematical knowledge. In section 5.3 I will set out the key elements of the virtue epistemology literature, followed in section 5.4 by its application to the case of mathematics as the Strong Proposal. In sections 5.5 and 5.6, I will apply to framework to show that the Strong Proposal also provides a well-motivated account of proof and rigour respectively. Furthermore, in section 5.7, I will show that the virtue approach on the Strong Proposal is very amenable to being extended to more notions in epistemology, in particular the under-explored

area of *mathematical understanding*. I will in section 5.8 examine a very current and ongoing case study concerning Shinichi Mochizuki's proposed proof of the *abc* conjecture, showing that there is indeed a close link to virtue theory: that the approach I advocate is in the best position to account for the surrounding controversy.

5.2 The Moderate Proposal for a Virtue Approach to Mathematics

Let us begin with a look at the Moderate Proposal for incorporating talk of virtues, vices and values more generally into our philosophy of mathematics. The central idea of this proposal is given above as:

Moderate Proposal Virtues and vices of mathematicians will be relevant to mathematical knowledge.

Of course, a great deal already rests on which aspects of mathematical practices we are seeking to explain, so this is something we must consider. Before I introduce the more weighty proposals drawing on virtue epistemology, the talk of virtue and vice can also be understood in different ways. One may talk about *theoretical* virtues and vices on the one hand and *personal* or *agential* virtues and vices on the other.²

The importance of the theoretical virtues to mathematics is already present in the philosophy of mathematical practice literature. To list but a few: elegance, simplicity, generality, unification, applicability, explanatoriness and beauty all affect how mathematics is developed, which areas we find interesting and worth pursuing, and how much we favour some given mathematical proof. Another way of putting this is that such theoretical virtues already play a role in mathematical practice, and already there is a blossoming literature surrounding a number of different theoretical virtues of this kind. Questions concerning mathematical beauty and mathematical explanation (and thereby explanatoriness) are already to be found in the philosophy of mathematical practice in particular. Meanwhile, questions concerning the applicability, simplicity, generality or unificatory power are

²A good reason to prefer the term 'agential' is that the virtues possessed by groups will be explanatorily important in the case of mathematics, and this term stays neutral on groups as agents, whereas 'personal' does not.

commonplace in mathematics and the philosophy of mathematics more generally. If we are interested in the nature and purpose of proofs, theoretical virtues immediately come to the fore. For example, proofs are evaluated for things like elegance, simplicity, intricacy, rigour, explanatoriness and beauty.³ The attribution of theoretical virtues and vices is certainly not limited to proofs either. Mathematical ideas come in many varieties, all of which may be evaluated according to what we value and the like. Besides proofs, these may include concepts, notations, theories, definitions, symbolisms, techniques, ideas etc. For instance, a well-constructed definition can be bountiful for mathematical theorising and streamlining for the proofs and lemmas it is deployed in. Consider, for example, the $\epsilon - \delta$ definitions of convergence and continuity, discussed in chapter 3. Meanwhile cumbersome notations can reduce our understanding or even inhibit mathematical breakthroughs. Nonetheless, let us set aside theoretical virtues to focus on agential virtues. This is not to diminish the importance of theoretical virtues, just that these are not central to the case to be made which focuses on how the mathematician is primary in obtaining mathematical knowledge.

We also can observe that agential virtues are regularly discussed in relation to their importance for mathematics; for instance, in praise of particular mathematicians who have made major contributions to their fields. A prime example of this would be the kind of language used to describe John Conway; no discussion of him seems to be complete without attributing his mathematical breakthroughs to his playfulness, curiosity and light-heartedness. For example:

But the truly amazing thing about the surreal numbers is how Conway found them: by playing and analysing games. Like an Escher tessellation of birds morphing into fish—focus on the white and you see the birds, focus on the red and you see fish—Conway beheld a game, such as Go, and saw that it embedded or contained something else entirely, the numbers. And when he found these numbers, he walked around in a white-hot daydream for weeks. (Roberts 2015)

The article quoted here is an edited selection from Siobhan Roberts's biography of Conway. Not only is his famous playfulness linked explicitly to

³And many more besides. See (Inglis & Aberdeen 2015) for experimental results on how the different descriptors may be linked.

the discovery of the surreal numbers, it is furthermore connected to his discovery of the Conway groups, where the Conway groups are three of the twenty-six sporadic cases of finite simple groups, discovered by looking at automorphisms of the Leech lattice. Playfulness as exhibited by Conway is seen as an epistemic virtue, in which he takes joy in the complex mathematical structuring to be found in games. By translating difficult mathematical problems and ideas into games, Conway managed to invoke his playfulness to be deeply fruitful for mathematical discovery. The virtue of playfulness is clearly relevant both to Conway's own knowledge and the collective knowledge of the mathematical community, since it was certainly instrumental in the discovery. The identification of the Conway groups also feeds in to the triumphal proof of the classification of finite simple groups, thereby demonstrating its connection to the rich tapestry of mathematical achievements.

That virtues are relevant, in an interesting sense, to mathematical knowledge is thus shown by the connection between Conway's playfulness and the mathematical discoveries it resulted in. The point could equally well have been made of numerous other mathematicians and their respective virtues which have contributed to mathematical breakthroughs and discoveries. As such, I take the Moderate Proposal to be vindicated by the example. The Moderate Proposal, nonetheless, is maybe best seen as a "foot in the door", in that it is readily acceptable but nonetheless brings to the foreground the importance of virtues in mathematics. By making it plain that our interests as philosophers of mathematics are bound up with the virtues displayed by mathematical theories, proofs, ideas, concepts, methods *and practitioners*, the Moderate Proposal is suggestive of the fact that it will be a fruitful project to further investigate the relationship between the two. This is my intention for the remainder of the chapter, beginning in the following section with a description of the main strands of thought in virtue epistemology.

5.3 Virtue Epistemology

The central idea that unites the diverse approaches under the banner of 'virtue epistemology' is that individual agents and groups of agents must be considered in approaching the core issues of epistemology. The way in which the knowers themselves figure in theories of knowledge (and other epistemological concepts) is through the exercise or failure of their epistemic

virtues and vices. Where these theories come apart is on the nature of these epistemic virtues. The two main camps are the virtue reliabilists and the virtue responsibilists, which I shall outline in turn, followed by a brief discussions of further ways in which one might be a virtue epistemologist, namely by taking on some kind of hybrid view or by focusing on epistemic agents while rejecting the reductive and definitional projects in epistemology.

5.3.1 Virtue Reliabilism

Virtue reliabilism is the approach proposed by Ernest Sosa (Sosa 1980, 1991, 2007), John Greco (Greco 2010), Christoph Kelp (Kelp 2011) and Alvin Goldman (Goldman 2000) which takes virtues to be stable, reliable faculties, abilities, skills or competences. The epistemic variety of virtues, the *intellectual* virtues, are then the faculties, skills or competences which are aimed at epistemic ends, such as acquiring true beliefs and avoiding false ones. Reliability is understood in terms of how well those ends are achieved. The faculties that reliably produce true beliefs in this way are quite broad, including “faculties of sense perception, memory, induction, and deduction” (Battaly 2008, p. 645).⁴

Given the picture of intellectual virtues as skills, abilities, faculties or competences which reliably maximise truth over falsity, we may now see the account of knowledge that this provides. In (Greco 2010), Greco argues that knowledge is a kind of success through ability, meaning that it amounts to an achievement rather than a merely lucky success (and thereby avoiding Gettier cases):

S knows that *p* if and only if *S* believes the truth (with respect to *p*) because *S*'s belief that *p* is produced by intellectual ability.
(Greco 2010, p. 71)

Similarly, Sosa defends the following claim:

[K]nowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence. (Sosa 1991, p. 277)

⁴Given the current focus on the mathematical case, the last item on the list, deduction, will be of particular interest to us. However, I believe that such a broad notion does not suit our needs in appealing to the virtue approach, so the account of mathematical deduction I am proposing will be more fine-grained.

This develops the intuitive idea that knowledge comes about through a process which isn't lucky, but also emphasises the cognitive role that is played by the knowers themselves.

The most famous way that Sosa explains his view is through the example of the archer and three evaluative measures of how well she does: accuracy, adroitness and aptness. *Accuracy* is an evaluation of the success of the shot at actually hitting the target. *Adroitness* is whether or not shooting the bow manifested the relevant skill at archery on the archer's part. However, the big observation is that a shot can be both accurate and adroit, but fail to be accurate-*because*-adroit, due to double luck situations analogous to Gettier cases where good luck cancels out bad luck. For example, the archer may line up the perfect shot, but a gust of wind blows the arrow astray, only for it to deflect off a tree and finally hit the target. *Aptness*, then, is Sosa's term for those performances which are successful because they manifest the correct skills or competences, and it is these which are creditable to a skilful and virtuous agent. The idea, then, is that obtaining knowledge requires the same manifestation of skills or competences, but furthermore, the success at obtaining true beliefs needs to be because of the skilful or competence performance. Sosa calls this the AAA-structure: knowledge requires being accurate, adroitness in manifesting the relevant skill or competence, and being apt.

The reliabilist virtues listed above are not the end of the picture, since many skills which reliably attain truths are acquired and developed through practice. Thus complex intellectual skills and abilities can play a key part in acquiring true beliefs, as we shall see in the mathematics case. This is closely connected to the difference between *low-grade* knowledge, the more immediate knowledge gained through channels such as sensory experience, and *high-grade* knowledge, the more reflective, systematic and inquiry-focused knowledge.

5.3.2 Virtue Responsibilism

Virtue responsibilism, as championed by Linda Zagzebski (Zagzebski 1996), Lorraine Code (Code 1987) and James Montmarquet (Montmarquet 1993), sees intellectual virtues as needing to be understood in a way that is broadly continuous with moral virtues in the virtue ethics tradition dating back to

Aristotle.⁵ On this view, virtues are acquired excellences or *traits of character*, examples of which include open-mindedness, intellectual courage, intellectual autonomy, intellectual humility, reliance on trustworthy authority, perseverance and thoroughness.

The intellectual virtues as character traits have two major components on Zagzebski's account.⁶ Firstly, they have a motivational component, which are dispositions to be motivated towards particular ends, something which wasn't present for the reliabilists. In the case of intellectual virtues, the motivation is generally epistemic and is described by Zagzebski as all being "forms of the motivation to have cognitive contact with reality" (Zagzebski 1996, p. 167), where this broad heading covers the desires for true beliefs, certainty, understanding etc. Secondly, virtues also have a success component, for Zagzebski argues that to be virtuous means to be reliably successful in securing the ends you are motivated towards. Putting these together, character traits which are made up of these two components are then enacted:

An act of intellectual virtue *A* is an act that arises from the motivational component of *A*, is something a person with virtue *A* would (probably) do in the circumstances, is successful in achieving the end of the *A* motivation, and is such that the agent acquires a true belief (cognitive contact with reality) through these features of the act. (Zagzebski 1996, p. 270)

Such acts of intellectual virtue are key in the definition of knowledge Zagzebski gives:

Knowledge is a state of cognitive contact with reality arising out of acts of intellectual virtue. (Zagzebski 1996, p. 270)

The fact that cognitive contact with reality here is broader than just having true beliefs is very relevant to the mathematical case we will get to later, for consider:

⁵It is pointed out in (Greco & Turri 2011) that this is better called Neo-Aristotelian rather than Aristotelian because Aristotle does not claim such a unified account of moral and intellectual virtues.

⁶While Code and Montmarquet have similar attitudes to what the virtues amount to, they are less committed to recovering a definition of knowledge in terms of the virtues. I'll return to this in the next section.

[U]nderstanding is also a form of contact with reality, one that has been considered a component of the knowing state in some periods of philosophical history... [I]t is a state that includes the comprehension of abstract structures of reality apart from the propositional. (Zagzebski 1996, p.167)

This will be useful in approaching the wider aspects of mathematics which are usually left aside in philosophical and epistemological accounts, such as mathematical understanding, visualisation and diagrams.⁷ We will return to understanding later in section 5.7.

5.3.3 Hybrid Virtue Approaches

Under the general heading of virtue epistemology we are not limited to the two approaches sketched above. In this section I will briefly set down a further distinct way of taking the virtue turn towards epistemology, by adopting a hybrid approach which combines aspects of the other two theories. The hybrid approach in virtue epistemology is endorsed by both Heather Battaly in (Battaly 2008) and Nenad Miscevic in (Miscevic 2007). The guiding idea is that both the virtue responsibilist and virtue reliabilist proposals have correct ideas which will take us forward in our epistemic theorising. Rather than seeing the two as rivals, the hybrid approach can partially endorse both, or take on aspects of both theories.

The first way to do this is by broadening the category of epistemic virtues to include both virtues as faculties and virtues as character traits, as is suggested by Battaly. Her reason is that both ways of filling out the concept of ‘virtue’ are equally legitimate, in that they both track normal uses of the term. The idea then is that the two kinds of virtue correspond to different kinds of knowledge. Contrasting the two categories of *high-grade knowledge* and *low-grade knowledge*, Battaly argues that these are achieved in different ways. Low-grade knowledge is the sort of quick and immediate knowledge we get from sensory experience, while high-grade knowledge is knowledge which requires greater cognitive effort or reflection, such as scientific knowledge. The hybrid view being proposed, then, can make use of the fact that for low-grade knowledge all that seems to be required is the correct and reliable functioning of the relevant faculties while high-grade knowledge seems

⁷Though excellent work in this area is done by (Giaquinto 2015).

to need character traits such as inquisitiveness, open-mindedness etc. Conversely, requiring responsibilist levels of motivation seems to fit poorly with the fact that we can come to know things by seeing them without any motivation towards knowing, while the reliabilists seem to leave an incomplete story if their account of higher-order knowledge fails to involve the relevant traits required by the process of inquiry. As such, the hybrid view can endorse different kinds of virtues as necessary for different kinds of knowledge.

The other way of endorsing a hybrid view, as presented by Miscevic, is to divide up the aims of the approach to allow the different types of virtue to satisfy those different aims. Miscevic's idea, entitled the *integrated virtue-based view*, is to endorse a virtue reliabilist account of knowledge (motivated by the idea that this does better at truth-tracking), while taking the value of knowledge to be tied to the character trait of inquisitiveness.

Indeed, Zagzebski's own virtue responsibilist view readily acknowledges that skills are importantly related to intellectual efforts, including mathematical ones:

Spatial reasoning skills, mathematical skills, and mechanical skills are important for effectiveness in many of life's roles, and the person who is virtuous in such roles would be ineffective without the associated skills. (Zagzebski 1996, p. 115)

Nowhere will this be clearer than in mathematics itself! There is certainly a close connection between intellectual virtues and intellectual skills, something which the hybrid views want to employ to develop a full account of both.

5.3.4 Epistemic Vices

Very little has been said so far about epistemic vices, echoing the literature where these have mostly played a secondary role, something which is only recently being rectified. Let us quickly survey how intellectual vices might be understood.

There is an Aristotelian line on the nature of moral vices, as is well-known, taking virtues to be intermediate between vices of excess and vices of deficiency:

Now [virtue] is a mean between two vices, that which depends on excess and that which depends on defect; and again it is a mean

because the vices respectively fall short of or exceed what is both right in passions or actions, while virtue both finds and chooses that which is intermediate. (Aristotle 2009, II.6.1107a.2–5)

Aristotle illustrates with various examples such as the virtue of proper pride being an intermediate between empty vanity and undue humility, or the virtue of courage being the mean between cowardice and rashness. However, there is no obvious reason to take this to be generally true, so while it might fit in particular cases this will not do as a characterisation of vice generally, nor of epistemic vice, nor of mathematical vice, without some further argument to that effect.⁸

Starting with the reliabilist approach, we do find some mentions of vices. For example, Goldman lists the following as intellectual vices:

The vices include intellectual processes like forming beliefs by guesswork, wishful thinking, and ignoring contrary evidence. (Goldman 2000, p. 6)

These appear to be intellectual processes which are actively misleading, taking us generally towards false beliefs rather than true ones. Similarly, for the virtue responsibilist approach, Zagzebski says:

Some examples of intellectual vices are as follows: intellectual pride, negligence, idleness, cowardice, conformity, carelessness, rigidity, prejudice, wishful thinking, closed-mindedness, insensitivity to detail, obtuseness, and lack of thoroughness. (Zagzebski 1996, p. 152)

The difference between the two quotes corresponds to the difference in their account of virtues, the reliabilist seeing vices as processes which are unreliable in delivering true beliefs and the responsibilist taking vices to be negative traits of character.

In (Battaly 2014), Heather Battaly points out that actually we can be more precise about the concepts of vice on offer. Firstly, reliabilist vices can be distinguished between those which deliver negative epistemic ends and

⁸Not to take the Aristotelian picture of virtue and vices too seriously as a modern framework was a point impressed on me by Brendan Larvor. Rightly so it seems, for even Zagzebski who is explicitly trying to apply Aristotle's picture of moral virtues to intellectual virtues does not take this to be correct.

those which fail to produce positive ones. Similarly, responsibilist vices can either be filled out in terms of having negative motivations or in terms of the failure to have positive motivations. The last of these four options may seem initially implausible, but Battaly argues that there are ways in which this can be more convincing. For example, we might be blameworthy of vice simply by failing to consider epistemic value, such as through dedicating ourselves to collecting trivia about soap operas; there is nothing wrong with soap trivia per se, so this cannot be a case of bad epistemic motivation, but the problem is the failure to pursue valuable epistemic ends.⁹

Epistemic vices, then, can be cashed out in different ways, according to one's other theoretical motivations. For the sake of this chapter, we can leave it open whether any one of these is the correct account of epistemic vice and continue now with an exploration of how virtue epistemology will apply to mathematical knowledge.

5.4 The Strong Virtue Proposal for Mathematical Knowledge

Above I argued for the correctness of the Moderate Proposal, where the investigation of theoretical or agential virtues is relevant to the philosophy of mathematics. Now I shall take on the Strong Proposal, which was the following:

Strong Proposal Virtues and vices of mathematicians will be explanatory of mathematical knowledge. In other words, virtue epistemology should be adopted to give the correct epistemology for mathematics.

Such a claim is far from obvious. Nonetheless, I shall propose that we should adopt a virtue-epistemological approach to mathematical knowledge and that this will be successful in settling difficult problems in the philosophy of mathematics. The virtue account of mathematical knowledge will be a genuine rival theory of mathematical knowledge which can draw on an established tradition in the realm of epistemology to solve problems pertaining

⁹An interesting conclusion of the analysis of vice is that Miranda Fricker's discussion of epistemic injustice (Fricker 2007) is best understood in terms of the bad outcomes it leads to, so aligns more naturally with the reliabilist conception of vice despite Fricker's responsibilist framework.

to mathematics, such as those concerning the nature of proof and mathematical rigour. In this section I will set out the virtue-epistemic approach to mathematical knowledge.

One place we might look for a prior articulation of the Strong Proposal would be in the literature on virtue epistemology. Indeed, we should expect any proponents of virtue epistemology to be immediately inclined to agree that their preferred epistemology extends to mathematical knowledge too. Yet there seems to be little previous consideration of how such a theory might go or how it might benefit mathematics. Reliabilists often list *deduction* as a virtue, in that it is a reliable faculty or competence that leads to knowledge, but it strikes me that this is far too coarse to deal with the subtle and complex issues going on in mathematics. Meanwhile, responsibilists have not had anything explicit to say about the mathematical case, although Zagzebski does quote Moravcsik relating mathematical proof to understanding at (Zagzebski 1996, p. 47).

Applying virtue epistemology to a hard case like mathematics should be an appealing undertaking for people already convinced of the correctness of virtue epistemology, not just as a straightforward application but also as an important test case. It might be that the mathematical applications of virtue epistemology will favour reliabilism over responsibilism or vice versa, or indeed a hybrid view such as Battaly's over both. Even more important is that mathematics has a number of peculiar epistemological difficulties, whereby we treat mathematical knowledge as *special*. Mathematical knowledge might be considered special in any number of ways such as it being a priori, necessary, deductive, objective, infallible, certain, analytic etc. Many of these are central to the Traditionalist take on mathematical knowledge and explaining them is seen as one of the main projects that we should be engaged in. The questions that might thus arise in applying virtue epistemology concern which of these properties of mathematical knowledge are maintained and defended, or whether we join Lakatos in storming the dogmatist fortress and reject some or all of these properties. For now I am just blazing a path for this work to be carried out so will not be answering all of these questions here, but it is important to note that there are such major questions which need to be developed in giving a theory of mathematical knowledge and that filling out these details for the virtue approach is a project which needs carrying out.

The first major question for the virtue project is: what exactly are the virtues at play in mathematics? Setting aside the theoretical virtues discussed above, in virtue epistemology the virtues are primarily those possessed by agents (or possibly groups of agents). The answer to what they are exactly will depend on the flavour of virtue epistemology one favours so let me take these in turn.

On the virtue reliabilist account, virtues are taken to be faculties, competences, skills or abilities, varying amongst the different proponents of the view. For the current case, then, mathematical virtues will be mathematical faculties, competences, skills or abilities. It strikes me that these may be separable into two levels. Firstly, mathematical skills or abilities may be about being able to implement particular mathematical techniques, at what we can call the *particular* level. These might range from the most basic mathematical skills such as counting, mental arithmetic or finding the roots of a quadratic equation, all the way to advanced mathematical techniques such as forcing in set theory, stochastic modelling or finding saddle points in dynamical systems. With the techniques at the particular level, there is a notable granularity issue in how precisely we define some given technique and when we attribute such a skill to a mathematician. For example, mental arithmetic above could have been broken down into skill at adding, subtracting, multiplying, etc. where we could count these as distinct skills. I don't believe anything major hangs on how exactly we specify such skills, but it does seem to indicate that certain contextual parameters will be in play when employing mathematical virtues for epistemological purposes. This would certainly fit with the discussion of Löwe & Müller from the previous chapter. This is also in line with how things go in virtue reliabilism more generally.

However, it seems that just deploying particular skills is not sufficient for many mathematical activities, such as developing and checking proofs. As a second level of mathematical virtues for the reliabilist, what we may call the *general* level, we might identify reflective, higher-order mathematical skills. For example, giving a proof isn't just about deploying particular mathematical techniques but also about picking the right method for the situation and combining it with other techniques in the correct manner. The higher-level virtues seem to be those we intend when we talk of 'mathematical thinking' generally, which includes being able to construct and follow mathematical

arguments, accurately check for errors and solve technical problems on the fly.

Let us proceed to the virtue responsibilist take on mathematical virtues. Mathematical virtues, here, would form a subcategory of intellectual virtues, which are defined as acquired character traits. There are many such intellectual virtues that are already present in the general case but seem to apply directly to the realm of mathematical activities. For example, Zagzebski lists thoroughness, perseverance, “the teaching virtues: the social virtues of being communicative, including intellectual candor and knowing your audience and how they respond” (Zagzebski 1996, p. 114) and reliance on trustworthy authority as intellectual virtues. Thoroughness and perseverance are clearly important for discovering and developing proofs; communication is not limited to mathematical teaching but also relevant to tailoring proofs to their intended audience; and relying on testimony from trustworthy sources will be vital to engaging with the mathematical community and developing collaborative mathematics. On the other hand, there may be intellectual virtues which are particular to mathematical endeavours. Indeed, I will later defend the idea that, for the virtue responsibilist take on mathematics, *mathematical rigour* is a specifically mathematical virtue.

Now the big idea behind the strong virtue proposal is that mathematical knowledge should be explained in terms of the possession and enacting of certain mathematical virtues, whichever account of them one prefers. Notably, though, an important aspect of the move towards the evaluation of mathematicians’ virtues as key to the notion of mathematical knowledge, is that this is intended to move away from a misguided, overly-idealised standard account towards a more accurate picture of mathematics and mathematical knowledge in the real world. As such, I am careful to point out that actual mathematical knowledge can be obtained in a whole range of different ways. The major traditional way is, of course, through deductive proofs, which will be the focus of the next section.

But there are many others too. Testimonial knowledge has a huge literature in epistemology (see Adler 2012), and can and does provide a lot of mathematical knowledge out there too. In mathematics education, there is a great deal of reliance on the teacher’s word, at least initially. One might think that once we get into the domain of mathematics proper, maybe university and beyond, that we should be familiar with the proofs behind all

of our mathematical knowledge, but here we might be relying on our memory to provide them; that having once seen/produced/understood a proof the knowledge remains. The virtue reliabilists actually categorise memory as an intellectual virtue, so admitting its use is already requiring greater philosophical resources. Even still, the idea that we have even encountered all of the proofs for theorems we can claim to know, rather than relying on testimony, might not even be so plausible. Indeed, it seems to be common knowledge that Fermat's last theorem and Poincarè's Conjecture are both true, yet undoubtedly most mathematicians have only a rough knowledge at best of how proofs of these proceed. Furthermore, mathematicians are social builders, disseminating and using results produced by others to make progress in their own work. While we might idealise that everyone will trace back all the results they use through the tree of dependencies of results those in turn rely on, it seems wildly unlikely. Consider further the classification of finite simple groups, a proof of such grand scale that likely nobody knows all of it. Such massive collaborative projects are grand successes of mathematics and undoubtedly involve testimonial knowledge.

We have only scratched the surface of the epistemic complexities involved in mathematical knowledge in practice. To list some more possible ways to gain mathematical knowledge, consider mathematical knowledge through sensory experience (Giaquinto 2015); through probabilistic justification, like cases of primality testing (which may be problematic in their resemblance to lottery cases); through computational verification, from mere calculators to complex models run on super-computers; or mathematical know-how which might be acquired through training and being inducted into particular practices (as discussed in the previous chapter). A natural response to the diversity of sources of mathematical knowledge is to concentrate on some particular special properties of mathematical knowledge which are accorded only to knowledge gained deductively from proofs and thereby isolate the interesting case. In the next section I take on the strong proposal's account of mathematical proofs, but I want to note that the virtue turn thrives on the diversity of sources of knowledge because there are a great number of intellectual virtues which are important for mathematics. The benefit, then, is that applying virtue epistemology to mathematical knowledge does not find itself limited to pure mathematics and deductively proven theorems, but conversely provides a framework for giving a fuller, richer account of

the whole spectrum of mathematical knowledge. A move away from taking all mathematical knowledge to be a priori is advocated in (Kitcher 1984) and the virtue approach allows us to do this while still adopting a unified account of mathematical knowledge. This is already a huge advantage of the virtue approach to mathematical knowledge being endorse on the Strong Proposal. Furthermore, in opening up this diversity of types of mathematical knowledge, we see the materialisation of the fact that deploying the virtue epistemological approach to mathematical knowledge is more than just the flat-footed project of answering Traditionalist questions with a new set of answers, but rather opens us up to the wider project of investigating diverse mathematical practices and the interrelated virtues and vices at play in them.

I will now proceed to a discussion of mathematical knowledge from proofs on the virtue account.

5.5 Virtues and Proving

The Strong Proposal to see mathematical knowledge as best explained through a virtue epistemological lens can provide a better alternative to the Formalist-Reductionist account of proofs. Rather than being hindered and undermined by the issues we have seen raised against the Formalist-Reductionist account throughout the thesis, the virtue epistemological view on proofs I will offer positively thrives on them, as we shall soon see.

The first move that we make on the virtue approach needs us to observe that in the virtue epistemology literature the key to knowledge is *virtuous intellectual activity*. It is through virtuous acts, acts in which we exhibit or manifest relevant virtues, that we gain knowledge. The Formalist-Reductionist approach, meanwhile, emphasises proofs as construed as objects, where we can then study these objects to discover what makes them good/correct/rigorous/etc. But then the account fares poorly with respect to our proving practices, having tried to abstract away from their material instantiation. In contrast to the proofs-as-objects view, then, the virtue account should focus instead on *proving as an activity*, as argued for in the previous chapter.

Let us now proceed to locate the role of virtues in this account of proofs and the knowledge it secures. The virtue epistemological view is that in

order for the proving activities to secure our mathematical knowledge, they need to be virtuous in the appropriate way. Once again, how the details of this are filled out will depend on the particular virtue epistemology one favours.

Starting with the reliabilist, for someone producing a proof to have secured mathematical knowledge therewith, the proving must be completed with the relevant skills and competences. The mathematician needs to correctly deploy the particular mathematical skills involved in whatever the proof is, rightly observing any limitations or restrictions on the domains of application. Furthermore, these skills need to be tied together to form a coherent whole, which delivers the final theorem as the result of the manoeuvres combined correctly, that is, avoiding errors and mistakes, which will be the result of general level competences. Similarly, in the case of checking or learning from a previous proof, the proof on the page (or wherever it may be) acts as a guide or recipe as to how one should carry out the actions of the proof, as described in chapter 4. Still, the person doing the checking must accurately follow the techniques being presented and see how one step follows the last in order to come to know the ultimate solution. Importantly, following steps in this sense does not need to be filled out as following the underlying formal moves, but instead is about the fitting together of the steps in the overall reasoning pattern and recognising what follows from what in the moves that are being made, moves which can certainly be informal in the operative sense.

Virtue responsibilists, on the other hand, require that the mathematician comes to know the proved theorem through acts of mathematical virtue, which is to say that their proving activities must be virtuous and free from vices. In particular, they gain knowledge through proving if the activity of this instantiates the necessary virtue of mathematical rigour. The nature of rigour is a major question of philosophy of mathematics, one which has not been done justice by the Formalist-Reductionist answer, so I will discuss this in greater depth in the following section as something additional that the Strong Virtue Proposal can offer besides the expected account of mathematical knowledge. Importantly, following the Zagzebski framework, the virtue of rigour will have an epistemic motivational component, usually aimed at establishing the truth of the theorem, or making cognitive contact with mathematical reality (whatever form that takes), and also a success

component of actually doing so. Once again, this builds in the fact that our right reasoning must track the contours and landscape of mathematics.

Let us illustrate these both using the Stallings example we began the chapter with, where these virtues go astray to lead to a failed proof. Recall that Stallings's proof of the Poincaré Conjecture failed because it relies on deploying a theorem in a case for which it doesn't hold. The mathematical reliabilist account of what has gone on here is that what Stallings describes as his "psychological problem, a blindness, an excitement, an inhibition of reasoning by an underlying fear of being wrong" (Stallings 1966, p. 88) represents a failure of his usual skills and competence at putting together a proof to form a complete argument, check for errors and, in particular, observe the domain of application of theorems being used. The talk of developing "[t]echniques leading to the abandonment of such inhibitions" (Stallings 1966, p. 88) can be taken seriously; the development of such skills is paramount in securing correct proofs and further mathematical knowledge. For the mathematical responsibilist, the focus will be on the misapplication of theorem 0 as a failure on his part to be rigorous in his proving, rigour being a mathematical character trait to be discussed shortly. Additionally, the responsibilist would also be more interested in the idea that developing the techniques for avoiding this is something desirable to "every honest mathematician" (Stallings 1966, p. 88) as the intellectual honesty will work alongside rigour to ensure that our motivations in our mathematical proving are of the correct sort, rather than simply being directed at fame and fortune.¹⁰ For both approaches, the failure to enact the relevant virtues in the creation of the proof are how the proof came to be mistaken and its reasoning flawed, with the upshot that the activities did not deliver mathematical knowledge to Stallings of the truth of the Poincaré Conjecture.

What has been given here is a description of how proving as an activity can deliver or fail to deliver mathematical knowledge. However, the challenge on which I began was to account for proofs, rather than the activities surrounding creating and verifying them. In the previous chapter, I put forward a re-orientation of priority, where we focus on proving activities first and see proofs as objects or arguments as of secondary concern. *Proofs*

¹⁰Of course, this does not mean that this cannot be a part of your motivation, just that it cannot come to the detriment of correctly tracking mathematical truth.

don't operate in a vacuum was a slogan form of the idea that a proof considered in some idealised sense will miss out on the crucial connection to the provers who create, know, understand and employ them, a connection I have claimed is fundamental to the mathematical knowledge that proofs are aimed at. Nevertheless, there is an emerging worry here which must be addressed: that prioritising the individual activities appears to take us too far in an individualistic direction. The fact is that we do have a seemingly robust notion of what is sufficient for a proof, a set of standards taught in classrooms and lecture halls which is enforced by teachers and even the referees for journals. One might think that it cannot be simply down to the individual whether a proof is enough to secure knowledge, as this seems hopelessly subjective with respect to mathematics, which should be held up as objective.

In response, though, we can observe that this does not really cut to the core of the issue and that this instead merely reintroduces certain Traditionalist attitudes to mathematical proof. For one thing, it seems that this response seeks to re-idealise proofs as something which are 'out there' to discover, thereby conflating the contours, structures and relationships of mathematics, on one hand, with the proofs themselves on the other. Crucially, proponents of the Strong Virtue Proposal can and should accept that there are many operative canons in mathematical practice and that proof is standardised in numerous ways. None of this poses a difficulty, however. The point is that the standards are set for how to best structure and communicate proofs, thereby also helping to inculcate mathematical reasoning and problem-solving into students by demonstrating how to set out reasoning in clear and cogent ways. That each individual needs to go about actually carrying out the reasoning in order to gain the specific type of knowledge associated with proving (as opposed, say, to testimonial knowledge of its correctness) does not necessarily impede the objectivity of the mathematics at stake. To return to the map metaphor: we can agree that there are important map-making conventions, while asserting that properly knowing the route it describes involves traversing it.

Under the current attitude, we can also do better than the Traditionalist in discussing the conventions and standards surrounding proving. For example, the claim that there are well-guarded standards of mathematical proofs must come with some major qualifications. The fact is that the demands

such conventions place on us vary from context to context, with more details demanded for proofs for students and less for discussions between colleagues etc. These differences seem to concern the granularity of the proof, or how much can be assumed on the part of the reader. For instance, in Pettigrew's review (Pettigrew 2016) of (Burgess 2015) he points out that informal proofs must communicate the key ideas which deliver the truth of the theorem, and that it is this rather than the mere convincing which is important for a proof to be successful. For the Traditionalist, such a statement would fit poorly with the constant standards of rigour which are assumed, but for the Virtue approach this is only natural since virtues can be assessed across contexts and situations. Besides the virtue of rigour which can be displayed in proving, many other virtues are relevant to mathematical work, including virtues (in whatever sense they are taken) which pertain to communicating and collaborating. It may well be that very coarse proofs which only include the main ideas might not fully display whether the thinking underlying them is rigorous or not, in that they leave substantial gaps in between these main ideas, but could still be sufficient to communicate how to go about rigorously proving something if we address them to the right people. The point is that a virtue approach embraces also the diversity of purposes for which we employ proofs, as set out at the start of the introduction to this thesis. While one such purpose is to fully set out how to deduce some conclusion, another may be merely to communicate how this is done to a fellow researcher who does not need the full explanation to arrive at mathematical knowledge, by virtue of their existing knowledge and abilities. The virtue picture does give us a way to account for how these different facets come together and what to say when they come apart, in particular that divergent purposes might need to be assessed with respect to different virtues.

Ultimately, I take the point here to be that communal standards reflect something about the way in which we systematise our communication of mathematical ideas, approaches and proofs. This does diverge from the question of how proofs and proving relate to mathematical knowledge, but has been constantly conflated in Traditionalist and Formalist-Reductionist approaches. Let us now proceed to see what can be said concerning mathematical rigour from the virtue perspective.

5.6 Rigour

Mathematical rigour has already played a role in the discussion above, where it has been proposed that we should see it as a specifically mathematical virtue. Let us investigate this in greater detail.

Again, there is a strong claim put forth on the nature of rigour given by the Formalist-Reductionist, taking rigour to amount to little more than formality or straightforward formalisability. The ‘straightforward’ component here is often cashed out as algorithmic formalisability, a process of filling in the gaps until we hit a bedrock of formal inferences. That the formalisation process is nowhere near so straightforward has been frequently pointed out, such as in (Antonutti Marfori 2010, p. 266) and back in chapter 1, and this cannot be seriously maintained by anyone with experience of actually formalising proofs. In Marfori’s view:

[A]s a matter of fact it cannot be denied that in ordinary mathematical practice, standards of rigour are constantly appealed to. However, these very much differ from standards of formal rigour. Formalisation is seldom called for, and the mathematical community seems to widely converge on what to count as an adequate proof for the truth of a theorem and adjudicates controversies often without the need to formalise informal arguments. (Antonutti Marfori 2010, pp. 270–271)

She argues that we should thus separate our investigations of formal rigour from a separate endeavour of investigating informal rigour, amounting to the community standards converged on as securing the right degree of mathematical certainty. However, even these communal standards seem insufficient to me for the reasons outlined at the end of the previous section, namely their variability and shifting with respect to the purposes, creator and audience of the proof. I claim, though, that investigating these will not solve the issue concerning the relationship between informal proofs and mathematical knowledge, due to still not appreciating the role of the knower in the proving which has been emphasised throughout this chapter as a way of casting off the problems facing the Traditionalist.

The proposal is that, in addition to the notions of formal correctness and the ‘informal rigour’ of communal standards, we also investigate the notion

of *rigour as a virtue*, analogous to intellectual thoroughness or meticulousness but concerned specifically with mathematics. In this way, mathematical rigour will have a clear connection to mathematical knowledge through proofs, as outlined above, as well as explaining how we can do rigorous mathematics outside of proofs, namely by manifesting the virtue in other mathematical activities. Let us examine the details of how this works for the virtue approaches to mathematical knowledge.

Let us begin this time with the responsibilist, as the story here is more straightforward. We can take rigour to be a mathematical virtue, that is, rigour is an acquired character trait and excellence specific to mathematical practices.¹¹ As with other intellectual virtues, rigour will have a motivational component and a success component. The motivation is, in the terminology of Zagzebski, to make cognitive contact with mathematical reality, where this involves one's proving and other mathematical work tracking the relationships, contours and dependencies of mathematics, while avoiding errors, substantial gaps and wrong turns. Success, of course, is success at doing so, that your proving is in actual fact error-free, fallacy-free and gap-free. Rigour is acquired in the sense that we need to train and habituate ourselves to be rigorous mathematicians, through schooling, university and ongoing practice, as discussed in the previous chapter in relation to Rylean know-how. This is done both through learning the actual mathematical facts and techniques themselves, as well as through constant feedback cycles of what needs to be reasoned out versus what counts as a gap in reasoning, with respect to a practical context we find ourselves in. We try to weed out related vices of sloppiness, guesswork and unrigorous thinking in order to avoid gaps, errors and leaps of reasoning.

Now there are two major worries that can be raised against the division of rigour into the three components: formal rigour of derivations in formal systems; informal rigour of communal conventions, standards and norms; and rigour as a virtue (as I have introduced it). The two related problems for the inclusion of rigour as a virtue are the following. Firstly, usual talk of rigour does not seem to match up with it being a character virtue of the sort described here, but rather is applied to proofs themselves. Secondly,

¹¹We might want to take the virtue here to be 'rigorousness' instead of simply 'rigour', to distinguish the character trait from the communal standards of rigour, but the point is that these are not sufficiently robust to ground proofs, while I argue that the character trait is and so should be taken as primary.

the virtue as described above for the responsibilist directly relies on notions of correctness, gaps and errors, for which one can argue we need to appeal to the existing notions of rigour, making this third aspect redundant.

In response, it would certainly be wrong to deny that we do talk of proofs themselves being rigorous or unrigorous. However, this does not rule out that the character virtue is at play in such assessments. For example, we also speak of proofs as being ‘creative’, ‘inventive’ or ‘ingenious’, all of which are best understood as assessing the mathematician who authored the proof in so doing. After all, to take a proof in the abstract to be creative would be a category mistake, with the real meaning being that it displays mathematical creativity. Even closer to the virtue of rigour, we can also call proofs ‘meticulous’ without raising a fuss and, again, this seems to be better read as saying that the proof shows the meticulousness of its author. We should take the virtue of rigour to be of the same type as this, except with a vitally important role to play in the connection between proofs, proving and mathematical knowledge.

On the second objection, the point isn’t that the proponent of the virtue of rigour needs to deny that there are a number of conventions and social norms as to how we do mathematics and assess its correctness, or decide whether something is an error or a substantial gap; certainly these do exist. The point is just that these require an extra step to secure mathematical knowledge, and that rigour as a virtue of mathematicians (or groups of mathematicians) can make this step, thus is far from redundant. The way I have laid out the proof as a record and guide to activity versus proving as the activity itself, makes clear that the conventions of informal rigour apply foremost to the proofs themselves while virtues will be more concerned with the dynamics of the activity of proving. As such there are clear and definite theoretical roles to be played for both notions.

Let us turn now to the virtue reliabilist account of rigour. To achieve mathematical knowledge through a proof, one must deploy particular mathematical skills. Yet, for the reliabilist most skills require being reliable, but in the case of reasoning we want the skills of mathematics to not lead us astray and not permit fallacies, gaps and errors. Sosa discusses fallacious reasoning in (Sosa 2007), with the example of affirming the consequent:

[...] what denies justification to the fallacious reasoner might just be his carelessness or inattention or blundering haste. With

rare exceptions, normal, rational humans do not affirm the consequent when they are careful, attentive and deliberate enough. [...] Fallacies can thus be viewed as performance errors chargeable against the subject, by contrast with deliverances of a competence. (Sosa 2007, p. 59)

The idea seems to be that rigour, as in being careful, attentive and deliberate enough, is not some special competence which needs to be added at the general level, but rather is already built into the reliability of the specific skills being deployed in our reasoning. Rigour is present when the reasoning is carried out correctly and not present when it goes awry. As such, by Sosa's way of speaking, rigour would not be a virtue itself, but rather a necessary feature of certain mathematical virtues. The communal standards of mathematics will set out criteria for correctness of moves in a proof, and then a mathematician will have been rigorous just in case the deployment of the particular moves are free from errors and gaps.

In summary, the responsibilist and reliabilist accounts might very well come apart on the issue of whether rigour is a mathematical virtue, or instead a condition on mathematical virtues. Either way, though, rigour has an analysis in the virtue approach which does not rely on unattractive appeals to formality and formal proofs, contra the Formalist-Reductionist view.

5.7 Mathematical Understanding

Let us turn now to a third issue for which adopting a virtue-theoretic philosophy of mathematics, in particular the Strong Proposal, can provide new answers and interesting insights: the nature of *mathematical understanding*.¹² While still being a major concept of epistemology, understanding seems to come apart from knowledge in numerous ways:

1. Understanding can come in degrees whereas knowledge does not.
2. Understanding is compatible with epistemic luck while knowledge is not.

¹²One thing to be careful about is the fact that we can be said to understand many different things, which might not be uniform in the properties that understanding amounts to. For example, understanding German might be conceptually different from understanding group theory, or indeed understanding a particular proof.

3. Understanding may be non-factive, while knowledge entails truth.

The first and second claims are defended in (Kvanvig 2003, Ch. 8). The third is the point made in (Elgin 2006) by focusing on the case of scientific theories:¹³

[S]cience is riddled with symbols that neither do nor purport to directly mirror the phenomena they concern. Purified, contrived lab specimens, extreme experimental situations, simplified models, and highly counterfactual thought experiments contribute to a scientific understanding of the way the world is. (Elgin 2006, p. 213)

Considering the case of mathematics and mathematical understanding, one can go either way on this. On one hand, one might think that there cannot be the same disconnect between mathematical theorising and mathematical reality, such that only true mathematical theories are the sort of things one can understand.¹⁴ Indeed, this might fit well with a narrower view of mathematics which focuses on pure mathematics and understanding proofs. On the other hand, there are plenty of mathematical areas which permit similar arguments to Elgin's. I already touched on this while discussing Rav's *Pythagora* thought-experiment, but mechanics, applied mathematics generally and statistics all seem to make frequent use of modelling which do not even purport to be literally true of the target phenomenon, but nonetheless contribute to understanding. Given the broader conception of mathematics used in the latter view, I see no reason to be unnecessarily restrictive in taking understanding to be factive in mathematics generally.

So what is understanding? This is the big question in giving an account of understanding and therefore fairly important if we want to give an account of the more specific phenomenon of mathematical understanding. Speaking loosely to get started, understanding seems to involve recognising patterns, structures, relations etc. As Riggs puts it:

[U]nderstanding [...] is the appreciation or grasp of order, pattern, and how things 'hang together'. Understanding has a

¹³Worthy of note is that Kvanvig disagrees on this last item, seeing understanding to be factive. I take Elgin's point to establish that this is not correct.

¹⁴There are, of course, questions concerning what it is for a mathematical theory to be true. I shall not engage with these deep and complex issues here.

multitude of appropriate objects, among them complicated machines, people, subject disciplines, mathematical proofs, and so on. Understanding something like this requires a deep appreciation, grasp, or awareness of how its parts fit together, what role each one plays in the context of the whole, and of the role it plays in the larger scheme of things. (Riggs 2003, p. 217)

While not being very exact, this certainly points to the kind of phenomenon we are interested in. However, seeing understanding in this light adds two further differences from basic accounts knowledge, traditionally conceived. First, unlike knowledge, understanding does not obviously consist merely in holding certain beliefs, whereas knowledge is normally taken to be justified true beliefs plus whatever extra criterion one's theory prefers to avoid Gettier cases. Second, when dealing with the standard cases of knowledge it is assumed to be propositional, whereas understanding might well not be. For instance, understanding might well involve diagrammatic representations which are not reducible to propositional content, something which is particularly relevant for the case of mathematics where diagrammatic and visual thinking are not uncommon (see Giaquinto 2015).

Mathematical understanding itself has not received a great deal of attention in the philosophical literature in comparison to mathematical knowledge. Nonetheless, it does seem to be central in mathematics in practice. For example, in (Martin 2015), Martin draws her title from the following description by Andrew Wiles (famed for his proof of "Fermat's Last Theorem", better called 'Fermat's Conjecture' and now 'Wiles's Theorem') of the mathematical process of coming to understand new areas of mathematics:

Perhaps I can best describe my experience of doing mathematics in terms of a journey through a dark unexplored mansion. You enter the first room of the mansion and it's completely dark. You stumble around bumping into the furniture, but gradually you learn where each piece of furniture is. Finally, after six months or so, you find the light switch, you turn it on, and suddenly it's all illuminated. You can see exactly where you were. Then you move into the next room and spend another six months in the dark. (Martin 2015, p. 30)¹⁵

¹⁵Note the close similarity between this metaphor and the description of proving as

Let us thus investigate how the two main strands of virtue epistemology will be able to incorporate a theory of mathematical understanding in their accounts of mathematics. To do so I will take as a starting point the main account of specifically mathematical understanding in the literature, that of Jeremy Avigad, as found in (Avigad 2008). Interestingly, the two strands of virtue epistemology will come apart in whether or not they can take on the position Avigad endorses, as I will show after a description of Avigad's account.

Avigad draws heavily on the later Wittgenstein to present an account of mathematical understanding which is functionalist, in that he rejects any spooky theorising about understanding as pertaining to our inner mental lives, in favour of identifying understanding with the possession of particular *mathematical abilities*. Mathematical understanding is also, according to Avigad, closely connected to the ascriptions people make of mathematical understanding, which coheres with the main thesis because he argues that people ascribe understanding to each other just in virtue of their abilities to perform a selection of relevant mathematical activities. As such:

[W]hen we talk informally about understanding, we are invariably talking about the ability, or a capacity, to do something. It may be the ability to solve a problem, or to choose an appropriate strategy; the ability to discover a proof; the ability to discern a fruitful definition from alternatives; the ability to apply a concept efficaciously; and so on. When we say that someone understands we simply mean that they possess the relevant abilities. (Avigad 2008, p. 321)

As can be seen in the quote, Avigad readily moves between talking about understanding and talking about ascriptions of understanding. The picture is complicated somewhat by the fact that understanding is ascribed on the basis of a whole selection of interrelated abilities: for understanding a proof, Avigad lists eleven different abilities, any combination of which someone might have in mind when ascribing understanding—and the list isn't even meant to be exhaustive. The point, of course, is that understanding comes in degrees and so won't have some exact formula. Rather, the Avigadian account will be that understanding amounts to a cluster of abilities:

journeying through the mathematical landscape given in the previous chapter.

The claim I am making here is simply that the terrain we are describing is best viewed as a network of abilities, or mechanisms and capacities for thought. (Avigad 2008, p. 326)

Beginning with the virtue reliabilist version of the Strong Proposal, there is a great deal that fits well with Avigad's account of understanding. In particular, the virtue reliabilists take the virtues to be stable, reliable skills, faculties, competences or *abilities*. Putting the theory in terms of abilities, as noted above, is exemplified by Greco in (Greco 2010), taking knowledge to be true belief caused by intellectual abilities. As such there is a clear parallel between Greco's account of knowledge and Avigad's account of understanding.¹⁶ Thus, it is an available and attractive option for the virtue reliabilists to adopt an Avigadian line on understanding, since this would draw on the same source in abilities to explain both mathematical knowledge and understanding while simultaneously explaining their distinct features. Mathematical knowledge consists in the true beliefs which result from intellectual abilities; mathematical understanding consists in the possession of a broader but related network of abilities. Focusing on proofs once again, knowledge is obtained by using one's mathematical abilities to present reasoning from premises to a conclusion, while understanding comes in degrees relative to how many additional surrounding abilities the mathematician possess: abilities to explain, generalise, re-formulate, formalise etc.

Turning to virtue responsibilism, Zagzebski has a great deal to say about understanding that will be relevant to mathematical understanding. Indeed, she quotes Moravcsik on what it is to understand a proof:

What is it to understand a proof? It cannot be merely being able to reproduce it, or to know what it is, or to know lots of truths about it. [...] What elevates the above to understanding is the possessing of the right concepts, and the intuitive insight of the connection that makes the parts of the proof to be the proper parts of a sequence. (Moravcsik 1979, p. 55). Taken from (Zagzebski 1996, p. 47).

Zagzebski uses this to point towards an agreement with Kvanvig (and the line we saw in Riggs) that understanding amounts to grasping structures

¹⁶Greco does mention understanding, following the Kvanvig line that understanding amounts to a knowledge of causes, broadly construed.

and patterns of “a whole chunk of reality” (Zagzebski 1996, p. 46). A main point she develops also in (Zagzebski 2001) is that understanding involves the grasp of structures that are not necessarily propositional. As mentioned previously in section 5.3.2, understanding for Zagzebski sits beside knowledge as making cognitive contact with reality, thereby expanding the domain of epistemology on the virtue account to include such contact with non-propositional structures too. This certainly seems desirable in the case of mathematics, to account not only for diagrams but also the fact that understanding does seem to involve the more holistic features of seeing how mathematical theories, concepts, proofs, definitions and methods fit inside large structures.

As to the Avigadian picture, the Zagzebskian view will not fit anywhere near so well as the reliabilist could. While certainly the responsibilist view Zagzebski defends will hold that understanding entails certain abilities, it will not endorse the functionalist leanings which have Avigad rejecting talk of the inner mental life. Instead, the account is one on which we see understanding as coming through intellectual character virtues making cognitive contact with mathematical reality. For example, the mathematician might understand a proof by possessing insight into the whole structure and how it fits together as described by Moravcsik above. Hereby the responsibilist can also reject the conflation between understanding and ascriptions of understanding, such that the ascriptions do not take on the same importance they did in the Avigadian picture.

5.8 Case Study: Mochizuki and the *abc* Conjecture

In this section I will give a case study one of the most fascinating current episodes in mathematics, that of Shinichi Mochizuki and his controversial proof of the *abc* conjecture. In it we will see the explicit presence of virtue-theoretic terminology in mathematical practice, thus enriching and exemplifying the account given above. I shall begin by setting out some of the background and details of the case, after which I will draw on Mochizuki’s reflections on the status of his proof in order to show that the virtue-theoretical approach is also the most effective explanation of the controversy surrounding it. Indeed, this case has not finished playing out: there is at the time

of writing no settled consensus of whether the proof will be accepted by the community or not. However, I don't think this matters to the point I am making with it: whether or not the proof is accepted as correct or some irreparable error or gap is found, the clearest and best explanation of what is going on in the controversy is the virtue-theoretic one.

Some background: the *abc* conjecture, otherwise known as the Oesterlé-Masser conjecture, can be stated as follows:

***abc* conjecture** For every $\epsilon > 0$, there are only finitely many triples (a, b, c) of coprime positive integers where $a + b = c$, such that $c > d^{1+\epsilon}$, where d denotes the product of the distinct prime factors of $a * b * c$.

So, to give an example, try the triple $(5, 8, 13)$ which are coprime positive integers and form the sum $5 + 8 = 13$. The distinct prime factors are then 2, 5 and 13, so then $d = 2 \times 5 \times 13 = 130$. So for this choice of c and d , $c < d$ and thus for all $\epsilon > 0$ we have that $c < d^{1+\epsilon}$, meaning this is not going to be one of the finitely many exceptions.¹⁷ The general mathematical interest of the conjecture lies in the huge number of consequences it has in establishing other theorems.

At the end of August 2012, Shinichi Mochizuki uploaded four papers containing a decade's worth of his solitary work developing ideas on elliptic curves in what he calls *Inter-Universal Teichmüller Theory*, known as IU Teich for short. In particular, one result this leads to is a proof of the *abc* conjecture. The problem is, however, that a huge amount of material is covered in the run-up to the proof—fifteen hundred to two thousand pages of dense technical work, or down to five hundred if one is already an expert in anabelian geometry—and in it Mochizuki has developed a whole new area of mathematics with its own terminology, structuring, deep ideas, novel tools and original ways of thinking. As such, the proof has still not been widely accepted because nobody has been able to independently penetrate and understand the proof.

After an initial wave of excitement, both within mathematics and from the media such as in (Ball 2012), the daunting task of reviewing and verifying the proof became clear. While Mochizuki is described as forthcoming in responding to emails from other mathematicians, he has as of yet not lectured on his work outside of Japan, which has created quite a large barrier

¹⁷If we were trying to find an example of one of these exceptions, choosing our c as prime would've been a pretty poor starting point.

to acceptance by the wider mathematical community. As a result, there has been some frustration towards the proof and aimed at Mochizuki himself, such as in (O’Neil 2012) and with further examples given in (Castelvecchi 2015). At the end of 2015, a workshop was held in Oxford to try to get a larger group familiar with the central themes of Mochizuki’s work and the general structuring of the proof of the *abc* conjecture, with some success and other fresh expressions of frustration, such as are given in (Knudson 2015) and described in detail in (Conrad 2015). The nature of the surrounding controversy, then, concerns the correctness of the proof and the respective duties in the verification process of both the originator of the proof and the community at large who have to carry out this verification.

One could say a great deal about this situation following the template set out for the Moderate Virtue Proposal above, focusing on the virtues and vices both of individuals and the mathematical community at large that have lead to the difficulties in communicating and verifying the proof. For example, we could look at the balance required between intellectual autonomy in developing new mathematics and the collaborativity in bringing others along with you. However, the virtue proposal can be of even greater benefit in examining Mochizuki’s own reflections on the status of the proof and its verification. In response to the widespread discussion about the difficulties in verifying the *abc* proof, Mochizuki has produced two reports on his personal endeavours in trying to disseminate and communicate the proof (Mochizuki 2013, 2014). What is of particular interest to us is that Mochizuki’s opinions on what it will take for others to verify his proof appeal directly to the characters and motivations of those doing the verifying. It is not so common or straightforward to find mathematicians’ opinions and reflections on their practices without a pre-established philosophical agenda.¹⁸ As such, the fact that on this occasion—where a significant mathematical break-through is at stake—the language and ideas are unmistakably virtue-theoretical seems to be worthy of note, and furthermore offers the opportunity to demonstrate the effectiveness of a virtue-theoretic philosophy of mathematics.

Let us examine what Mochizuki has to say in some detail. The central contrast Mochizuki is trying to draw is between his collaborators, who he has worked through the proof with and now seem to have agreed on the

¹⁸We might be less concerned here about Mochizuki’s obvious mathematical and personal agenda, since this does not seem to automatically introduce a philosophical bias.

correctness of his work (namely the three mathematicians Go Yamashita, Mohamed Saïdi and Yuichiro Hoshi), and the detractors in the mathematical community. Ultimately, the point that he seems to want to make is that the level of verification performed by the former group should be sufficient to satisfy the community at large:

[T]he verification activities on the part of the three researchers discussed above *already exceed, by a quite substantial margin* — i.e., in their content, thoroughness, and meticulousness — *the usual level of refereeing for a mathematical journal*. (Mochizuki 2014, p. 7)

[I]t seems to me that the degree of meticulousness and attention to detail exhibited in the verification activities [...]— which, as noted above, exceed, by a substantial margin, the scope of a typical referee’s report for a mathematical journal — together with the wealth of refereeing experience of Yamashita and Saïdi [...] should be regarded as lending *quite substantial weight* to the *extremely positive evaluation* that I received from both of them in the course of these activities. (Mochizuki 2013, p. 4)

In both cases, observe how the concern is with the *thoroughness* and *meticulousness* involved in the checking of the proof. I take these to be a main component of the virtue of rigour, as the responsibilist would put it, as described above. Mochizuki thus directly links whether or not we know a proof is correct to these aspects of virtuous behaviour and character. Furthermore, he insists the following plausible point about offering up opinions on mathematical work:

[T]he essential significance of such an opinion concerning IUTeich lies [...] in the issue of whether or not the opinion *reflects a rigorous and appropriate mathematical understanding of the topic under consideration*. (Mochizuki 2014, p. 13)

Note again the discussions of the rigorousness of the understanding involved in the evaluation of the proof. Importantly, the not-so-veiled criticism of the detractors of his work is that they lack the motivation and dedication required to properly engage with the mathematics, because of the following claimed feature it has:

Yamashita warned that if you attempt to study IU_{Teich} by skimming corners and “occasionally nibbling” on various portions of the theory, then you will not be able to understand the theory even in 10 years; on the other hand, if you study the theory systematically from the beginning, then you should be able to understand it in roughly half a year. (Mochizuki 2014, p. 2)

The upshot is that the normal approach of searching for overlap with prior knowledge they possess and thereby assimilating the new results won't work, according to Mochizuki. In order to be able to pass a substantial judgement on the proof, one must have a proper and full mathematical understanding of the way it works. The way he puts this point is:

[E]very researcher in arithmetic geometry [...] throughout the world is a complete novice with respect to the mathematics surrounding IU_{Teich}, and hence, in particular, is simply not qualified to issue a definitive (i.e., mathematically meaningful) judgement concerning the validity of IU_{Teich} on the basis of a “deep understanding” arising from his/her previous research achievements. (Mochizuki 2014, p. 9)

This quote fits very well with the virtue reliabilist approach: the suggestion being that the researchers who are failing to penetrate the proof are lacking the necessary skills and expertise for grasping the proof or judging its validity. Indeed, Mochizuki's opinion is that established skills may actually be a hindrance:

[T]he most essential stumbling block lies not so much in the need for the *acquisition of new knowledge*, but rather in the need for researchers (i.e., who encounter substantial difficulties in their study of IU_{Teich} and related topics) *to deactivate the thought patterns* that they have installed in their brains and taken for granted for so many years [...] (Mochizuki 2014, p. 11)

We might translate this to the moral that skills in one mathematical domain do not automatically transfer to another, nor are they guaranteed to generate knowledge there. Equally, the talk of “deactivating thought patterns” is also clearly linked to the virtue of open-mindedness, in that one should

be flexible in one's thinking and not reject new mathematics because it is hard, unfamiliar or requires learning new ways of thinking.

Mochizuki also adds more general allusions to the virtues of his colleagues and the link between these and correct mathematics. For example, he refers to the "serenity" of Saïdi's demeanour (p. 3) , as well as the "exception zeal and teamwork" of all three (p. 15).¹⁹ The best quote has been saved for last though:

[...] I have always been a strong advocate of the need, in the case of both domestic and international interaction activities, to *maintain a humble stance dedicated to uncovering the ultimate truth of things* (Mochizuki 2014, p. 14)

Not only does this echo the Stallings quote, in that arriving at mathematical truths requires one to set aside ego, personal excitement and fear, but furthermore it explicitly acknowledges the role of humility as a virtue which directly benefits mathematical inquiry.

So here is the main point that is illustrated in looking at this case. Adopting a virtue-theoretic perspective towards mathematics, as I have been advocating in this chapter, allows us to give the best account of what is going on in the controversy in general and Mochizuki's perspective on it in particular. The issue is one of the division of labour between the creator of a proof and those who have to check it. Critics, believe that a major part of mathematics and mathematical proving is the communication of your results to others. For example, O'Neil says:

It only constitutes a proof if I can readily convince my audience, i.e. other mathematicians, that something is true. Moreover, if I claim to have proved something, it is my responsibility to convince others I've done so; it's not their responsibility to try to understand it (although it would be very nice of them to try).
(O'Neil 2012)

¹⁹I should state again that I don't have a horse in this race. However, if one is concerned about the correctness or success of IUTeich, it seems like zealotry on the part of its proponents might well be counter-productive in comparison to open-mindedness or the humility I will come to momentarily. Similarly, the fact that they are working as a team with Mochizuki might well be seen to detract from the earlier suggestions that they are suitable evaluators of the correctness of the theorems.

For Mochizuki, the view is that the unusually large body of work required for the *abc* proof, and its novelty, difficulty etc., requires a corresponding increase in effort to understand and corroborate the results. The virtue interpretation offers the following analysis of the problems on both sides. For Mochizuki, as we have seen, the verification of the proof will require immersion in the theory of IUTeich rather than the application of previous knowledge to get to an understanding. To be properly immersed, according to Mochizuki, requires a mathematician to be strongly motivated and virtuous in a number of other ways, say through serenity, humility and rigour, which will be necessary for the verification of IUTeich and the *abc* conjecture. Conversely, in demanding so much from his referees, opponents can make the case that Mochizuki has breached a social convention of doing all one can to make the understanding of a new theorem smooth and painless. They could argue that this shows a deficit of virtues of communication and collaboration, or worse that there are vices in play in making the theory impenetrable. Either way, there are huge barriers to understanding for even the practitioners who are actively engaging with the theory, placing epistemic limits on what they can get out of the proof.

In summary, while the controversy is one concerning the verification of a potentially ground-breaking contribution to mathematics, the best way to explain the details of the case is to appeal to the epistemic and mathematical virtues needed in coming to understand a proof.

5.9 Conclusion

In this chapter I have defended a virtue-theory for the philosophy of mathematics, in terms of the Moderate and Strong Proposals:

Moderate Proposal Virtues and vices of mathematicians will be relevant to mathematical knowledge.

Strong Proposal Virtues and vices of mathematicians will be explanatory of mathematical knowledge. In other words, virtue epistemology should be adopted to give the correct epistemology for mathematics.

Not only do these provide a robust and plausible account of mathematical knowledge, but I have shown that these offer rich and fruitful answers to other major related issues in the philosophy of mathematics. In particular,

I have shown that both virtue reliabilists and virtue responsibilists can do better than the Formalist-Reductionist approach in dealing with the nature of proofs and rigour. Indeed, the virtue approach is a natural companion for the switch of emphasis from proofs-as-objects to proving-as-an-activity which I have also been advocating. Furthermore, deploying virtue epistemology has the potential to be bountiful as a major school of thinking in mathematical epistemology. Not only are there well-motivated extensions of the central ideas presented here to the many other ways of gaining mathematical knowledge, those the Traditionalist trend of thinking widely ignores, but additionally the virtue approach expands the relevant epistemological concepts which can be accommodated in our theorising, as I have shown in the previously under-explored area of mathematical understanding.

One thing I have not covered is the Radical Proposal:

Radical Proposal Virtues and vices of mathematicians explain mathematical knowledge and extend to provide alternative answers to other kinds of questions in the philosophy of mathematics, e.g. those concerning ontology, access, metaphysics etc.

The difference was that the Strong Proposal aims to remain neutral on a great deal of traditional philosophy of mathematics and would fit comfortably alongside it. For example, different stances in previous debate might still help to settle how it is that contact with mathematical reality is possible, something which has been left open above. In the Radical Proposal the role of virtues is meant to extend beyond the realm of epistemic concepts such knowledge, understanding and proof, to also provide a new set of answers to traditional questions, potentially in direct conflict with a greater range of the literature. Expanding on how such a project could go, however, would require a great deal of further work, while embroiling us in many more controversies.²⁰

²⁰One view which could be fairly amenable to being adopted to a Radical Proposal-style virtue mathematics would be a mathematical version of Peregrin's inferentialism (Peregrin 2014) (which builds on work by (Sellars 1954) and (Brandom 1994)). Peregrin offers us a broad picture of inferentialism, intending to give an account of meaning in terms of inferential role. He covers both inferentialism in logical systems, defining logical constants through their inference rules, and inferentialism in natural language, seeing meaning as deeply related to rules of usage which over time come to have normative force through corrective behaviour. Peregrin does not engage with the difficulties of informal mathematics, but one could propose the application of his inferentialist theory of natural language (rather than that of formal logic) to the mathematical case. This would allow

To conclude, then, let me emphasise that in many respects the virtue-theoretic approach to mathematical knowledge is a grand departure from orthodox thinking in mathematical epistemology. Nonetheless, what I have shown is that there are interesting and novel answers which it can give to both traditional questions and the difficulties coming out of the philosophy of mathematical practice, difficulties which threaten to undermine a great deal of the orthodox conception anyway. While virtue epistemology is now receiving major attention as a theory of epistemology, it has hitherto not been adopted in the realm of mathematics, leaving an unexplored gap. What I have done here is to begin mapping out some of the terrain of this bountiful new land.

us to incorporate a virtue-theoretic perspective in that we could tie the correctness of practices to the behaviours and virtues of the agents participating in making the rules normative. This would also be radical, in our sense, in that the meanings of mathematical terms would come down to the rules for their manipulation. One might term a view like this *game informalism*, where the acceptable moves are tied to the virtues of agents and their communities. Again, a great deal of additional work would need to be done to make it tenable, but it does not seem that we should rule out the possibility of having answers to questions of ontology in mathematics be secondary to a theory of virtue.

Conclusion

In the last three chapters, I have put essentially the same strategy to use three times to make progress on three different topics. The strategy has been one of showing that there is some significant open area of philosophical questions pertaining to proofs, then connecting it up to an existing literature to demonstrate that there are resources to draw on to answer them. Firstly, I argued that the difficulties encountered in relation to mathematical concepts can draw on recent work in conceptual engineering. Secondly, I argued that the important shift in what we consider to be the important knowledge in mathematics should build on the literature on knowledge-how, particularly anti-intellectualism. Finally, I presented an account of mathematical knowledge, proofs and rigour which was based on virtue epistemology. The moral to draw here is that the philosophy of mathematical practice is not alone. While the ‘maverick’ approach of studying mathematical practice is doing something new in the arena of the philosophy of mathematics, I have shown that there is no shortage of allies in philosophy more broadly.

In the first chapter, I presented a new argument for the failure of the Formalist-Reductionist view of informal proofs, based on a dilemma between agent-independence and agent-dependence, as well as an overgeneration argument to the effect that the Formalist-Reductionist’s picture of how the correspondence between informal and formal proofs justifies rigour and correctness is incompatible with there being many substantially different and equally good formal correspondents for a given informal proof. What the argument brings out is the fact that if the account of informal proofs that the Formalist-Reductionist offers is so dependent on the correspondence between informal proofs and their formal counterparts then the details of how they are matched up matter, and as it turns out these details do not support the claims being made.

In the second chapter, I showed that Beall and Priest's arguments that paradoxes about the notion of informal provability lead to the conclusion that mathematics is inherently inconsistent were unsuccessful. The central criticism I had of these arguments rejected the premise that all of mathematics could be straightforwardly formalised. I suggested that this was too simplistic to be true, and that ways of filling it out that are accurate pictures of the formalisation process no longer support the dialetheist arguments. In particular, I suggested that formalisation is a complex procedure which might lead to many different results, rather than the single super-system that the argument depends on. Priest uses Gödel's results to claim that the trade-off between the axes of completeness and consistency will favour completeness, but I instead argued that we need to see a third axis of formality to properly evaluate the arguments.

In chapter 3, I investigated the relationship between proofs and mathematical concepts. I first compared the ideas underlying Waismann and Shapiro's notions of open texture, showing the difference in the phenomena they are intended to pick out. Next, I worked through Lakatos's *Proofs and Refutations* in detail, looking at how proofs develop the concepts that feature in them according to Lakatos's dialectical philosophy of mathematics. I took a strong stance on the correctness of different interpretations of Lakatos's view of the role of formal proofs, favouring the interpretations of Davis & Hersh and Larvor over those of Worrall & Zahar and Corfield. However, from this discussion it emerged that there was a significant open question for dialectical approaches to philosophy of mathematics concerning the role of formal results, proofs, theories and methods as they are used in modern mathematics. After a brief sojourn showing that several Lakatosian ideas were also found earlier in the overlooked works of Kneebone, I argued that one of the big morals we should draw is that mathematical concepts do display open texture. In the latter half of the chapter I turned to the literature on conceptual engineering as it applies to the development of mathematical concepts. I mainly concentrated on two different proposals, those of Haslanger and Scharp. I deployed Haslanger's distinction between manifest and operative concepts for the mathematical cases of set-theoretic foundations and the difference between formal and informal proofs. I then discussed Scharp's replacement strategy as it works for solving the tension between formal methods and informal mathematics.

The purpose of chapter 4 was to link the literature in epistemology on knowing-how to the arguments by Rav that the important mathematical knowledge is embedded in proofs. I argued that the thought-experiment of *Pythiagora* was fine for Rav's motives of emphasising the importance of proofs, but that his idea of what would happen in such a scenario was incorrectly pessimistic. I argued that a better alternative picture would emphasise the interrelation between the methods, techniques and know-how in proofs and the propositional truths of mathematics. I reviewed the literature on knowledge-how and knowledge-that, covering Ryle, Stanley & Williamson, Hornsby and Wiggins, arguing against the Intellectualist position, and argued that Wiggins's picture of the interdependence of knowledge-how and knowledge-that fits best with case of knowledge of proofs. I backed this up with discussion and examples of mathematical knowledge-how and its importance in mathematical practice. I examined Löwe & Müller's contextualist account of mathematical knowledge, which employs the idea of mathematical know-how as skills, but showed that ultimately it does not succeed because it allows one to know some mathematical fact while believing the opposite. Finally, I examined Larvor's view of informal proofs as depending on inferential actions and expanded on the importance for a view of mathematical knowledge to prioritise proving as an activity over proofs as objects.

In the final chapter, I proposed a virtue-epistemic account of mathematical knowledge. I suggested three different levels of engagement with the literature on intellectual virtues, and argued in favour of the Strong Proposal, that virtue epistemology is the best account of mathematical knowledge. I demonstrated that this view does better than the Formalist-Reductionist family of approaches in accounting for mathematical knowledge through proving, the nature of rigour and mathematical understanding, while also naturally generalising to other types of mathematical knowledge. I also gave a case study of Mochizuki's proof of the *abc* conjecture, showing the fruitfulness of deploying virtue theory in the philosophical study of mathematical practice.

Bibliography

- Abramsky, S. (2008) “Information, Processes and Games” in van Adriaans, P. & van Benthem, J. (eds.) *Handbook on the Philosophy of Information*, Oxford, Elsevier, pp. 483–550.
- Adler, J. (2012) “Epistemological Problems of Testimony” in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition), <http://plato.stanford.edu/archives/sum2015/entries/testimony-episprob/>, accessed May 2016.
- Antonutti Marfori, M. (2010) “Informal Proofs and Mathematical Rigour”, *Studia Logica* 96, pp. 261–272.
- Aristotle (2009) *The Nicomachean Ethics*, Ross, D. (trans.), 2nd Edition revised and with notes by Brown, L., *Oxford World’s Classics*, Oxford, Oxford University Press.
- Avigad, J. (2008) “Understanding Proofs” in Mancosu, P. (ed.) *The Philosophy Of Mathematical Practice*, Oxford, Oxford University Press, pp. 317–353.
- Avigad, J.; Dean, E. & Mumma, J. (2009) “A Formal System for Euclid’s *Elements*”, *The Review of Symbolic Logic* 2, pp. 700–768.
- Azzouni, J. (2004a) “The Derivation-Indicator View Of Mathematical Practice”, *Philosophia Mathematica (III)* 12, pp. 81–105.
- Azzouni, J. (2004b) “Proof And Ontology In Euclidean Mathematics” in Kjeldsen, T. H.; Pedersen, S. A. & Sonne-Hansen, L. M. (eds.) (2004): *New Trends in the History and Philosophy of Mathematics*, Denmark, University Press of Southern Denmark, pp. 117–133.
- Azzouni, J. (2005a) “How To Nominalize Formalism”, *Philosophia Mathematica (III)* 13, pp. 135–159.
- Azzouni, J. (2005b) “Is There Still A Sense In Which Mathematics Can Have Foundations?” in Sica, G. (ed.) *Essays on the Foundations of Mathematics and Logic*, Milan, Polimetria, pp. 9–47.

- Azzouni, J. (2009) “Why Do Informal Proofs Conform To Formal Norms?”, *Foundations of Science* 14, pp. 9–26.
- Azzouni, J. (2013) “That We See That Some Diagrammatic Proofs Are Perfectly Rigorous”, *Philosophia Mathematica (III)* 21, pp. 323–338.
- Ball, P. (2012) “Proof Claimed for Deep Connection Between Primes”, *Nature*, available online: <http://www.nature.com/news/proof-claimed-for-deep-connection-between-primes-1.11378>, accessed May 2016.
- Bancerek, G. (1995) “The Mutilated Chessboard Problem— checked by Mizar” in Matuszewski, R. (ed.) *The QED Workshop II, Technical Report No. L/1/95*, pp. 37–38.
- Bartha, P. F. A. (2010) *By Parallel Reasoning: The Construction and Evaluation of Analogical Arguments*, New York, Oxford University Press.
- Battaly, H. (2008) “Virtue Epistemology”, *Philosophy Compass* 3/4, pp. 639–663.
- Battaly, H. (2014) “Varieties of Epistemic Vice” in Matheson, J. & Vitz, R. (eds.) *The Ethics of Belief*, Oxford, Oxford University Press, pp. 51–77.
- Beall, J. (1999) “From Full Blooded Platonism to Really Full Blooded Platonism”, *Philosophia Mathematica (III)* 7, pp. 322–325.
- Beall, J. (2009) *Spandrels of Truth*, Oxford, Clarendon Press.
- Bell, E. T. (1937) *Men of Mathematics*, New York, NY, Simon and Schuster.
- Benacerraf, P. (1965) “What Numbers Could Not Be”, *The Philosophical Review* 74, pp. 47–73.
- Benacerraf, P. & Putnam, H. (eds.) (1983) *Philosophy of Mathematics: Selected Readings*, 2nd edition, Cambridge, Cambridge University Press.
- Bengson, J. & Moffett, M. A. (2011) “Nonpropositional Intellectualism” in Bengson, J. & Moffett, M. A. (eds.) *Knowing How: Essays on Knowledge, Mind, and Action*, Oxford, Oxford University Press, pp. 161–195.
- Black, M. (1946) *Critical Thinking*, Englewood Cliffs, NJ, Prentice Hall.
- Blackburn, S. (1999) *Think*, Oxford, Oxford University Press.
- Boolos, G. (1987) “A Curious Inference”, *Journal of Philosophical Logic* 16, pp. 1–12.
- Brandom, R. (1994) *Making It Explicit*, Cambridge, MA, Harvard University Press.

- Brown, J. A. (2013) “Knowing-How: Linguistics and Cognitive Science”, *Analysis* 73, pp. 220–227.
- Buldt, B.; Löwe, B. & Müller, T. (2008) “Towards A New Epistemology Of Mathematics”, *Erkenntnis* 68, pp. 309–329.
- Burgess, A. & Plunkett, D. (2013) “Conceptual Ethics I”, *Philosophy Compass* 8/12, pp. 1091–1101.
- Burgess, A. & Plunkett, D. (2013) “Conceptual Ethics II”, *Philosophy Compass* 8/12, pp. 1102–1110.
- Burgess, J. P. (2015) *Rigor and Structure*, New York, Oxford University Press.
- Carnap, R. (1945) “The Two Concepts Of Probability”, *Philosophy and Phenomenological Research* 5, pp. 513–532.
- Carnap, R. (1950) *Logical Foundations of Probability*, Chicago, University of Chicago Press.
- Carroll, L. (1895) “What the Tortoise Said to Achilles”, *Mind* 4, pp. 278–280.
- Castelvechi, D. (2015) “The biggest mystery in mathematics: Shinichi Mochizuki and the impenetrable proof”, *Nature*, available online: <http://www.nature.com/news/the-biggest-mystery-in-mathematics-shinichi-mochizuki-and-the-impenetrable-proof-1.18509>, accessed May 2016.
- Chihara, C. (1979) “The Semantic Paradoxes: A Diagnostic Investigation”, *The Philosophical Review* 88, pp. 590–618.
- Chomsky, N. (1965) *Aspects of the Theory of Syntax*, Cambridge, MA, The MIT Press.
- Clark, P. (2009) “Mathematical Entities” in Le Poidivin, R.; Simons, P.; McGonigal, A. & Cameron, R. P. (eds.) *The Routledge Companion to Metaphysics*, Abingdon, Routledge, pp. 346–356.
- Code, L. (1987) *Epistemic Responsibility*, Hanover, NH, University Press of New England.
- Conrad, B. (2015) “Notes on the Oxford IUT workshop”, available online: <https://mathbabe.org/2015/12/15/notes-on-the-oxford-iut-workshop-by-brian-conrad/>, accessed May 2016.
- Corfield, D. (1997) “Assaying Lakatos’s Philosophy of Mathematics”, *Studies in History and Philosophy of Science* 28, pp. 99–121.

- Corfield, D. (2003) *Towards a Philosophy of Real Mathematics*, Cambridge, Cambridge University Press.
- Davis, P. J. & Hersh, R. (1981) *The Mathematical Experience*, Brighton, Harvester Press Limited.
- Detlefsen, M. (2008) “Proof: Its Nature and Significance” in Gold, B. & Simons, R. (eds.) *Proof and Other Dilemmas: Mathematics and Philosophy*, Washington, Mathematical Association of America, pp. 3–32.
- Ditter, A. (2016) “Why Intellectualism Still Fails”, *The Philosophical Quarterly* 66, pp. 500–515.
- Došen, K. (2003) “Identity Of Proofs Based On Normalization And Generality”, *Bulletin of Symbolic Logic* 9, pp. 477–503.
- Dutilh Novaes, C. (2007) *Formalizing Medieval Logical Theories*, Dordrecht, Springer.
- Dutilh Novaes, C. (2011) “The Different Ways In Which Logic Is (Said To Be) Formal”, *History and Philosophy of Logic* 32, pp. 303–332.
- Elgin, C. Z. (2006) “From Knowledge to Understanding” in Hetherington, S. (ed.) *Epistemology Futures*, Oxford, Clarendon, pp. 199–215.
- Enderton, H. B. (1977) *Elements of Set Theory*, London, Academic Press Inc.
- Ernest, P. (1997) “The Legacy of Lakatos: Reconceptualising the Philosophy of Mathematics”, *Philosophia Mathematica (III)* 5, pp. 116–134.
- Etchemendy, J. (2008) “Reflections on Consequence” in Patterson, D. (ed.) *New Essays on Tarski and Philosophy*, Oxford, Oxford University Press, pp. 263–299.
- Fallis, D. (2003) “Intentional Gaps In Mathematical Proofs”, *Synthese* 134, pp. 45–69.
- Fantl, J. (2012) “Knowledge How”, in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition), <http://plato.stanford.edu/archives/fall2014/entries/knowledge-how>.
- Feferman, S. (1978) “The Logic of Mathematical Discovery vs. The Logical Structure of Mathematics”, in *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, Volume 2: Symposia and Invited Papers*, pp. 309–327.
- Feferman, S. (2012) “And so on... Reasoning with Infinite Diagrams”, *Synthese* 186, pp. 371–386.

- Fricker, M. (2007) *Epistemic Injustice*, Oxford, Oxford University Press.
- Gardner, M. (1988) *Hexaflexagons And Other Mathematical Diversions*, Chicago, The University of Chicago Press.
- Gasarch, W. I. (2002) “The P=?PN Poll”, *SIGACT News Complexity Theory Column* 36, pp. 34–47, available online at <http://www.cs.umd.edu/~gasarch/papers/poll.pdf>, accessed July 2016.
- Giaquinto, M. (2015) “The Epistemology of Visual Thinking in Mathematics”, *The Stanford Encyclopedia of Philosophy (Winter 2015 Edition)*, Zalta E. N. (ed.), available online: <http://plato.stanford.edu/archives/win2015/entries/epistemology-visual-thinking/>, accessed November 2015.
- Geist, C.; Löwe, B. & Van Kerkhove, B. (2010) “Peer Review and Knowledge by Testimony in Mathematics” in Löwe, B. & Müller, T. (eds.) *Philosophy of Mathematics: Sociological Aspects and Mathematical Practice*, London, College Publications, pp. 155–178.
- Ginet, C. (1975) *Knowledge, Perception, and Memory*, Dordrecht, D. Reidel Publishing Company.
- Glick, E. (2011) “Two Methodologies for Evaluating Intellectualism”, *Philosophy and Phenomenological Research* 83, pp. 398–434.
- Gödel, K. (1935) “Review of Carnap 1934: The Antinomies and the Incompleteness of Mathematics” in Feferman, S.; Dawson (Jr.), J. W.; Kleene, S. C.; Moore, G. H.; Solovay, R. M. & van Heijenoort, J. (eds.) (1986) *Kurt Gödel Collected Works, Vol. I*, Oxford, Oxford University Press, pp. 389.
- Gödel, K. (1953) “Is Mathematics Syntax Of Language?” in Feferman, S.; Dawson (Jr.), J. W.; Kleene, S. C.; Moore, G. H.; Solovay, R. M. & van Heijenoort, J. (eds.) (1995) *Kurt Gödel Collected Works, Vol. III*, Oxford, Oxford University Press, pp. 334–364.
- Goldman, A. (2000) “Epistemic Folkways and Scientific Epistemology” in Axtell, G. (ed.) *Knowledge, Belief and Character*, Lanham MD, Rowman & Littlefield Publishers Inc.
- Gonthier, G. (2008) “Formal Proof— The Four-Color Theorem”, *Notices of the American Mathematics Society* 55, pp. 1382–1393.
- Gowers, W. T. (2006) “Does Mathematics Need a Philosophy?” in Hersh, R. (ed.) *18 Unconventional Essays on the Nature of Mathematics*, New York, Springer, pp. 182–200.

- Greco, J. (2000) “Two Kinds Of Intellectual Virtue”, *Philosophy and Phenomenological Research* 60, pp. 179–184.
- Greco, J. (2010) *Achieving Knowledge*, Cambridge, Cambridge University Press.
- Greco, J. & Turri, J. (2011) “Virtue Epistemology”, *Stanford Encyclopedia of Philosophy* (Fall 2015 Edition), Zalta, E. N. (ed.), available online: <http://plato.stanford.edu/archives/fall2015/entries/epistemology-virtue/>
- Greenough, P. (2005) “Contextualism about Vagueness and Higher-Order Vagueness”, *Proceedings of the Aristotelian Society* 79, pp. 167–190.
- Hamkins, J. (2012) “The Set-Theoretic Multiverse”, *Review of Symbolic Logic* 5, pp. 416–449.
- Hanna, G. (1989) “More than Formal Proof”, *For the Learning of Mathematics* 9, pp. 20–23.
- Hanna, G. & Jahnke, H. N. (1996) “Proof and Proving” in Bishop, A. J.; Clements, K.; Keitel, K.; Kilpatrick, J. & Laborde, C. (eds.) *International Handbook of Mathematics Education*, Dordrecht, Kluwer Academic Publishers, pp. 877–908.
- Hanna, G. & Barbeau, E. (2008) “Proofs as the Bearers of Mathematical Knowledge”, *Mathematics Education* 40, pp. 345–353.
- Hanna, G. (2014) “The Width of a Proof”, *PNA* 9, pp. 29–39.
- Hardy, G. H. (1929) “Mathematical Proof”, *Mind* 38, pp. 1–25.
- Haslanger, S. (2000) “Gender and Race: (What) Are They? (What) Do We Want Them to Be?”, *Noûs* 34, pp. 31–55.
- Haslanger, S. (2005) “What Are We Talking About? The Semantics and Politics of Social Kinds”, *Hypatia* 20, pp. 10–26.
- Haslanger, S. & Saul, J. (2006) “Philosophical Analysis and Social Kinds”, *Proceedings of the Aristotelian Society* 80, pp. 89–143.
- Haslanger, S. (2012) *Resisting Reality: Social Construction and Social Change*, Oxford, Oxford University Press.
- Hawley, K. (2011) “Success and Knowledge-How”, *American Philosophical Quarterly* 40, pp. 19–31.
- Hersh, R. (1993) “Proving is Convincing and Explaining”, *Educational Studies in Mathematics* 24, pp. 389–399.

- Hersh, R. (1997) “Prove— Once More And Again”, *Philosophia Mathematica (III)* 5, pp. 153–165.
- Hornsby, J. (1980) *Actions*, London, Routledge and Kegan Paul.
- Hornsby, J. (2012) “Ryle’s Knowing-How, and Knowing How to Act”, in Bengson, J. & Moffett, M. A. (eds.) *Knowing How: Essays on Knowledge, Mind and Action*, Oxford, Oxford University Press, pp. 80–98.
- Inglis, M. & Aberdein, A. (2015) “Beauty Is Not Simplicity: An Analysis of Mathematicians Proof Appraisals”, *Philosophia Mathematica (III)* 23, pp. 87–109.
- Kelp, C. (2011) “In Defence of Virtue Epistemology”, *Synthese* 179, pp. 409–433.
- Kitcher, P. (1984) *The Nature of Mathematical Knowledge*, Oxford, Oxford University Press.
- Kitcher, P. & Aspray, W. (1988) “An Opinionated Introduction” in Aspray, W. & Kitcher, P. (eds.) *History and Philosophy of Modern Mathematics*, Minneapolis, University of Minnesota Press, pp. 3–57.
- Kleiner, I. (1991) “Rigor and Proof in Mathematics: A Historical Perspective”, *Mathematics Magazine* 64, pp. 291–314.
- Kneebone, G. T. (1955) “Abstract Logic and Concrete Thought”, *Proceedings of the Aristotelian Society* 56, pp. 25–44.
- Kneebone, G. T. (1957) “The Philosophical Basis of Mathematical Rigour”, *The Philosophical Quarterly* 7, pp. 204–223.
- Knudson, K. (2015) “A purported new mathematics proof is impenetrable — now what?”, available online: <http://phys.org/news/2015-12-purported-mathematics-proof-impenetrable.html>, accessed May 2016.
- Kunen, K. (1980) *Set Theory: An Introduction To Independence Proofs*, Amsterdam, North Holland Publishing Company.
- Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*, Cambridge, Cambridge University Press.
- Lakatos, I. (1963) “Proofs and Refutations (I)”, *The British Journal for the Philosophy of Science* 14, pp. 1–25.
- Lakatos, I. (1963) “Proofs and Refutations (II)”, *The British Journal for the Philosophy of Science* 14, pp. 120–139.
- Lakatos, I. (1963) “Proofs and Refutations (III)”, *The British Journal for the Philosophy of Science* 14, pp. 221–245.

- Lakatos, I. (1963) “Proofs and Refutations (IV)”, *The British Journal for the Philosophy of Science* 14, pp. 296–342.
- Lakatos, I. (1976) *Proofs And Refutation*, Cambridge, Cambridge University Press.
- Larvor, B. (1998) *Lakatos: An Introduction*, London, Routledge.
- Larvor, B. (2001) “What is Dialectical Philosophy of Mathematics?”, *Philosophia Mathematica (III)* 9, pp. 212–229.
- Larvor, B. (2012) “How to think about informal proofs”, *Synthese* 187, pp. 715–730.
- Larvor, B. (2016) “Why the Naïve Derivation Recipe Model Cannot Explain How Mathematicians’ Proofs Secure Mathematical Knowledge”, *Philosophia Mathematica (III)*, available via early access.
- Leitgeb, H. (2009) “On Formal and Informal Provability” in Bueno, O. & Linnebo, Ø. (eds.) *New Waves in Philosophy of Mathematics*, Basingstoke, Palgrave Macmillan, pp. 263–299.
- Levy, A. (1979) *Basic Set Theory*, London, Springer-Verlag.
- Löwe, B. (forthcoming) “Philosophy or Not? The Study of Cultures and Practices of Mathematics” in Ju, S.; Löwe, B.; Müller, T. & Xie, Y. (eds.) *Cultures of Mathematics and Logic: Selected papers from the conference in Guangzhou, China, 9-12 November 2012*, Basel, Birkhaeuser.
- Löwe, B. & Müller, T. (2008) “Mathematical Knowledge Is Context-Dependent”, *Grazer Philosophische Studien* 76, pp. 91–107.
- Löwe, B. & Müller, T. (2010) “Skills and Mathematical Knowledge”, in Löwe, B. & Müller, T. (eds.) *Philosophy of Mathematics: Sociological Aspects and Mathematical Practice*, London, College Publications, pp. 265–280.
- Mac Lane, S. (1986) *Mathematics: Form and Function*, New York, Springer-Verlag.
- Mac Lane, S. (1998) *Categories for the Working Mathematician*, 2nd Edition, Ann Arbor, MI, Springer.
- MacFarlane, J. (2000) *What Does It Mean To Say That Logic Is Formal?*, *PhD Thesis*, University of Pittsburgh.
- MacKenzie, D. (2001) *Mechanizing Proof*, London, MIT Press.
- Maddy, P. (2007) *Second Philosophy*, Oxford, Oxford University Press.

- Mancosu, P. (ed.) (2008) *The Philosophy Of Mathematical Practice*, Chippenham, Oxford University Press.
- Manders, K. (2008) “The Euclidean Diagram” in (Mancosu 2008), pp. 80–133.
- Martin, U. & Pease, A. (2013) “Mathematical Practice, Crowdsourcing and Social Machines” in Carette, J.; Aspinall, D.; Lange, C.; Sojka, P. & Windsteiger, W. (eds.) *Intelligent Computer Mathematics*, Volume 7961 of the series *Lecture Notes in Computer Science*, pp. 98–119.
- Martin, U. (2015) “Stumbling Around in the Dark: Lessons from Everyday Mathematics” in Felty, A. P. & Middledorp, A. (eds.) *Automated Deduction: CADE-25, Lecture Notes in Artificial Intelligence*, Cham, Springer International Publishing Switzerland, pp. 29–51.
- McCarthy, J. (1995) “The Mutilated Checkerboard in Set Theory” in Matuszewski, R. (ed.) *The QED Workshop II, Technical Report No. L/1/95*, pp. 37–38.
- Mejia-Ramos, J. P. & Weber, K. (2014) “How and Why Mathematicians Read Proofs: Further Evidence from a Survey Study”, *Educational Studies in Mathematics* 85, pp. 161–173.
- Milne, P. (2007) “On Gödel Sentences and What They Say”, *Philosophia Mathematica (III)* 15, pp. 193–226.
- Miscevic, N. (2007) “Virtue-Based Epistemology and the Centrality of Truth (Towards a Strong Virtue-Epistemology)”, *Acta Analytica* 22, pp. 239–266.
- Mochizuki, S. (2013) “On The Verification of Inter-Universal Teichmüller Theory: A Progress Report (As of December 2013)”, available online at <http://www.kurims.kyoto-u.ac.jp/~motizuki/research-english.html>, accessed May 2016.
- Mochizuki, S. (2014) “On The Verification of Inter-Universal Teichmüller Theory: A Progress Report (As of December 2014)”, available online at <http://www.kurims.kyoto-u.ac.jp/~motizuki/research-english.htm>, accessed May 2016.
- Montmarquet, J. (1993) *Epistemic Virtue and Doxastic Responsibility*, Lanham MD, Rowman & Littlefield Publishers Inc.
- Moore, A. W. (1997) *Points of View*, Oxford, Oxford University Press.
- Moravcsik, J. (1979) “Understanding and Knowledge in Plato’s Philosophy”, *Neue Hefte für Philosophie* 15/16, pp. 53–69.

- Moschovakis, Y. (2006) *Notes on Set Theory: Second Edition*, New York, NY, Springer.
- Muntersbjorn, M. M. (2003) “Representational Innovation and Mathematical Ontology”, *Synthese* 134, pp. 159–180.
- Mumma, J. (2010) “Proofs, Pictures, and Euclid”, *Synthese* 175, pp. 255–287.
- Nelsen, R. B. (1993) *Proofs Without Words: Exercises in Visual Thinking*, Washington, DC, The Mathematical Association of America.
- Nelsen, R. B. (2000) *Proofs Without Words II: More Exercises in Visual Thinking*, Washington, DC, The Mathematical Association of America.
- Noë, A. (2005) “Against Intellectualism”, *Analysis* 65, pp. 278–290.
- O’Neil, C. (2012) “The ABC Conjecture has not been proved”, available online: <https://mathbabe.org/2012/11/14/the-abc-conjecture-has-not-been-proved/>, accessed May 2016.
- Paulson, L. C. (2001) “A Simple Formalization And Proof For The Mutilated Chess Board”, *Logic Journal of the IGPL* 9, pp. 477–485.
- Pelc, A. (2009) “Why Do We Believe Theorems?”, *Philosophia Mathematica (III)* 17, pp. 84–94.
- Peregrin, J. (2014) *Inferentialism: Why Rules Matter*, Basingstoke, Pallgrave Macmillan.
- Pettigrew, R. (2016) “Review of John P. Burgess’s *Rigor and Structure*”, *Philosophia Mathematica (III)* 24, pp. 129–146.
- Poincaré, H. (1907) *The Value of Science*, Halsted, G. B. (trans.), New York, NY, The Science Press.
- Pólya, G. (1945) *How to Solve It: A New Aspect of Mathematical Method*, Princeton, NJ, Princeton University Press.
- Priest, G. (1987) *In Contradiction: A Study of the Transconsistent*, Oxford, Oxford University Press.
- Priest, G. (2012) “Mathematical Pluralism”, *Logic Journal of the IGPL* 21, pp. 4–13.
- Rav, Y. (1999) “Why Do We Prove Theorems?”, *Philosophia Mathematica (III)* 7, pp. 5–41.

- Rav, Y. (2007) “A Critique Of A Formalist-Mechanist Version Of The Justification Of Arguments In Mathematicians’ Proof Practices”, *Philosophia Mathematica (III)* 15, pp. 291–320.
- Rayo, A. (2012) *The Construction of Logical Space*, Oxford, Clarendon Press.
- Read, S. (2014) “*Replacing Truth* Book Review”, *The Philosophical Quarterly* 64, pp. 535–537.
- Ripley, D. (2014) “Review of Kevin Scharp’s *Replacing Truth*”, *Notre Dame Philosophical Reviews*.
- Riggs, W. (2003) “Understanding ‘Virtue’ and the Virtue of Understanding” in DePaul, M. & Zagzebski, L. (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford, Oxford University Press, pp. 203–227.
- Roberts, S. (2015) “A Life in Games: The Playful Genius of John Conway”, *WIRED Magazine*, available online: <http://www.wired.com/2015/09/life-games-playful-genius-john-conway/>, accessed May 2016.
- Robinson, J. A. (1991) “Formal and Informal Proofs” in Boyer, R. S. (ed.) *Automated Reasoning*, Dordrecht, Kluwer Academic Publishers, pp. 267–282.
- Robinson, J. A. (2005) “Informal Rigor and Mathematical Understanding”, *Computational Logic and Proof Theory*, Volume 1289 of the series *Lecture Notes in Computer Science*, pp. 54–64.
- Rosefeldt, T. (2004) “Is Knowing-How Simply a Case of Knowing-That?”, *Philosophical Investigations* 27, pp. 370–379.
- Rudnicki, P. (1995) “The Mutilated Checkerboard Problem In The Lightweight Set Theory of Mizar”, available online at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.49.596>, accessed July 2014.
- Rumfitt, I. (2003) “Savoir Fair”, *Journal of Philosophy* 100, pp. 158–166.
- Ryle, G. (1946) “Knowing How and Knowing That: The Presidential Address”, *Proceedings of the Aristotelian Society* 46, pp. 1–16.
- Ryle, G. (1949) *The Concept of Mind*, Harmondsworth, Middlesex, Penguin. Republished 1986.
- Ryle, G. (1971) “Thinking and Self-Teaching”, *Journal of Philosophy of Education* 5, pp. 216–228.

- du Sautoy, M. (2015) “How Mathematicians are Storytellers and Numbers are the Characters”, *The Guardian*, 23rd January 2015.
- Scharp, K. (2013) *Replacing Truth*, Oxford, Oxford University Press.
- Scharp, K. & Shapiro, S. (forthcoming) “Revising Inconsistent Concepts” in Amour-Garb, B. (ed.) *The Relevance of the Liar*, Oxford, Oxford University Press.
- Schlimm, D. (2011) “On the Creative Role of Axiomatics: The Discovery of Lattices by Schröder, Dedekind, Birkhoff, and Others”, *Synthese* 183, pp. 47–68.
- Schlimm, D. (2012) “Mathematical Concepts and Investigative Practice” in Feest, U. & Steinle, F. (eds.) *Scientific Concepts and Investigative Practice*, Berlin, de Gruyter GmbH, pp. 127–147.
- Sellars, W. (1954) “Some Reflections on Language Games”, *Philosophy of Science* 21, pp. 204–228.
- Shapiro, S. (2006) *Vagueness In Context*, New York, Oxford University Press.
- Shapiro, S. (2006) “Computability, Proof, And Open-Texture” in Olszewski, A., Woleński, J. & Janusz, R. (eds.) *Church’s Thesis after 70 Years*, Frankfurt, Ontos Verlag, pp. 420–455.
- Shapiro, S. (2013) “The Open Texture of Computability” in Copeland, B. J.; Posy, C. J. & Shagrir, O. (eds.) *Computability: Turing, Gödel, Church, and Beyond*, Cambridge MA, MIT Press, pp. 153–181.
- Simard Smith, P. & Moldovan, A. (2011) “Arguments as Abstract Objects”, *Informal Logic* 31, pp. 230–261.
- Simmons, G. F. (1991) *Differential Equations with Applications and Historical Notes*, 2nd Edition, New York, NY, McGraw-Hill Book Co.
- Sjögren, J. (2010) “A Note on the Relation Between Formal and Informal Proof”, *Acta Analytica* 25, pp. 447–458.
- Sosa, E. (1980) “The Raft and the Pyramid: Coherence Versus Foundations in the Theory of Knowledge”, *Midwest Studies in Philosophy* 5, pp. 3–26.
- Sosa, E. (1991) *Knowledge In Perspective*, Cambridge, Cambridge University Press.
- Sosa, E. (2007) *A Virtue Epistemology*, Oxford, Oxford University Press.

- Stallings, J. (1965) “How Not To Prove The Poincaré Conjecture” in Bing, R. H. & Bean, R. J. (eds.) *Topology Seminar, Wisconsin, Ann. of Math. Studies* 60, Princeton, NJ, Princeton University Press, pp. 83–88.
- Stanley, J. & Williamson, T. (2001) “Knowing How”, *Journal of Philosophy* 98, pp. 411–444.
- Steiner, M. (1975) *Mathematical Knowledge*, London, Cornell University Press.
- Steinhart, E. (2002) “Why Numbers Are Sets”, *Synthese* 133, pp. 343–361.
- Subramanian, S. (1994) “A Mechanically Checked Proof Of The Mutilated Checkerboard Theorem”, available on-line at <ftp://ftp.cs.utexas.edu/pub/boyer/nqthm/nqthm-1992/examples/subramanian/mutilated-checkerboard.pdf>, accessed July 2014.
- Subramanian, S. (1996) “An Interactive Solution to The $n \times n$ Mutilated Checkerboard Problem”, *Journal of Logic and Computation* 6, pp. 573–598.
- Tanswell, F. S. (2015) “A Problem with the Dependence of Informal Proofs on Formal Proofs”, *Philosophia Mathematica (III)* 23, pp. 295–310.
- Tanswell, F. (forthcoming) “Saving Proof from Paradox: Gödel’s Paradox and the Inconsistency of Informal Mathematics”, in Andreas, H. & Verdée, P. (eds.) *Logical Studies of Paraconsistent Reasoning in Science and Mathematics*, *Trends in Logic* series, Springer.
- Thurston, W. (1994) “On Proof and Progress in Mathematics”, *Bulletin of the American Mathematical Society* 30, pp. 161–177.
- de Toffoli, S. & Giardino, V. (2014) “An Inquiry into the Practice of Proving in Low-Dimensional Topology”, in Lolli, G.; Panza, M. & Venturi, G. (eds.) *From Logic to Practice: Italian Studies in the Philosophy of Mathematics*, Cham, Switzerland, Springer, pp. 315–336.
- Tymoczko, T. (1979) “The Four-Colour Problem and its Philosophical Significance”, *The Journal of Philosophy* 76, pp. 57–83.
- Van Bendegem, J. P. (2014) “The Impact of Philosophy of Mathematical Practice on the Philosophy of Mathematics” in Soler, L.; Zwart, S.; Lynch, M. & Israel-Jost, V. (eds.) *Rethinking Science after the Practice Turn*, Abingdon, Routledge, pp. 215–226.
- Van Kerkhove, B. & Van Bendegem, J. P. (2008) “Pi on Earth, or Mathematics in the Real World”, *Erkenntnis* 68, pp. 421–435.

- Waismann, F. (1951) “Analytic-Synthetic III”, *Analysis* 11, pp. 49–61.
- Waismann, F., (1968) “Verifiability”, in Flew, A. (ed.) *Logic and Language*, Oxford, Basil Blackwell, pp. 118–144.
- Wiggins, D. (2012) “Practical Knowledge: Knowing How To and Knowing That”, *Mind* 121, pp. 97–130.
- Woodin, H. (2011) “The Continuum Hypothesis, the Generic-Multiverse of Sets, and the Ω -Conjecture”, in Kennedy, J. & Kossak, R. (eds.) *Set Theory, Arithmetic, and Foundations of Mathematics: Theorems, Philosophies*, Cambridge Cambridge University Press, pp. 13–42.
- Zagzebski, L. T. (1996) *Virtues Of The Mind*, Cambridge, Cambridge University Press.
- Zagzebski, L. T. (2001) “Recovering Understanding” in Steup, M. (ed.) *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, Oxford, Oxford University Press, pp. 235–251.