

RIGOUR AND PROOF

OLIVER TATTON-BROWN

University of Bristol

Abstract. This paper puts forward a new account of rigorous mathematical proof and its epistemology. One novel feature is a focus on how the skill of reading and writing valid proofs is learnt, as a way of understanding what validity itself amounts to. The account is used to address two current questions in the literature: that of how mathematicians are so good at resolving disputes about validity, and that of whether rigorous proofs are necessarily formalizable.

§1. Introduction. What's going on in mathematical proofs? How do they establish the truth of their conclusions? By *proof* here I mean the kind of proof mathematicians actually write and exchange with each other and accept as valid.

One way to attempt to understand proof is via derivations, the formal objects that logicians use to model deduction. Gentzen intended his system of natural deduction to be

a formalism that reflects as accurately as possible the actual logical reasoning involved in mathematical proofs (Gentzen, 1969, p. 74)

and he described its derivations as having

[a] close affinity to actual reasoning (Gentzen, 1969, p. 80).

One can read this as meaning that all valid mathematical inferences should be instances of the logical rules of natural deduction, or closely related to them.¹ If mathematicians did actually explicitly work according to certain fixed formal rules, then proof would be unmysterious: understanding proof would just come down to understanding the relevant formal rules. However as many authors point out, in reality mathematical proofs do not proceed according to any list of rules that one could specify in advance: there is far too much variety of inferences for that, and new proofs will often contain inferences that are somewhat (or completely) novel.²

Received: August 25, 2020.

2020 *Mathematics Subject Classification*: 00A30, 00A35, 03A05.

Key words and phrases: mathematical rigour, mathematical proof, validity, formalizability, mathematical practice, length of proofs..

¹ This is apparently the reading of Goethe & Friend (2010, pp. 274–275). It is also possible that Gentzen merely intended that the logical substructure of proofs should be expressible in natural deduction—aspects such as introducing and eliminating premises, moving from a statement about an arbitrary x to a statement about all x , and so on.

² These points are argued for instance by Tragesser (1992), Celluci (2009), Leitgeb (2009), Goethe & Friend (2010), and Larvor (2012). A nuanced version of the rule based approach to proof is given by Hamami (2019). Due to the arguments of the previously mentioned

More plausible than a naive rules based view of proof is one on which the correctness of a proof consists in its being *formalizable*—translatable into a derivation according to some given system of formal rules. What exactly this description amounts to will depend on what notion of translation is employed (and on the underlying system of formal rules). Though historically often implicitly assumed by philosophers of mathematics to be correct, numerous authors have recently objected to this view. There have thus been many calls to develop a new, more plausible account of mathematical proof and its epistemology.³

This paper aims to provide such an account, or at least the beginnings of it—found largely in §3 and §4.⁴ The account is not actually of proof in general, but only of *rigorous* proof, rigour being the standard of acceptable proof in much of modern mathematics. I do not view this as a significant limitation. Indeed I argue in §2 that questions like “what is mathematical proof?”, asked in full generality, are unlikely to receive a satisfying answer: there is no univocal notion of proof in mathematics, or at least not one we can expect to obtain a substantial philosophical analysis of. The account of rigorous proof given here is not the last word on the subject, and I note places where it could be expanded on.

The account put forward here takes a somewhat novel approach: instead of focusing exclusively on proofs, considering also the ability of mathematicians to produce and recognize valid inferences.⁵ Where does this ability to recognize validity come from? What can be said about it? Answering these questions is one way to gain insight into what is going on in proofs themselves—how they justify their conclusions.

The resulting account is used to address two questions about proof that have been raised in the literature. The first is that of how mathematicians are so able to generate agreement about the validity of proofs. This phenomenon—that if a mathematician thinks a proof is valid, they can generally convince others, or be convinced themselves of a flaw in it—has been noted by various authors, including Azzouni (2004, pp. 83–84) and Antonutti Marfori (2010, pp. 267, 270–271). Explaining it is one of the major motivations for the “derivation–indicator” view of proof that Azzouni (2004) puts forward. Azzouni’s analysis of proof has met with controversy, with for instance Tanswell (2015) raising what appear to be valid criticisms. Using the account of rigour put forward here, in §5 I give an explanation of this agreement about validity that aims to be simpler and more plausible than Azzouni’s.

The second question is that of whether rigorous proofs are necessarily formalizable. In §7 it is argued not that formalizability is directly required of valid proofs (as is sometimes thought), but that it is a consequence of the norm of rigour in mathematics—as spelled out in §3 and §4—that valid proofs are necessarily formalizable. The worries of Rav (1999) and Weir (2016) about whether the process of filling in intermediate

authors, I do not think an account like Hamami’s can ultimately be successful, though I do not have space to discuss it in detail here.

³ As illustrations of the objections, and the calls for improvement, see for instance Rav (1999), Rav (2007), Celluci (2009), Detlefsen (2009), Leitgeb (2009), Pelc (2009), Goethe & Friend (2010), Antonutti Marfori (2010), Larvor (2012), Weir (2016), De Toffoli & Giardino (2016), and Larvor (2019).

⁴ The account developed here is a fuller version of that sketched in Tatton-Brown (2019).

⁵ Hamami (2019) also focuses on how the skill of assessing rigour results from mathematical training, though from a more formal perspective.

steps in an argument will ever terminate are addressed, and a possible response to the related worry of Pelc (2009) is sketched.

One cannot properly discuss formalizability without discussing what formal system one is targeting, and first §6 considers the argument that mathematicians should really be regarded as working in naive set theory, with its unrestricted comprehension scheme. This has been claimed for instance by Jones (1998) and Leitgeb (2009). If correct this would render the formalizability claim empty of interest, since set theory with unrestricted comprehension is inconsistent and so any argument is trivially formalizable in it (with every inference justified by your favorite set theoretic paradox). This position is evaluated, and dismissed.

These are not all the worries about formalizability that have been or could be raised, and consideration of further qualms—such as the concern expressed by for instance Larvor (2012) that informal arguments may undergo some sort of violence, or essential loss, when formalized—will have to await consideration in further work.

§2. Initial remarks on rigour. The main account of rigour is found in §3 and §4, but first there are some preliminary remarks worth making. To begin with, we will see some examples of nonrigorous mathematics; this helps illustrate the distinctive nature of mathematical rigour, and is also used to argue against the idea that there is a unified notion of “proof” in mathematics that is worth conceptually analyzing. Indeed one intended moral of this paper is that it is rigorous proof, not proof in general, which is the philosophically interesting concept. The latter part of this section summarizes the discussion of rigour given by Burgess (2015), which makes a number of valid, significant points, but which does not address the questions discussed in §1 that this paper attempts to answer.

First, the examples of nonrigorous mathematics. One good example consists of manipulations involving infinitesimals in the 17th and 18th centuries, which—before the introduction of limits into analysis—were not generally rigorous. For a toy example of how they often worked, we can determine the derivative of the function $x \mapsto x^2$. If we let dx be small, then we have

$$\frac{(x + dx)^2 - x^2}{dx} = \frac{x^2 + 2x \cdot dx + dx^2 - x^2}{dx} = \frac{2x \cdot dx + dx^2}{dx} = 2x + dx$$

and then since dx is small we discard it, obtaining $2x$ as the derivative of $x \mapsto x^2$ at x . Arguments along these lines (and more complicated versions) were carried out by various authors, with Fermat perhaps being the first to give this particular kind of calculation (Kline, 1990b, pp. 344–345). These methods met with criticism however, as it was not clear what the status of “small” quantities such as dx was, or what was allowed when manipulating them. Indeed if dx is small but nonzero then the result is only approximate; for the result to be exact, we require that $2x + dx = 2x$, but then we obtain by subtraction that $dx = 0$ and so we cannot divide by dx to begin with. Rolle pointed this out (Mancosu, 1989, pp. 230–231), followed more famously by Berkeley who complained that infinitesimals appeared to be the “ghosts of departed quantities” (Berkeley, 1999, pp. 80–81).

Other common methods in the 17th and 18th centuries also lacked rigour. Often arguments proceeded by assuming that what held for the finite also held for the infinite, with infinite series being manipulated as though they were finite sums, without worrying

about issues of convergence. For instance Jacob Bernoulli argued (essentially) that

$$\begin{aligned} 1 &= \left(1 + \frac{1}{2} + \frac{1}{3} + \dots\right) - \left(\frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \dots\right) \\ &= \left(1 - \frac{1}{2}\right) + \left(\frac{1}{2} - \frac{1}{3}\right) + \left(\frac{1}{3} - \frac{1}{4}\right) + \dots = \frac{1}{2} + \frac{1}{6} + \frac{1}{12} + \dots \end{aligned}$$

despite knowing that the sum $1 + \frac{1}{2} + \frac{1}{3} + \dots$ is infinite (Kline, 1990b, p. 443).⁶ As illustrated by many further examples in Kline (1990b, Chap. 20), infinite series were freely manipulated in this period without worrying about convergence, despite these methods sometimes leading to false or contradictory conclusions.

Arguments like these gradually came to be seen as unacceptable. In the 19th century both the calculus and the study of infinite series were rephrased in terms of the concept of limit, putting them on a firm footing (Kline, 1990a, chap. 40). Infinitesimals were then largely eschewed in analysis until Robinson demonstrated how one could in fact reason rigorously about them, via the logical concept of a nonstandard model (Robinson, 1996).

Other ways of reasoning formerly regarded as valid also came to be shunned, such as appeals to intuition. Proofs were demanded of even very intuitive statements, with the Jordan curve theorem being a famous example. This states that if $\phi : S^1 \rightarrow \mathbb{R}^2$ is continuous and injective then $\mathbb{R} \setminus \phi(S^1)$ consists of exactly two connected components, one of which is bounded and the other unbounded, and $\phi(S^1)$ is the boundary of each component. On an intuitive understanding of continuity, this can appear obvious, and it takes some effort to imagine a curve ϕ in such a way that the conclusion is not immediate. However on the ε - δ definition of continuity, the fact is not obvious at all—and providing a rigorous proof turned out to be very difficult. Bolzano had already noted that this fact required proof (Coulston, 1970, p. 274) and this proof was only provided by Jordan (published in Jordan, 1887) more than 30 years after Bolzano's death.⁷

There are a number of lessons to be drawn from these examples. First, they can be used to assess the adequacy of an account of rigour, such as that given in §3 and §4—it needs to be able to explain why they are not considered rigorous. Second, such examples illustrate that an argument does not have to be rigorous to be reliable, or explanatory, or valuable. Arguments involving infinitesimals in the 18th century could plausibly be all three (such as the differentiation example above), as could manipulations of infinite series. Much of modern nonrigorous physics and engineering is also presumably reliable, explanatory and valuable. Nonetheless rigour does bring benefits, some of which will be discussed in this paper.

These examples also illustrate the broadness of the notion of mathematical proof in times past. This, I believe, tells against the desire to seek a philosophical account of

⁶ One can adjust this argument to make it rigorous by telescoping the partial sums:

$$\frac{1}{2} + \frac{1}{6} + \dots + \frac{1}{n(n+1)} = \left(1 - \frac{1}{2}\right) + \left(\frac{1}{2} - \frac{1}{3}\right) + \dots + \left(\frac{1}{n} - \frac{1}{n+1}\right) = 1 - \frac{1}{n+1}$$

⁷ Some controversy followed Jordan's proof, with Veblen (1905) claiming it was flawed, and claiming to give the first rigorous proof. However Hales (2007) argues that Jordan's proof was basically valid, though perhaps not as polished as it could be.

the general notion of proof—to discover what proof in general “really is”, or where its boundaries are drawn. Indeed in the above examples, reasoning of various kinds all distinct from the usual mathematical paradigm of deduction is seen. To make arguments involving infinitesimals one postulates a new manner of calculation, in which a quantity is treated at one stage as nonzero, and later as small enough that it can be neglected. This can otherwise be viewed as the postulation of a new kind of entity with these apparently odd properties. Either way, it is essentially a form of abductive reasoning: reasoning of a kind which is not justified by anything that has gone before, but instead by its immense success in solving all manner of differential problems.⁸ It is not so different to the postulation of new principles or entities in physics, except that it is confirmed by mathematical applications and deductions, rather than by experiments. Then the manipulation of infinite series as though they were finite is essentially a case of argument by analogy, again backed up by its apparent success in solving problems. Finally we have appeals to intuition, delivering conclusions that one finds very hard to doubt because of one’s intuitive grasp of the concepts involved—not so different from the intuition that philosophical zombies could exist, or that nothing can cause itself.

I do not think there is much to be gained by seeking to discover what these disparate forms of reasoning “have in common”. They are all taken to justify high credence in their conclusions, and they all concern abstract, mathematical subject matter, but beyond that there may not be much more to say. Certainly one could conduct a fruitful investigation into abductive reasoning, or intuition as a form of evidence, but there is unlikely to be much distinctive to say about either in the context of mathematical proof that does not apply to more general contexts. There is also not much I think to say about why such methods were accepted, beyond that they were felt to be reliable. The best assessment of proof in general may just be that there are different kinds of permissible actions that one may carry out, as Larvor (2012) puts it. It is *rigorous* proof which appears to be more deserving of philosophical attention: this is where the ideal of flawless deduction that Euclid aspired to takes its purest form (Burgess, 2015, pp. 36–38), and where mathematical reasoning is found at its most distinctive, and epistemologically robust.

The variety of kinds of inference allowed in proofs historically makes it clear that it is only for rigorous mathematics that one could defend the formalizability of arguments. For instance what would a formal system for arguments like Bernoulli’s, involving manipulation of divergent series, look like? Would it formalize analogies between the finite and the infinite? Similarly there is no reason to think that irreducibly intuitive reasoning would be formalizable.

Given the major differences between the historical standard of proof and what modern rigour permits, an obvious question is how and why the shift to rigour came about. Considering that would take us too far afield, but it is one topic which Burgess discusses (Burgess, 2015, chap. 1) in his account of rigour, to which we now turn. Burgess is mainly concerned with implications of the norm of rigour for the debate over structuralism, and does not give explicit arguments concerning the kinds of questions raised in §1—the epistemology of proof, the ability of mathematicians to

⁸ As an anonymous referee remarks, part of the confidence in the use of infinitesimals stemmed from how in many cases one could use them to solve a problem—like the volume of a sphere—whose solution was already known, and thus check that the infinitesimal methods gave the correct results.

reach agreement about the validity of proofs, the issue of formalizability and so on. Nonetheless some observations he makes are worth highlighting.

Burgess emphasizes that any piece of rigorous mathematics takes place in a context of existing results and definitions which can be appealed to Burgess (2015, pp. 149–158). One can then extend the boundaries of knowledge with a new argument, stringing together definitions and proofs of propositions, appealing to existing results and using existing concepts where needed. Burgess also emphasizes that often it does not matter whether a fact one is appealing to is a basic principle or a consequence thereof, or how concepts used were actually defined as long as the properties one needs of them do hold. This is the basis of his critique of structuralism as a metaphysical position.⁹

Both components of the rigorous process—definitions, and proofs of propositions—merit some attention. Burgess notes that when introducing a new concept, one is required a clear definition in terms of existing ones (Burgess, 2015, p. 7). This definition does not have to be completely formal—for instance one can state that a vector space is a set equipped with an abelian group structure and a scalar multiplication operation, without specifying exactly how this is coded set theoretically: as a triple $(V, +, \cdot)$, or $((V, +), \cdot)$, or as $(+, \cdot)$, or in some other way. It just needs to be clear that the definition could be made completely precise in such a way that all uses made of the concept would be valid.

Burgess also discusses what the standard of rigour requires for proofs of propositions. He considers various possibilities, and ultimately comes to the (tentative) conclusion that:

What rigor requires is that each new result should be obtained from earlier results by presenting enough deductive steps to produce conviction that a full breakdown into obvious deductive steps would in principle be possible (Burgess, 2015, p. 97)

This I think is basically right. However there is more to say before this can be brought to bear on the issues discussed in §1. If we are interested in the epistemology of proof, then this is only a sketch rather than a full account. How is this conviction generated? How is it reliable? If mathematicians are judging formalizability in principle (which is roughly what “full breakdown into obvious deductive steps” might amount to), how are they able to judge this? As discussed in §1, most mathematicians have no experience of or interest in formalization, after all. There is also the question of how mathematicians are so good at resolving disputes, discussed in §1, which we could hope to answer. Addressing these issues is the purpose of the remainder of the paper.

§3. A rigorous eduction. As mentioned in §1, the account of rigorous proof put forward here uses a somewhat novel approach: to try to understand the skill of mathematicians in judging and producing rigorous proofs by thinking about how this skill is acquired. For this we will start at the beginning. There are many different universities around the world that teach rigorous mathematics, and they may teach it

⁹ He argues that exactly how the concepts one uses were defined is often irrelevant, as all one will need are certain derived properties. Thus one need not care about how things were defined when doing mathematics. It is this irrelevance of definitions that Burgess argues has been mistaken by structuralists for a metaphysical truth about the nature of mathematical structures, with structuralists hoping to infer for instance that mathematical objects have only general structural properties (Burgess, 2015, chap. 3).

- 2.1.8 Theorem** (a) If $a \in \mathbb{R}$ and $a \neq 0$, then $a^2 > 0$.
 (b) $1 > 0$.
 (c) If $n \in \mathbb{N}$, then $n > 0$.

Proof. (a) By the Trichotomy Property, if $a \neq 0$, then either $a \in \mathbb{P}$ or $-a \in \mathbb{P}$. If $a \in \mathbb{P}$ then by 2.1.5(ii), $a^2 = a \cdot a \in \mathbb{P}$. Also, if $-a \in \mathbb{P}$, then $a^2 = (-a)(-a) \in \mathbb{P}$. We conclude that if $a \neq 0$, then $a^2 > 0$.

(b) Since $1 = 1^2$, it follows from (a) that $1 > 0$.

Copyright © 2000 John Wiley & Sons, Inc. All rights reserved.

Fig. 1. “Introduction to Real Analysis”, Bartle and Sherbert, p. 26 (proof of (c) is omitted).

in somewhat different ways, but there are some common features that can be pointed to. Students are generally taught the basics of rigorous proof by seeing and working through examples, paired with descriptions of how and why the reasoning involved works. An example of a basic early result students might see is displayed in Figure 1.

This demonstrates a fact commonly assumed without question: that the square of a nonzero real is positive. Probably the only part of the argument that requires explanation is the symbol \mathbb{P} , which denotes the set of (strictly) positive real numbers. Axioms concerning this set have been stated in the textbook a few pages previously. The relevant axioms are that:

- (i) For any $a \in \mathbb{R}$, either $a \in \mathbb{P}$ or $(-a) \in \mathbb{P}$ or $a = 0$, with exactly one of these holding.
- (ii) If $a, b \in \mathbb{P}$ then $a \cdot b \in \mathbb{P}$.

From these the above proof proceeds straightforwardly, arguing by cases.

The main thing to note about this proof is just how incredibly detailed it is. Virtually all of the logical structure of the argument is right there on the page. There are places where one could be even more explicit, in particular in the assertion that $a^2 = (-a)(-a)$, and indeed this follows immediately from the facts that $1 = (-1)(-1)$ and that $(-a) = (-1)a$, which are both given as exercises (Bartle & Sherbert, 2000, p. 29). Nonetheless the proof is very close to the formal level and would be no challenge to formalize.

We can call this level of very great detail that proofs can be carried out at the “week 2 level of detail”. Of course students may not see this particular argument in week 2, or at all; it is just a convenient name. We are not defining the “week 2 level of detail” here in terms of what is comprehensible to certain students—instead we give examples of basic arguments at this level of incredible detail, such as the above and also for instance basic number theoretic results (Taylor & Garnier, 2014, theorem 6.2; Silverman, 2012, Lemma 7.1) or basic results from linear algebra (Axler, 1997, Propositions 1.1–1.6).

As students learn the subject they won’t just be passively reading proofs like this. They will also typically (and importantly) be proving these kinds of basic facts themselves, demonstrating them with arguments written out at this very explicit level of detail. The hope is that by doing this they will gain what we can call “proficiency at week 2 detail”, the ability to prove simple facts like this one by chaining together these kinds of very basic steps.

A bit more will be said about how this basic level of proficiency is gained later in this section. For now we proceed onwards through the curriculum. As time passes the

6.2.1 Interior Extremum Theorem *Let c be an interior point of the interval I at which $f: I \rightarrow \mathbb{R}$ has a relative extremum. If the derivative of f at c exists, then $f'(c) = 0$.*

Proof. We will prove the result only for the case that f has a relative maximum at c ; the proof for the case of a relative minimum is similar.

If $f'(c) > 0$, then by Theorem 4.2.9 there exists a neighborhood $V \subseteq I$ of c such that

$$\frac{f(x) - f(c)}{x - c} > 0 \quad \text{for } x \in V, x \neq c.$$

If $x \in V$, and $x > c$, then we have

$$f(x) - f(c) = (x - c) \cdot \frac{f(x) - f(c)}{x - c} > 0.$$

But this contradicts the hypothesis that f has a relative maximum at c . Thus we cannot have $f'(c) > 0$. Similarly (how?), we cannot have $f'(c) < 0$. Therefore we must have $f'(c) = 0$. Q.E.D.

Copyright © John Wiley & Sons, Inc. All rights reserved.

Fig. 2. “Introduction to Real Analysis”, Bartle and Sherbert, p. 168.

arguments the students are presented with will gradually get faster, and have fewer of the details filled in. Sometime later—perhaps a few months, or a term—they may encounter an argument like that in Figure 2.

This theorem presents another fundamental fact: that if a function on an interval has a “relative extremum”—a local maximum or a local minimum—at an interior point, and is differentiable there, then the derivative must be zero. This is clear by visualizing the situation, but we are doing rigorous mathematics so are not satisfied with that, and we demand a proof.

The argument given is again fairly detailed, but is slightly less explicit than the previous example: not all the details are there. It only actually covers the case where $f'(c) > 0$, showing that this cannot happen, and the task of showing that $f'(c) < 0$ cannot occur is left to the reader. If the reader has understood the argument they should have no problem seeing how this would go, or writing it out. This aspect of the proof is fundamental to the way we learn rigorous mathematics. Students will hopefully not be treating proofs like the deliverances of some oracle: lecturers will ideally encourage them to engage with the proofs, to see if they could have proved the results themselves, to see if they can prove similar results by similar methods, and to see if they can fill in any parts where the proof is sketchy, and check any parts of the proof they are not sure about in more detail.

The course did not start by teaching students the week 2 level of detail material just to pad the schedule. The hope is that now when they meet an argument like this which is a little bit faster, strung together out of inferences that are simple but not necessarily completely basic, they can check any inference they are not sure of by proving it in more detail, using their “proficiency at week 2 detail” that they have hopefully already attained. Thus they can sharpen their judgment of which simple (but not completely basic) inferences are valid, checking such inference whenever necessary by seeing if they can be proved.

Students won’t just be seeing theorems like this however. They will also be proving these kinds of slightly higher level statements themselves, by stringing together inferences that are simple (but not necessarily completely basic). By doing so they

9.4.12 Differentiation Theorem *A power series can be differentiated term-by-term within the interval of convergence. In fact, if*

$$f(x) = \sum_{n=0}^{\infty} a_n x^n, \quad \text{then} \quad f'(x) = \sum_{n=1}^{\infty} n a_n x^{n-1} \quad \text{for} \quad |x| < R.$$

Both series have the same radius of convergence.

Proof. Since $\lim(n^{1/n}) = 1$, the sequence $(|na_n|^{1/n})$ is bounded if and only if the sequence $(|a_n|^{1/n})$ is bounded. Moreover, it is easily seen that

$$\limsup (|na_n|^{1/n}) = \limsup (|a_n|^{1/n}).$$

Therefore, the radius of convergence of the two series is the same, so the formally differentiated series is uniformly convergent on each closed and bounded interval contained in the interval of convergence. We can then apply Theorem 9.4.4 to conclude that the formally differentiated series converges to the derivative of the given series. Q.E.D.

Copyright © John Wiley & Sons, Inc. All rights reserved.

Fig. 3. “Introduction to Real Analysis”, Bartle and Sherbert, p. 270.

will hopefully gain what we can call “proficiency at term 2 detail”, the ability to prove these slightly higher level inferences by stringing together simple inferences, and to reliably judge the validity of simple inferences (checking whenever necessary by proving them at the week 2 level of detail).¹⁰ Again we define the “term 2 level of detail” by giving examples, such as the above and also for instance from Bartle & Sherbert (2000, Theorem 5.2.1), Axler (1997, Propositions 2.9 and 2.13) and Silverman (2012, Lemma 9.2).

As the terms go by the students are gradually exposed to more and more condensed arguments. After another year or so they might meet an argument like that in Figure 3.

Here the students see a proof that power series can be differentiated term by term. The proof is another step up in terms of compression, in terms of relying on the intelligence of the reader. This can be seen in the first line, where the reader is expected to see that a certain sequence is bounded if and only if another sequence is. Also in the second sentence, the reader is expected to “easily see” that a certain equation holds. These statements are extremely plausible; and if a student has any doubt, they can check them by proving them in more detail, using the “proficiency at term 2 detail” ability they have hopefully gained. They do not need to take these statements on trust, and they do not need to guess.

Again we can talk roughly about this “year 2 level of detail”, giving further examples of arguments at about this level of detail to help explicate it, again for instance from analysis (Bartle & Sherbert, 2000, Theorems 9.3.2 and 10.1.3), from number theory

¹⁰ A note on terminology. I find it natural to speak of *proving inferences*, as in proving them in greater detail, though strictly speaking this may be a category error: inferences are things that we draw, assess, or justify, and we normally only speak of proving statements and propositions. Nonetheless I think it is clear what is meant—replacing a given inference by a chain of intermediate inferences, which collectively constitute a proof of the conclusion from the premises—and it is a convenient and expressive idiom.

(Silverman, 2012, Theorem 42.1), and also from ring theory (Aluffi, 2009, Propositions III.3.11 and III.4.5) and complex analysis (Bak & Newman, 2010, Proposition 3.1).

This process continues in the obvious way. As the years progress a student is exposed to gradually faster and faster arguments, arguments where gradually more and more of the details are left out and more is left to the reader's intelligence. We can pick out further levels of detail a student will encounter, in the same way as above. First, we define a "year 3 level of detail" by giving examples, now with a wider variety: from functional analysis (Rudin, 1987, Theorems 4.6–4.12), complex analysis (Conway, 1978, sec. IV.2), measure theory (Fremlin, 2010, chap. 12), general topology (Munkres, 2000, sec. 33), algebraic topology (Lee, 2000, chap. 13), differential geometry (Lee, 2012, chap. 3), commutative algebra (Aluffi, 2009, sec. V.1), representation theory (James & Liebeck, 2001, chap. 6), number theory (Niven, Niven, & Zuckerman, 1991, sec. 1.2), combinatorics (Szemerédi, 1975, Facts 1 and 2), logic (Cori & Lascar, 2000, sec. 1.1), and category theory (Awodey, 2010, Proposition 2.10). Again these may not be arguments a given student actually sees in their third year, but the level of detail is intended to be one that competent third year students will be gaining proficiency at, for both reading and writing proofs.

Detail here is not the same thing as accessibility. An argument can be very detailed but still difficult, for instance because it involves difficult concepts, or relies on difficult results, or because the result is poorly motivated and the proof strategy unexplained—or just because the argument is too long. Detail here means explicitness, and proximity to definitions, and how much the proof says of what could be said. It is the antonym of "how much is left out".

Naturally these predicates "week 2 level of detail", "year 2 level of detail" and so on will be vague: we may not be able to always determine precisely whether an argument is at the year 2 level of detail or not, just as we may not be able to decide whether a jumper is red, or perhaps orange instead. That does not undermine these predicates' validity or usefulness. Although there will be borderline cases, there will also be cases where we can in fact state with confidence that an argument is at around the year 2 or year 3 level of detail, rather than the week 2 or graduate level of detail (defined shortly). Of course an argument may not all take place at the same level of detail, so that describing the different levels of detail its parts take place at may be more appropriate than trying to shoehorn the whole argument into one category—as with a multicolored jumper. One issue with these levels of detail that does not have such a comparison with jumpers is that it can potentially be quite difficult to compare the detail of pieces of mathematics from very different areas, where the reasoning is of a very different style. We can mitigate that as here by giving examples from a wide range of areas when characterizing levels of detail.

We now continue in the same way, defining further levels of detail a student will encounter. It should perhaps be emphasized that this terminology of levels of detail is new terminology I am introducing, and not a standard part of mathematical discourse. There are times one might see something like it used—for instance if a mathematician presented an unconvincing argument to a colleague, and after some questioning the colleague asked them to explain it more slowly, like they were talking to a grad student. Also, concepts like these can perhaps be seen as implicitly underlying some mathematicians' talk of detail in mathematics, an example of which will be seen in §5 when discussing how these levels of detail can help mathematicians resolve disputes about the validity of proofs.

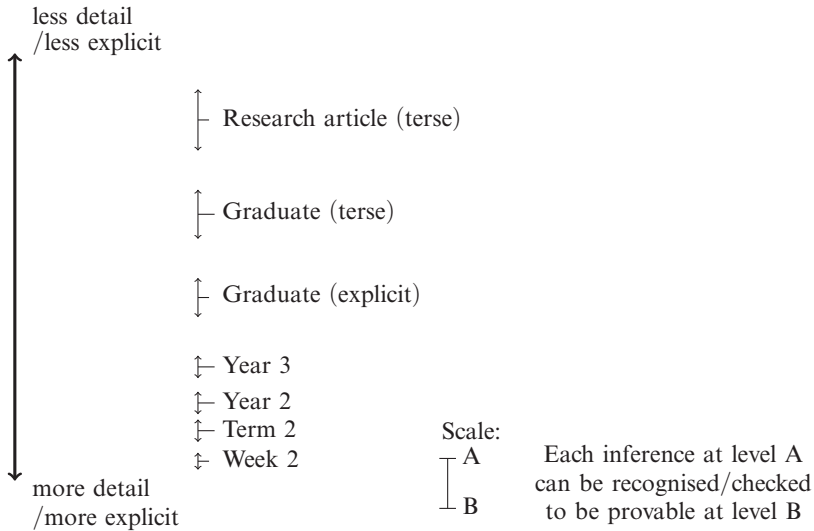


Fig. 4. Levels of detail.

We will actually now pick out two different graduate levels of detail. First, we define the “graduate level of detail (explicit)” by giving examples, from functional analysis (Banach, 1987, Theorem II.1), complex analysis (Conway, 1978, sec. IV.6), measure theory (Schwartz, 1954), general topology (Walker, 1974, secs. 1.1–1.16), algebraic topology (Switzer, 2002, chap. 4), differential geometry (Hirsch, 1976, sec. 1.3), algebraic geometry (Eisenbud & Harris, 2000, sec. I.1.4), commutative algebra (Eisenbud, 1995, chap. I.2), representation theory (Fulton & Harris, 1991, Lecture 4), number theory (Niven *et al.*, 1991, sec. 5.7), combinatorics (Erdős, 1947, Theorem 1), logic (Prawitz, 1965, chap. I–III), and category theory (MacLane, 1998, secs. II.3–II.6).

The basic idea is hopefully now clear, but we can keep going and pick out a “graduate level of detail (terse)” by giving some examples: from algebraic topology (Hatcher, 2001, sec. 2.2), differential geometry (Thurston, 1997, sec. 2.7), algebraic geometry (Hartshorne, 1977, sec. II.3) and category theory (MacLane, 1998, chap. IX). As seen in the examples above, some research mathematics is written out at a level of detail already covered: for instance Szemerédi (1975, Facts 1 and 2) at the year 3 level of detail, and Banach (1987, Theorem II.1), Schwartz (1954), Erdős (1947, Theorem 1) and Prawitz (1965, chap. I–III) at the graduate level of detail (explicit). However plenty of research mathematics does take place at a level of greater compression, and one could keep going and pick out levels of this, named perhaps “research article level of detail (terse)”, “research article level of detail (very terse)”, and maybe one or two more. A range of levels of detail we can obtain in this way is seen in Figure 4.

To be clear, there is nothing privileged about the levels of detail listed in Figure 4. There is a continuum of levels of detail that we could potentially pick from (perhaps idealizing somewhat, given that statements in proofs are finite objects), and nothing to mark out those in Figure 4 as special; they are just useful examples for the purposes of this paper.

A little more should be said about the upper reaches of mathematics, at the top of Figure 4 and beyond. In fact it is clear that nothing too different happens as one

approaches the research frontier—the gradual ascent to levels of greater and greater compression continues. That nothing radically different is happening can be seen in the ability of professors to take the most important, beautiful or useful results in a field—once it has matured—and collate them into a textbook with proofs accessible to graduate students. In the process proofs may be simplified or altered, but there is never any great obstacle to writing out what was once research level mathematics at a level of detail that graduate students can follow. One can see this in for instance the titles in the Springer Graduate Texts in Mathematics series.

A particularly extreme example of this is given by the case of Perelman's proof of the Poincaré conjecture. This was one of the most major conjectures in mathematics, and the subject of a Clay Millennium Prize. In 2002 and 2003 Perelman uploaded three papers containing a claimed proof (Perelman, 2002, 2003a, 2003b). The papers were written at a very high level, containing mathematics sketchy enough that despite only totally 70 pages, it took teams of mathematicians 3 years—working in correspondence with Perelman—to verify the argument as correct. Perelman was then offered a Field's medal, and subsequently a Millennium Prize, both of which he declined. Since 2006 more detailed expositions of his argument have been produced, such as Morgan & Tian (2007), which is a textbook intended to be accessible to graduate students. Indeed it looks to be at about the graduate level of detail (explicit) or graduate level of detail (terse). It comes in at 521 pages—about 8 times longer than Perelman's original papers. This is quite an increase, but even so it shows that there is not too dramatic a leap in terms of compression from maths at a level graduate students can understand to some of the most concise mathematics acceptable as a proof.

Now to say a little more about the lower end of Figure 4. First, a potential issue is that some areas of maths are generally only studied at a high level, because they have substantial mathematical prerequisites or typically involve subtle or complex arguments. This is the case with harmonic analysis, the theory of functions of several complex variables, and modern algebraic geometry, amongst other areas. This presents a potential problem with regards to the lower levels of detail in Figure 4, since for instance no-one does algebraic geometry in week 2 of their degree. In some cases this is unproblematic since expanding an argument in great detail will lead to basic inferences like those found in other areas—for instance combinatorics or analysis, in the case of harmonic analysis. In other cases reasoning used is more *sui generis*. Nonetheless in these cases I think it still makes sense to talk about what an inference carried out at say the term 2 level of detail would look like. In fact we sometimes see this happen: when a new manner of argument is introduced, even to advanced students, a few very explicit examples are often given of how it works. This is so that students have a sense of what is underlying more complex arguments, and know what to fall back on if they ever find more complex arguments hard to follow or produce. An instance of this from differential topology—a fairly advanced subject, done in full generality—is seen in Lee (2012, proposition 2.4). Here Lee is giving an example of how to use the smooth charts on a manifold to prove local facts about them, and it is written up in great detail to make clear to students how this works, at around the week 2 or term 2 level of detail. One could do the same for other advanced subjects, for instance writing out the arguments of Eisenbud & Harris (2000, sec. I.1.4) from algebraic geometry mentioned above at the term 2 level of detail.

A second issue about levels of greater detail is one that was rather passed over in the discussion above: what takes place in the initial stages of learning rigorous mathematics,

before the ascent up the levels of compression can start. To begin with, students will be taught how the basics of proof work, and what to do with the logical vocabulary $\wedge, \exists, \neg, \forall$ and so on, by a combination of examples and informal descriptions of what is going on and why. For instance they will (hopefully) learn by seeing and working through examples that to prove a statement $\forall n \phi(n)$ about all numbers, one can take an unknown number n , prove that $\phi(n)$ holds without assuming anything special about n , and then deduce that indeed $\forall n \phi(n)$ holds—essentially the \forall -introduction rule from natural deduction. The logical workings of proof can be stated in simple, clear, precise form—in terms of the hypotheses active at each stage, and how to introduce and exploit them—and when students can grasp how this works from examples it is not such a great surprise. Anyway not all students do manage to learn the rules from examples, and how to best teach the logic of proof is much discussed in the mathematics education literature (see for instance Epp, 2003). As well as the logical vocabulary students will learn the basics of set theory, including how to determine if two sets are equal via extensionality (whether this is an axiom or an inference rule is not important to the students, and the distinction may not be clear to them), and they will be shown various acceptable ways of forming new sets—power sets, cartesian products and so on.

There is a further basic aspect of proof that students are expected to infer from examples, and this is the ability to prove results by describing algorithms or procedures to achieve some desired mathematical goal, with Euclid's algorithm being an early example students often encounter (Silverman, 2012, pp. 33–34), and further examples coming in linear algebra (Axler, 1997, propositions 2.6 and 2.7; Artin, 1991, pp. 14–15) and other areas. Though some such proofs can be rephrased as arguments by induction, sometimes they may implicitly require the definition of functions by recursion, and this is not usually something early undergraduates will be in a position to justify formally—the set theoretic treatment of recursion is typically taught later on in a more general form that applies to recursions on all ordinals (such as in Jech, 2006, theorem 2.15). In fact this does not present any sort of problem, and is not so different to the cases above where students grasp principles from examples; in this case the general principle implicitly underlying these kinds of recursive arguments is the axiom of dependent choice. This states that if X is a set and R is a binary relation on X such that for all $x \in X$ there is $y \in X$ with xRy , then for all $x \in X$ there is a sequence (x_0, x_1, \dots) of elements of X where $x_0 = x$ and for all i , $x_i R x_{i+1}$. That this is a statement rather than an inference rule, and not a basic axiom of set theory (it is deduced from the axiom of choice) is not important here. It is intuitive, and can be stated clearly, simply and precisely, and it is not a surprise that students can grasp from examples what kinds of arguments are in line with it.

The basic axioms governing sets are exempted from the general requirement that inferences be justifiable by proofs in greater detail. One can rightly assert this without actually deciding on which the basic axioms are; for instance it does not matter whether one regards the statement that function sets B^A exist as a basic axiom, or as justified by an argument that appeals to more basic axioms (union, separation, pair set and power set perhaps). Whatever the basic axioms are, they need not be justified by a proof, and other basic properties of sets are justified in terms of them (perhaps out of sight of students). There has been some discussion of exactly what means of set formation are allowed in mathematics—in particular, whether the unrestricted comprehension principle of naive set theory is used—and this issue will be addressed in §7.

With these points addressed, we have the essentials of how rigour is learnt. We can pick out different levels of detail maths is done at by giving examples, and students proceed upwards through these levels of detail as described above: once they have gained proficiency at a certain level of detail they are in a position to engage with more concise arguments, with a tutored sense of what less explicit are valid—tutored by their experience at proving such inferences. If they are ever unsure, they can use their existing proficiency to check a difficult inference and see if it can in fact be justified by a proof; and if so, they can ask why they were suspicious about it, and consider how to adjust their instincts to recognize such inferences as valid in future. It is essential to the normal process of learning rigorous mathematics that students are in a position to check inferences for themselves in this way, rather than just being presented with high level arguments they are intended to imitate. This is the most significant difference between mathematics as taught rigorously, and mathematics as taught in a physics degree (for instance). As discussed above there are also cases early on where students are expected to infer general principles from examples. There are only a few such cases though, and the reasoning each general principle encompasses can be characterized simply and precisely.

As mentioned in §1, there are various places where this account could be expanded on. This could be a task for further investigations on the subject. For instance, one might seek a better understanding of exactly how students “adjust their instincts” to recognize a wider variety of inferences as valid, having seen particular cases to be provable in greater detail; or how, and with what degree of success, they “infer general principles from examples” in the early stages. We leave such investigations for the future.

§4. The concept of rigour. The key feature of rigorous mathematics, on the view put forward here, is that there is a range of levels of detail that rigorous mathematics can take place at, where nontrivial inferences at a more compressed level can necessarily be proved at an appreciably greater level of detail. This is implicit in the process of learning rigorous mathematics described in §3, where each time a student is trying to master a new level of greater compression, it is constitutive of inferences being at that level of compression that they be provable at a previous level of detail, a level the student is already comfortable at—so that there are no leaps in the process of learning rigorous mathematics where a student is unable to check inferences for themselves (apart from when grasping certain basic principles). Indeed if inferences at the level of greater compression didn't need to be provable in more detail, then “checking” them by seeing which inferences can be justified with a proof would be a mistake. As discussed above, the basic axioms of set theory are exempted (whatever exactly they are taken to be), and are intended to be accepted by students without argument, though possibly with the assurance that they are “obvious”.

The requirement that nontrivial inferences be provable at a greater level of detail is also key to the mechanism for resolving disputes in mathematics discussed in §5. For the purpose of resolving disputes, this requirement is of less practical importance at levels of very great detail, as mathematicians may agree immediately about sufficiently basic inferences; but equally, for these simpler inferences it is generally more obvious that they can be proved in greater detail, and how such a proof would go, so the requirement is no unnecessary burden. Again, the basic axioms of set theory are exempted from

the requirement of being justifiable in greater detail: their correctness is not up for debate (within mathematics) and is supposed as a precondition for the mathematical enterprise to get going.

This requirement that inferences be provable in more detail is why the examples discussed in §2 were not rigorous: the manipulations of infinitesimals and infinite series discussed could not be justified in greater detail, and could not be regarded as basic rigorous rules in themselves since it was not clear how to demarcate what reasoning was acceptable.¹¹ Similarly brute intuitions such as for the truth of the Jordan curve theorem are not in themselves suitable in proofs unless they can be backed up with more detailed arguments.

The necessity that inferences be provable in greater detail applies to levels of detail that students (and mathematicians) reach after progressing onwards from the basic level at which the subject is first taught—assumed above to be the week 2 level of detail. However not all students do manage to directly grasp how proof works at this level of detail, and some need more explicit demonstration of the rules that proof is implicitly following. There are courses and textbooks which provide this, such as Velleman (2006), which teaches how proof works essentially by teaching how to prove statements using the natural deduction rules: here the premises being used are explicitly tracked and calculated with, according to the rules governing the various bits of logical vocabulary. We can call this most basic, most explicit level of detail the “intro to proof level of detail”. Students can use this as a stepping stone to gain comprehension of how basic arguments at the week 2 level of detail work, and implicit in this (as above with later levels of detail) is that inferences at the week 2 level of detail be provable at the intro to proof level of detail—otherwise gaining a grasp of how arguments work at the intro to proof level would be misleading as to what is going on at the week 2 level. Thus we can extend the above argument that inferences need to be provable in greater detail all the way down until we reach the intro to proof level of detail, where arguments explicitly use the natural deduction rules.

A minor caveat to this ability to prove in greater detail is that some arguments may require rephrasing when proving in more explicit formal terms—for instance one would often justify the Euclidean algorithm as an informal recursive process of repeated division with remainder, described for integers $a, b > 0$ perhaps by saying:

$$\begin{aligned} &\text{Write } a = q_0b + r_0 \text{ with } 0 \leq r_0 < b, \\ &\text{write } b = q_1r_0 + r_1 \text{ with } 0 \leq r_1 < r_0, \\ &\text{write } r_0 = q_2r_1 + r_2 \text{ with } 0 \leq r_2 < r_1, \\ &\text{and so on, until we reach } r_n = 0 \end{aligned}$$

which if proved more explicitly would be transformed into some sort of formal recursive definition (of the form justified by the axiom of dependent choice—though there is no choice here—as discussed in §3). In these kinds of cases one is replacing a part of a proof with a similar more detailed version, rather than literally filling in inferences

¹¹ Reasoning with infinitesimals was put on a firm footing by Robinson (1996) in terms of the logical concept of nonstandard model, and Nelson (1977) showed how to give a rigorous axiomatization for the approach. One can also use a paraconsistent approach to embrace the contradictory nature of infinitesimals, as seen in Brown & Priest (2004) and Sweeney (2014).

in greater detail. However these kinds of cases are the exception rather than the rule, and do not make any essential difference to any of the discussion below, so we will generally include them under “proving in greater detail” (speaking a little loosely).

Now although the account above has focused on the ability to prove inferences in greater detail, it should be emphasized that one does not generally have to see how to prove an inference in greater detail to understand it, or accept it. For instance, to someone familiar with the notion of homeomorphism it probably feels obvious that the sphere $\{x \in \mathbb{R}^3 : \|x\| = 1\}$ is homeomorphic to the cube $\{x \in \mathbb{R}^3 : \max_i |x_i| = 1\}$, but sitting down and trying to write out a proof of this could well take a while. Such examples are not limited to topology. In logic—a subject where perhaps one would expect “intuition” would play less of a role—it might well feel obvious that substituting term t for variable x in a formula ϕ , when x is not free in ϕ , will just return ϕ , but again proving this in detail would take a bit of work (though probably less insight than the previous example). These kinds of higher level judgments about inferences—without a proof in mind—are an essential part of mathematics. Nonetheless it is important in rigorous mathematics that the option of proving inferences in more detail is there, to aid in gaining a firm grasp of any new concepts, and in guiding and sharpening one’s judgment in any difficult cases—not all homeomorphic spaces are as obviously homeomorphic as the two above.

Also, when checking a proof one does not always necessarily actually check that every inference in is valid. Indeed if an unsurprising claim in a proof is supported by reasoning that looks like the right kind of thing, and the right amount of effort, then an experienced mathematician may pass on without checking every single detail. This is seen in interviews conducted by Andersen (2017) with mathematics referees, and also mentioned by Thurston (1997, p. 32). It appears this can be a fairly good guide to the overall correctness of results, though numerous commentators have remarked on the unreliability of the mathematical literature and the pervasiveness of errors in proofs, to which this manner of refereeing may be a contributing factor (Jaffe & Quinn, 1993; Thurston, 1994, p. 33; Nathanson, 2008; Gracar, 2013, pp. 421–422).

§5. Disagreements about validity. It is traditional when studying deduction to think in terms of a single sharply defined notion of validity, that every inference either has or lacks. I think the above analysis of rigour rather tells against this conception.

Indeed having picked out various levels of detail that mathematical inferences can take place at, we can introduce a cumulative hierarchy of validity predicates—“valid at the week 2 level of detail”, “valid at the term 2 level of detail”, and so on, where for instance “valid at the term 2 level of detail” means an inference either at the term 2 level of detail, or at a level of greater detail. One can keep ascending in this way, defining validity predicates which allow more and more compressed inferences, inferences that are increasingly challenging for even an experienced mathematician to follow. At some point one will reach inferences compressed enough that they are well beyond the bounds of what mathematicians consider to be valid. However there appears to be no natural place on the continuum of levels of detail to draw a line, and say *this* is the limit of validity: that inferences at least that detailed are valid, while those less detailed are not. Asserting that there is a precise such limit seems to just be philosophical dogma, unsupported by the facts of the practice. Suppose for instance that a preprint of an article is uploaded to the arXiv, and read by two mathematicians experienced in the

field—one of whom concludes that it is perfectly rigorous, the other that the proof of a certain lemma is too sketchy and incomplete. Who are we to announce—without examining the preprint—that one is definitely right, and the other definitely wrong? I think it is more plausible to say that judgments of validity can—to some extent—be both vague and subjective.

Some philosophers are aghast at this suggestion, it is worth explaining why, given the account of rigour above, it is not as damaging a claim as might be thought. In as much as there is disagreement about the validity of a certain inference, the framework of rigour provides a mechanism for resolving it.

That mathematicians are good at resolving disagreements about validity—that if a mathematician believes a proof is valid, they are generally able to convince others of this, or become convinced themselves of a flaw in it—has been noted by various authors as a fact that deserves explanation, and is one of the main motivations of Azzouni's derivation indicator view of proof is to attempt to provide this (Azzouni, 2004, pp. 83–84; Antonutti Marfori, 2010, pp. 267, 270–271). The idea behind Azzouni's account is that the informal proofs mathematicians write serve to indicate formal derivations. This has met with criticism, with Tanswell (2015) pointing out that attempted proofs may have many different attempted formalizations, which poses a problem since Azzouni wants to characterize validity of the informal proof in terms of success of the indicated formal derivation. Additionally as seen in §3, mathematics generally proceeds at a much greater level of compression than is found in formal proofs, and (as mentioned in §1) most mathematicians have no experience of or interest in the activity of formalization, so Azzouni's account is rather far removed from how mathematicians typically engage with proofs in practice.

Instead of hoping for an explanation in terms of completely formal proofs, it is more promising to look to the process of proving in more detail itself. Indeed if a mathematician ever puts forward a purported proof in which an inference is not convincing, then a more detailed justification for that inference can be requested—an appreciable step up to a level of greater detail, perhaps from the research article level of detail (terse) to the graduate level of detail (terse), as discussed in §3.¹² At this level of greater detail inferences are more transparent and judgments of validity are more reliable, and this may already serve to resolve the controversy—with the new argument being acceptable, or an obvious error in it discovered. If not, and an inference in this more detailed argument is still controversial, a more detailed justification can be asked of it in turn, taking us to a still greater level of detail at which errors will be even more obvious—perhaps we now reach the third year level of detail. In principle this process will terminate when one reaches the level of complete formalization, though in practice if both sides are of sound mind and proceeding in good faith then the controversy will be resolved well before that.¹³

Thus if one believes a proof to be rigorous, in as much as this belief is correct one can always (in principle, and usually in practice) fill in the details of any inferences that are felt to be sketchy, to increase the level of detail to one which is found acceptable.

¹² Though this terminology of levels of detail is new, the process is not.

¹³ The resolution of disagreement in this manner is an example of the idea from argumentation theory that a debate goes down to the level of detail that will satisfy both parties (using the apt words of an anonymous referee). For more on argumentation theory and mathematics, see Aberdein & Dove (2013).

In the imagined case considered initially of two mathematicians disagreeing about an arXiv preprint—with one finding that a particular lemma was argued for too briefly—the dispute would normally be resolvable in this way, bringing the proof into a form acceptable to everyone.

Thus though in my view there may be some subjectivity and vagueness to where exactly the limits of rigour are drawn, much more significant is the strong form of objectivity afforded by rigour, in which there are always robust reasons available to win over an objector—provided one's assessment of a proof as valid is justified.

In fact this dispute resolution mechanism is essentially that described in the Princeton Companion to Mathematics:

[T]he fact that arguments can in principle be formalized provides a very valuable underpinning for the edifice of mathematics, because it gives a way of resolving disputes. If a mathematician produces an argument that is strangely unconvincing, then the best way to see whether it is correct is to ask him or her to explain it more formally and in greater detail. This will usually either expose a mistake or make it clearer why the argument works. (Gowers, Barrow-Green, & Leader, 2008, p. 74)

§3 and §4 can be seen as a clarification and elaboration of the process outlined in this quote. As an explanation of how disputes in mathematics can be resolved, this seems to be both more straightforward and more plausible than Azzouni's account, and better grounded in mathematical practice. However there may well be more to be said about the reality of disputes over validity in mathematics, and how well this simple account fits it.

§6. Naive set theory? The final topic of the paper is how this account of rigour impacts on the question of formalizability. Before the discussion of this in §7, we take a brief detour from the main course of the paper, to address a sceptical view about the basic principles used in mathematics: the idea that mathematicians should be viewed not as working in a system like ZFC(U), but in naive set theory, with its axiom scheme of unrestricted comprehension. This has been suggested by Leitgeb (2009), and some mathematicians have made similar claims about their own understanding (Jones, 1998, p. 205; Aluffi, 2009, p. 1). Indeed “Naive set theory” is actually the title of a set theory textbook by Halmos (2011). If mathematicians are best regarded as using unrestricted comprehension, this would make the claim that mathematical proofs can be formalized trivial, since set theory with unrestricted comprehension is inconsistent and so any argument can be immediately formalized in it (with every inference justified by a set theoretic paradox).

The key step when considering this possibility is to distinguish different senses of the term “naive set theory”. Certainly most mathematicians do not know what the axioms of ZFC are, but they do have a solid grasp of how to legitimately form new sets: by taking unions, subsets, power sets, Cartesian products, function sets, equivalence classes, and so on. This understanding may be “naive” in the sense that it is not accompanied by explicit awareness of how these operations are justified in terms of the basic axioms—but that is totally different to “naive” set theory in the logicians' sense,

in which the central principle is that of unrestricted comprehension, the scheme

$$\exists y (x \in y \Leftrightarrow \phi(x))$$

for all formulae ϕ in which y does not occur free.¹⁴ Indeed the above set forming operations are not generally justified any more directly by unrestricted comprehension than by ZFC; to form Cartesian products for instance, one still needs to go through the rigmarole of defining what ordered pairs are and what a family of objects is, and the availability of unrestricted comprehension does not significantly simplify this. Moreover, there is no evidence of mathematicians making essential appeals to unrestricted comprehension, and this being accepted as valid. In the normal course of mathematics, all classes one would like to be sets are easily seen to be set sized using the standard set forming operations. In category theory, where size issues are encountered, the axiom of universes was introduced specifically so that they could be dealt with in a rigorous way. There are occasional instances where classes are manipulated as though they were sets, for instance in the definition of the Grothendieck group as a quotient of the set of isomorphism classes of finitely generated modules over a ring R ; but at least in this case there is nothing genuinely troubling going on since one can easily define a set of representatives of the isomorphism classes instead (the quotient modules $\frac{R^n}{M}$ of powers of R , with isomorphic quotients identified), or one can appeal to the axiom of universes.

Moreover there are genuine mathematical cases where the distinction between sets and classes is crucial, and the use of unrestricted comprehension would be disastrous. For instance the general adjoint functor states that if $G : D \rightarrow C$ is a functor with D complete and locally small, then G has a left adjoint iff it preserves all limits and satisfies the solution set condition (MacLane, 1998, p. 121). The solution set condition states that a set of morphisms with a certain property exists, and in this case the fact that this be a set rather than a class is key, as there is always a class of morphisms with the required property. In the presence of unrestricted comprehension, the general adjoint functor theorem would become the claim that any limit preserving functor whose domain is complete and locally small has a left adjoint, which is false in general. The issue of which functors have adjoints is not some category theoretic curiosity—it is an important issue in various areas of mathematics including algebra and topology.

Thus the way mathematicians form sets may be naive in the sense that it need not be founded in explicit knowledge of the basic principles, but there is no indication that it is naive in the sense of relying on unrestricted comprehension. If it did, signs of this ought not to be too hard to find.

§7. Formalizability. Now to the question of formalizability itself. As advertised previously, it will be argued here that valid rigorous proofs are formalizable, in principle, though what this means requires clarification. The prospects for feasible formalization will also be touched on.

¹⁴ The set theory in Halmos's textbook is naive in an even weaker sense. Halmos does in fact state the usual basic axioms of set theory (with no mention of unrestricted comprehension), and he uses them to derive various set theoretic operations and facts, but he says the approach is naive in that "the language and notation are those of ordinary informal (but formalizable) mathematics".

For the purposes of this section we introduce the concept of “deductive grounding” between levels of detail. If L, L' are levels of detail, then we say that L' is *deductively grounded* in L if every inference valid at level of detail L' is provable at level of detail L .¹⁵ This concept of grounding is justificatory rather than metaphysical, and does not share all the properties of standard notions of metaphysical grounding (for instance it is reflexive). This concept of grounding uses a notion of provability, and the kind of modality employed here needs to be spelt out before we know what deductive grounding amounts to. The notion of provability we will use is one of “in principle” provability, abstracting from limitations in terms of time or other resources (and thus perhaps abstracting away from the limitations of our own physical universe). One could otherwise use a notion of what is actually feasibly provable, given the physical and biological constraints on us, and obtain thereby a notion of “feasible deductive grounding”.

We can relate this notion of deductive grounding to a notion of “in principle formalizability”. Indeed, we can fix some standard system of first order natural deduction (for instance that of Prawitz, 1965, chap. I), and we take a formal proof to be a complete derivation in this system in the language of set theory with all premises amongst the axioms of ZFCU. We then define the “formal level of detail” to consist just of these formal proofs. We characterize a proof as being *formalizable in principle* if every inference in it is provable at the formal level of detail. Thus the claim that every proof at a level of detail L is formalizable in principle is just the claim that L is deductively grounded in the formal level. This is a weak notion of formalizability in principle, and for instance if the Riemann hypothesis is a theorem of ZFC then the one line proof of the Riemann hypothesis from no premises is formalizable in principle by this definition. Nonetheless it is one available sense of formalizability in principle; the question of what we do and should mean by formalizability will be returned to later in this section.

The key property of this notion of in principle provability is that it satisfies a version of the converse Buridan formula. In general the converse Buridan formula is (the scheme of formulae) of the form

$$\forall x \diamond(\phi(x)) \Rightarrow \diamond(\forall x \phi(x)).$$

It is of the same general form as the converse Barcan formula, though with an existential instead of universal quantifier (Konyndyk, 1986, p. 94). This converse Buridan formula is not typically valid. For instance if I have a well stocked fridge, it may be the case that for every item in the fridge I can have that item as part of my dinner, but that it is impossible for me to have every item in the fridge as part of my dinner. Nonetheless in cases where one abstracts from resource constraints it can be valid. In particular if we have a finite set S of inferences, and write Prov to indicate that we have obtained a proof of $s \in S$, then we do have that

$$\forall s \in S \diamond(\text{Prov}(s)) \Rightarrow \diamond(\forall s \in S \text{Prov}(s))$$

¹⁵ As was noted in §4, some inferences at level L' may be part of an informal section of a proof which as a whole needs to be replaced by a more detailed and formal version at level L ; this was illustrated with the example of the Euclidean algorithm. This caveat makes no real difference to what follows.

since if each element of S is provable, then—given sufficient time—it will be possible to obtain coeval proofs of every element of S .

It follows that if level of detail L' is deductively grounded in level of detail L , and we have a proof p of result s valid at level of detail L' , then one can in principle obtain a proof of s at level of detail L . Indeed every inference in p is provable at level of detail L , so that (as just discussed) it is possible to obtain a simultaneous proof of every inference in p ; concatenating these then gives a proof of s at level of detail L , as claimed.

Thus we can obtain that the notion of deductive grounding is transitive. Indeed suppose we have levels of detail L, L', L'' with L' deductively grounded in L and L'' deductively grounded in L' . Let s be an inference valid at level of detail L'' . Then by the definition of deductive grounding, we can in principle obtain a proof of s at level of detail L' ; but then as just discussed, given such a proof one can in principle obtain a proof of s at level of detail L . Thus s is indeed provable at level of detail L .¹⁶

We will now argue that for the levels of detail that a student moves through on their way to mastering research level mathematics, each more compressed level is deductively grounded in its more detailed predecessors. Recall as discussed in §3 and §4 that it is crucial for learning rigorous mathematics that at each stage, one has the option of proving the more compressed inferences one is encountering at a greater level of detail that one has already mastered.

We will consider a student gradually moving through an education in mathematics, where at each time t there is the level of detail m_t of mathematics that they have mastered so far, and the level of detail l_t that they are learning at that point. We will assume that the collection of times t making up this period of education forms a complete totally ordered set, which we denote by I . I may or may not have endpoints—an initial point, a final point, or both. If L and L' are levels of detail we will write $L \leq L'$ to denote that L' consists of mathematical inferences at least as compressed as those at level L . We will assume that the function $t \mapsto m_t$ is monotonic, i.e., that if $t' \geq t$ then $m_{t'} \geq m_t$. For there to be no magical jumps in this process of learning, it needs to be the case that for all times t not initial in I , there is some $t' < t$ such that $m_t \leq l_{t'}$, so that the level of detail mastered at time t was actually learnt at some previous point in time. We also need that for all times t not final in I , there is some $t' > t$ such that $m_{t'} \leq l_t$, so that there is no magical jump after time t where at all subsequent times, a level of detail has been mastered that is greater than that which was being learnt at time t . Finally, as discussed above, it is constitutive of learning rigorous mathematics that at each stage, the level of detail one is learning is deductively grounded in a level of detail one has already mastered, i.e., that l_t is deductively grounded in m_t .

Now for the argument. For each t not initial in I there is some $t' < t$ such that $m_t \leq l_{t'}$, but $l_{t'}$ is deductively grounded in $m_{t'}$, so that by transitivity m_t is also deductively grounded in $m_{t'}$ (call this backwards grounding). Also, for each t not final in I , there is $t' > t$ such that $m_{t'} \leq l_t$, so that $m_{t'}$ is deductively grounded in m_t (call this forwards grounding). Now suppose for contradiction that there is $s > t$ such that m_s is not deductively grounded in m_t . We let r be the infimum of

$$\{s \mid m_s \text{ not deductively grounded in } m_t\}$$

¹⁶ This argument could be carried out more formally, and implicitly uses the principle $\diamond\diamond\phi \Rightarrow \diamond\phi$ of S4, which holds for the kind of metaphysical possibility being employed.

(a set which is bounded below by t). By forwards grounding, there is $t_1 > t$ such that m_{t_1} is deductively grounded in m_t , so that by monotonicity of m , we have that $r \geq t_1 > t$. Thus r is not initial in I , and so by backwards grounding for r , there is $t_2 < r$ such that m_r is deductively grounded in m_{t_2} . But by the definition of r , we must have that m_{t_2} is deductively grounded in m_t , and thus by transitivity that m_r is deductively grounded in m_t . But then forwards grounding for r gives that there is $t_3 > r$ such that m_{t_3} is grounded in m_r , and thus in m_t ; then by monotonicity of m , if $r \leq t' \leq t_3$ then $m_{t'}$ is deductively grounded in m_t , so that t_3 is actually a lower bound for

$$\{s \mid m_s \text{ not deductively grounded in } m_t\},$$

so that r is not the infimum of this set, giving the required contradiction. Thus from the assumptions we have made, it follows that if $s > t$, then m_s is deductively grounded in m_t .

Thus we obtain that all levels of detail that can be mastered rigorously, from a process starting at the week 2 level of detail, are deductively grounded in the week 2 level of detail. Then it is clear by inspection that arguments at the week 2 level of detail can be written out as formal derivations, as was noted in §3, and their premises are all amongst the axioms of ZFCU, so that the week 2 level of detail is deductively grounded in the formal level. Otherwise, as discussed in §4, one can define an “intro to proof level of detail”, in which the natural deduction rules are explicitly used, and which some students use as a stepping stone to master the week 2 level of detail; it follows (as above) that the week 2 level of detail is deductively grounded in the intro to proof level of detail, whose proofs are trivially formalizable—they are already natural deductions of a slightly nonstandard kind. Either way we obtain that the week 2 level of detail is deductively grounded in the formal level, and thus that all the levels of detail that can be mastered rigorously are deductively grounded in the formal level. In other words, all rigorous mathematics is, in principle, formalizable. This is not a mere empirical fact obtained by looking at examples of mathematical arguments (except perhaps for the step from the intro to proof or week 2 level to the formal level), but a direct consequence of the norm of rigour in mathematics as enunciated here in terms of levels of detail.

This account answers a potential worry about formalizability: if rigour requires proofs to be formalizable, how are mathematicians so good at judging this? The answer is that mathematicians are not directly judging formalizability, but are instead judging the rigour of inferences (as discussed in §3 and §4), and that formalizability is obtained as a consequence of rigour.

In the literature the main worry about this kind of in principle formalizability is raised by Rav (1999), reiterated by Weir (2016). Rav considers a situation where one has an inference $A \rightarrow B$ in a proof, and after some thought fills this in with intermediate inferences to obtain $A \rightarrow A_1, A_1 \rightarrow A_2, \dots, A_n \rightarrow B$. Perhaps one is then questioned by a student or nonspecialist as to why A_1 follows from A , and comes up with a new interpolation $A \rightarrow A', A' \rightarrow A_1$. Rav claims we can give no “theoretical” reason why this process of adding interpolations ought to ever terminate (Rav, 1999, pp. 14–15).

The basic problem with this description is the lack of any attempt to characterize the form the interpolating inferences must take. One cannot just write in any intermediate inferences that suit one’s fancy, whether justifying an inference to oneself or to a sceptic. For instance suppose we are trying to argue that the fact there are infinitely many primes (IP) is a consequence of the fundamental theorem of arithmetic (FTA),

and we write RH for the Riemann hypothesis and ST for Szemerédi's theorem. It is clearly nonsense to try to justify $\text{FTA} \rightarrow \text{IP}$ by filling in intermediate inferences of the form

$$\text{FTA} \rightarrow \text{RH}, \text{RH} \rightarrow \text{ST}, \text{ST} \rightarrow \text{RH}, \text{RH} \rightarrow \text{IP}.$$

These may be valid as material implications (in the presence of the axioms of ZFC), but they are totally useless as intermediate inferences for us; and to explain why they are the wrong kind of intermediate inferences, we have to start putting inferences on some sort of scale of plausibility, or simplicity, or evidentness, and require that adding intermediate inferences take us in the increasing direction on this scale. This is the first step towards thinking of mathematics in terms of levels of detail, as in §3 and §4, and to the requirement—implicit in Rav's description—that nontrivial inferences be provable in greater detail. Then we are on the road to the argument for formalizability just given.

This argument concerned in principle formalizability, but it is also possible to say a little from this perspective about how practical formalization might be. Indeed we can use the above framework to address a weaker version of Rav's worry put forward by Pelc (2009), who accepts that the process of filling in with intermediate inferences will necessarily terminate, but questions whether we have any reason to believe this process will result in a formal proof of feasible length (given some initial proof of reasonable size). Pelc defines a vast number M in terms of various fundamental constants, large enough so that there is no possibility of us ever (in practice) constructing a formal proof of this length. His number M is at least the number of particles in the universe divided by the Planck time (in seconds), and thus at least 10^{120} on standard estimates; so a special case of Pelc's worry is whether when formalizing a proof of reasonable length, we have any reason to believe the resulting formal proof will be less than 10^{120} symbols.

The framework of levels of detail developed above can also be useful when addressing this kind of worry. Indeed one can get from the week 2 level of detail to the research article level of detail (terse) in a small number of steps up: from the week 2 level of detail to the year 2 level of detail, from the year 2 level of detail to the graduate level of detail (explicit), and from the graduate level of detail (explicit) to the research article level of detail (explicit). For each of these steps one can estimate by considering examples what kind of factor increase in length one generally obtains, when writing out an inference from the more compressed level at the more detailed level. Though it would have to be confirmed by more careful investigation, my belief from considering examples is that this factor is not generally too large, with a factor of less than 5 being common, a factor of 10 being fairly rare and a factor of 20 being very rare (as a proportion of inferences). This is supported by the example of Perelman's proof of the Poincaré conjecture discussed in §3, where an extremely high level argument was written out at around the graduate level of detail with a factor of increase in length of 8. Thus though one may obtain exponential growth in proof length as one fills in details in a proof, to bring it down to the week 2 level of detail, the base and exponent are both fairly small: the former being the factor of increase with each step to greater detail, the latter being the number of such steps (4 in the case above). Even with a factor of 20 at each stage, and another factor of 20 (again a crude upper bound) to reach the formal level from the week 2 level, this gives us an overall factor of increased length of at most $20^5 = 3,200,000$ to formalize an argument at the research article level of

detail (terse). This is vastly less than the factor Pelc is concerned might be necessary. Though this could undoubtedly result in very unwieldy proofs (if this crude upper bound was attained), they would still be within the bounds of feasibility, as usually conceived—requiring perhaps a few gigabytes or tens of gigabytes of space to store on a hard drive.¹⁷

Even if one can fill out the above sketched argument to make a convincing case that all rigorous proofs of reasonable length will be feasibly formalizable, this does not meet one major kind of objection to formalizability: namely, that the process of formalization so dramatically changes a proof that the formal proof that results cannot rightly be regarded as the “same proof” as the original, and thus should not be regarded as a formalization of it. Larvor writes of the “violence or essential loss” that can result from formalization (Larvor, 2012, p. 717). A related question concerns what practical relevance formalization has, or could have, to mathematics; it played only an indirect role in the account of rigour from §3 and §4, and turned out to be inessential to the process of resolving disputes in mathematics described in §5. These questions will have to await possible consideration in future work.

§8. Final thoughts. To conclude, there are a few loose ends concerning rigorous proof to discuss.

First, to say a bit about a topic that has been neglected here: understanding proofs. Indeed instead of just checking each line of a proof, one generally also wants to understand the proof “as a whole”. This can be phenomenologically quite distinct from mere confidence that each line follows from the previous ones, as Tieszen (1992, pp. 58–59) emphasizes. Though it is tempting to think about understanding in terms of the distinctive subjective experience of grasping a proof, this characteristic sensation may not always be attainable—for instance it may be hard to gain the feeling of an immediate grasp of a proof taking longer than a page or so, or of a proof which relies on substantial previous results. An attractive alternative is to think of understanding as an ability, as advocated by Avigad (2011): for instance understanding a proof may mean that one can recreate it oneself, convey its ideas informally to another mathematician, use its ideas or techniques to prove similar results, generalize it, suggest how it could have been discovered, and so on. In many of these capacities that displays understanding of a proof, the ability to create or recognize valid proofs—the ability discussed throughout this paper—is key: recreating the proof means writing out a similar valid proof of the same result, generalizing it means writing out a similar valid proof of a more general result, and so on. Arguably, to convey the ideas behind the proof to another mathematician means to use words, gestures, diagrams and so on to

¹⁷ This sketched argument is intended to concern formalizability in a standard system of natural deduction, such as that of Prawitz (1965, chap. I), relative to the axioms of ZFCU, which I take to be the axioms underlying the usual practice of proof in modern mathematics. The question of proof speed up at a more compressed level of informal mathematics should be distinguished from the question of proof speed up in a stronger formal system (there is no reason to believe that any system of rules could be written down that govern informal mathematical inferences, as discussed in §1), though it is possible that if one required formalization in a weaker system, the lengths of the formal proofs obtained in practice would increase.

equip the recipient with the means to recreate for themselves a similar valid proof of the result in question. As Thurston (1994, pp. 31–32) emphasizes, it can often be much easier to convey mathematical ideas by informal communication than by embedding them in and then excavating them from rigorous proofs.

A second point is that in practice in mathematics it is not demanded that every inference in a proof actually be valid. Total validity is intended when many philosophers speak of proof (as has been the attitude in this paper), but in reality if an argument published in a mathematics journal contains a number of typos and minor logical errors, it may still be regarded as a perfectly acceptable proof. We could call the mathematicians' use of the word proof "proof in the weak sense", to distinguish it from the philosophical use of the term. The key feature I think for proof in this weak sense is I think that the proportion of valid inferences is high, or very high, and that those inferences which are invalid are each fixable relatively easily. Thus for instance the classification of finite simple groups may well be a proof in this weak sense, even though it is so enormously long that it is practically certain that it contains errors. This is more or less the view taken by Aschbacher (2005). A proof of a result in this weak sense still establishes that its conclusion is a logical consequence of the basic principles used—since in principle one could convert it into a proof in the strict sense, in which case the result would be a logical consequence of the relevant basic principles (as seen in §7), and whether or not the conclusion is a logical consequence of the basic principles is independent of whether any such conversion is actually carried out.

Finally, a note on the epistemology of mathematics. In §3 and §4, an attempt was made to understand how rigour is judged in mathematics by thinking about how rigorous proof is learnt—by a gradual ascent up levels of greater and greater compression, at each stage being able to tutor one's judgment of which inferences are valid by checking if they can be proved in greater detail. There is undoubtedly more that could be said about this process, but at any rate it is only part of a full epistemology of mathematics. Indeed it was argued in §7 that a valid proof shows its conclusion to be a logical consequence of the basic principles used; but this in itself does not imply that the conclusion is *true*. That would require further arguments, for instance arguments that the basic principles themselves are true. Thus the epistemology sketched here is one component of a full epistemology of proof, which would need to be supplemented by an epistemology of the basic principles themselves. Whether the basic principles generally used in mathematics—those of set theory—are true, and how we could know this, are exactly the kinds of epistemological questions that philosophy of mathematics has long wrestled with, and which some advocates of the shift to focusing on actual mathematics like to disparage, or describe as irrelevant to mathematics itself (for instance Rav, 1999, Goethe & Friend, 2010 and De Toffoli & Giardino, 2016). It is of note that thinking about mathematical practice and the epistemology of proof naturally leads us back to exactly this standard epistemological question.

Acknowledgements I am very grateful to Leon Horsten, Catrin Campbell-Moore, Marcus Giaquinto, Yacin Hamami and various anonymous referees for detailed and insightful comments on this paper, which improved it significantly, as well as to members of the audiences at talks given on the subject in Bristol, Mussomeli, London (Ontario) and Brussels for valuable comments and questions. I am also grateful to the AHRC for the grant AH/L503939/1 which supported this work.

BIBLIOGRAPHY

- Aberdein, A., & Dove, I. J., editors (2013). *The Argument of Mathematics. Logic, Epistemology, and the Unity of Science*. Netherlands: Springer.
- Aluffi, P. (2009). *Algebra: Chapter 0*. Providence, RI: American Mathematical Society.
- Andersen, L. E. (2017). On the nature and role of peer review in mathematics. *Accountability in Research*, **24**(3), 177–192.
- Antonutti Marfori, M. (2010). Informal proofs and mathematical rigour. *Studia Logica: An International Journal for Symbolic Logic*, **96**(2), 261–272.
- Artin, M. (1991). *Algebra* (first edition). Englewood Cliffs, NJ: Pearson.
- Aschbacher, M. (2005). Highly complex proofs and implications of such proofs. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **363**(1835), 2401–2406.
- Avigad, J. (2011). Understanding proofs. In Mancosu, P., editor. *The Philosophy of Mathematical Practice*. Oxford: Oxford University Press, pp. 317–353.
- Awodey, S. (2010). *Category Theory* (second edition). Oxford: Oxford University Press.
- Axler, S. (1997). *Linear Algebra Done Right* (second edition). New York: Springer.
- Azzouni, J. (2004). The derivation-indicator view of mathematical practice. *Philosophia Mathematica*, **12**(2), 81–106.
- Bak, J., & Newman, D. J. (2010). *Complex Analysis* (third edition). Undergraduate Texts in Mathematics. New York: Springer-Verlag.
- Banach, S. (1987). *Theory of Linear Operations*. Netherlands: Elsevier.
- Bartle, R. G., & Sherbert, D. R. (2000). *Introduction to Real Analysis*. Canada: John Wiley & Sons Limited.
- Berkeley, G. (1999). The analyst. In Ewald, W. B., editor. *From Kant to Hilbert Volume 1: A Source Book in the Foundations of Mathematics*. Oxford: Oxford University Press, pp. 60–92.
- Brown, B., & Priest, G. (2004). Chunk and permeate, a paraconsistent inference strategy. Part I: The infinitesimal calculus. *Journal of Philosophical Logic*, **33**(4), 379–388.
- Burgess, J. P. (2015). *Rigor and Structure*. Oxford: Oxford University Press.
- Celluci, C. (2009). Why proof? What is a proof? In Lupacchini, R., and Corsi, G., editors. *Deduction, Computation, Experiment*. Milano, Italy: Springer-Verlag GmbH, pp. 1–27.
- Conway, J. B. (1978). *Functions of One Complex Variable I* (second edition). Graduate Texts in Mathematics. New York: Springer-Verlag.
- Cori, R., & Lascar, D. (2000). *Mathematical Logic Part 1: Propositional Calculus, Boolean Algebras & Predicate Calculus: A Course with Exercises*. Oxford: Oxford University Press.
- Coulston, C., editor (1970). *Dictionary of Scientific Biography Volume II: Hans Berger—Christoph Buys Ballot*. New York: Charles Scribner's Sons.
- De Toffoli, S., & Giardino, V. (2016). Envisioning transformations—the practice of topology. In Larvor, B., editor. *Mathematical Cultures: The London Meetings 2012–2014* (first edition). New York: Birkhäuser, pp. 25–50.
- Detlefsen, M. (2009). Proof: Its nature and significance. In Gold, B., and Simons, R., editors. *Proof and Other Dilemmas: Mathematics and Philosophy* (UK edition). Cambridge: Cambridge University Press, pp. 3–32.

- Eisenbud, D. (1995). *Commutative Algebra: With a View Toward Algebraic Geometry*. Berlin, Germany: Springer Science & Business Media.
- Eisenbud, D., & Harris, J. (2000). *The Geometry of Schemes*. Graduate Texts in Mathematics. New York: Springer-Verlag.
- Epp, S. S. (2003). The role of logic in teaching proof. *The American Mathematical Monthly*, **110**(10), 886–899.
- Erdős, P. (1947). Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, **53**(4), 292–294.
- Fremlin, D. (2010). *Measure Theory*, Vol. 1. Torres Fremlin.
- Fulton, W., & Harris, J. (1991). *Representation Theory: A First Course*. Berlin, Germany: Springer-Verlag.
- Gentzen, G. (1969). *The Collected Papers of Gerhard Gentzen*. Amsterdam, Netherlands: North-Holland Publishing Company.
- Goethe, N. B., & Friend, M. (2010). Confronting ideals of proof with the ways of proving of the research mathematician. *Studia Logica*, **96**(2), 273–288.
- Gowers, T., Barrow-Green, J., & Leader, I., editors (2008). *The Princeton Companion to Mathematics*. Princeton, NJ: Princeton University Press.
- Grcar, J. (2013). Errors and corrections in mathematics literature. *Notices of the American Mathematical Society*, **60**(4), 418–425.
- Hales, T. C. (2007). Jordan's proof of the Jordan Curve theorem. *Studies in Logic, Grammar and Rhetoric*, **10**(23), 45–60.
- Halmos, P. R. (2011). *Naive Set Theory*. Mansfield, CT: Martino Fine Books.
- Hamami, Y. (2019). Mathematical rigor and proof. *The Review of Symbolic Logic*, 1–41. doi: [10.1017/S1755020319000443](https://doi.org/10.1017/S1755020319000443)
- Hartshorne, R. (1977). *Algebraic Geometry*. Graduate Texts in Mathematics. New York: Springer-Verlag.
- Hatcher, A. (2001). *Algebraic Topology by Allen Hatcher* (first edition). Cambridge: Cambridge University Press.
- Hirsch, M. W. (1976). *Differential Topology*. Graduate Texts in Mathematics. New York: Springer-Verlag.
- Jaffe, A., & Quinn, F. (1993). “Theoretical mathematics”: toward a cultural synthesis of mathematics and theoretical physics. *Bulletin of the American Mathematical Society*, **29**(1), 1–13.
- James, G., & Liebeck, M. (2001). *Representations and Characters of Groups* (second edition). Cambridge: Cambridge University Press.
- Jech, T. (2006). *Set Theory: The Third Millennium Edition* (revised and expanded third edition). Berlin, Germany: Springer.
- Jones, V. (1998). A credo of sorts. In Dales, H. G., and Oliveri, G., editors. *Truth in Mathematics*. Oxford: Clarendon Press, pp. 203–214.
- Jordan, C. (1887). *Cours d'analyse de l'école polytechnique*.
- Kline, M. (1990a). *Mathematical Thought from Ancient to Modern Times* (new edition), Vol. 1. New York: Oxford University Press.
- . (1990b). *Mathematical Thought from Ancient to Modern Times* (new edition), Vol. 3. New York: Oxford University Press.
- Konyndyk, K. (1986). *Introductory Modal Logic* (first edition). Notre Dame, IN: University of Notre Dame Press.
- Larvor, B. (2012). How to think about informal proofs. *Synthese*, **187**(2), 715–730.

- . (2019). From Euclidean geometry to knots and nets. *Synthese*, **196**(7), 2715–2736.
- Lee, J. (2012). *Introduction to Smooth Manifolds* (2012 edition). New York: Springer.
- Lee, J. M. (2000). *Introduction to Topological Manifolds*. Graduate Texts in Mathematics. New York: Springer-Verlag.
- Leitgeb, H. (2009). On formal and informal provability. In Bueno, O., and Linnebo, Ø., editors. *New Waves in Philosophy of Mathematics* (2009 edition). Basingstoke: Palgrave Macmillan, pp. 263–299.
- MacLane, S. (1998). *Categories for the Working Mathematician* (second edition). New York: Springer.
- Mancosu, P. (1989). The metaphysics of the calculus: A foundational debate in the Paris academy of sciences, 1700–1706. *Historia Mathematica*, **16**(3), 224–248.
- Morgan, J., & Tian, G. (2007). *Ricci Flow and the Poincaré Conjecture* (reprint edition). Providence, RI: American Mathematical Society.
- Munkres, J. (2000). *Topology* (second edition). Upper Saddle River, NJ: Pearson.
- Nathanson, M. (2008). Desperately seeking mathematical truth. *Notices of the American Mathematical Society*, **55**(7), 773.
- Nelson, E. (1977). Internal set theory: A new approach to nonstandard analysis. *Bulletin of the American Mathematical Society*, **83**(6), 1165–1198.
- Niven, I. M., Niven, A., & Zuckerman, H. S. (1991). *An Introduction to the Theory of Numbers* (fifth edition). New York: John Wiley & Sons.
- Pelc, A. (2009). Why do we believe theorems? *Philosophia Mathematica*, **17**(1), 84–94.
- Perelman, G. (2002). The entropy formula for the Ricci flow and its geometric applications. arXiv:math/0211159.
- . (2003a). Finite extinction time for the solutions to the Ricci flow on certain three-manifolds. arXiv:math/0307245.
- . (2003b). Ricci flow with surgery on three-manifolds. arXiv:math/0303109.
- Prawitz, D. (1965). *Natural Deduction: A Proof-Theoretical Study*. Mineola, NY: Dover Publications.
- Rav, Y. (1999). Why do we prove theorems? *Philosophia Mathematica*, **7**(1), 5–41.
- . (2007). A critique of a formalist-mechanist version of the justification of arguments in mathematicians' proof practices. *Philosophia Mathematica*, **15**(3), 291–320.
- Robinson, A. (1996). *Non-standard Analysis* (new edition). Princeton, NJ: Princeton University Press.
- Rudin, W. (1987). *Real and Complex Analysis*. New York: McGraw-Hill.
- Schwartz, J. (1954). The formula for change in variables in a multiple integral. *The American Mathematical Monthly*, **61**(2), 81–85.
- Silverman, J. H. (2012). *A Friendly Introduction to Number Theory* (fourth edition). Boston, MA: Pearson.
- Sweeney, D. J. (2014). Chunk and permeate: the infinitesimals of Isaac Newton. *History and Philosophy of Logic*, **35**(1), 1–23.
- Switzer, R. M. (2002). *Algebraic Topology—Homotopy and Homology*. Classics in Mathematics. Berlin, Heidelberg/Germany: Springer-Verlag.
- Szemerédi, E. (1975). On sets of integers containing no k elements in arithmetic progression. *Polska Akademia Nauk. Instytut Matematyczny Acta Arithmetica*, **27**, 199–245.

- Tanswell, F. (2015). A problem with the dependence of informal proofs on formal proofs. *Philosophia Mathematica*, **23**(3), 295–310.
- Tatton-Brown, O. (2019). Rigour and intuition. *Erkenntnis* (online only early view).
- Taylor, J., & Garnier, R. (2014). *Understanding Mathematical Proof* (first edition). Boca Raton, FL: Routledge.
- Thurston, W. P. (1994). On proof and progress in mathematics. *Bulletin of the American Mathematical Society*, **30**(2), 161–177.
- . (1997). *Three-Dimensional Geometry and Topology*, Vol. 1. Princeton, NJ: Princeton University Press.
- Tieszen, R. (1992). What is a proof? In Detlefsen, M., editor. *Proof, Logic and Formalization* (first edition). London: Routledge, pp. 57–76.
- Tragesser, R. (1992). Three insufficiently attended to aspects of most mathematical proofs: Phenomenological studies. In Detlefsen, M., editor. *Proof, Logic and Formalization* (first edition). London: Routledge, pp. 162–198.
- Veblen, O. (1905). Theory on plane curves in non-metrical analysis situs. *Transactions of the American Mathematical Society*, **6**(1), 83–98.
- Velleman, D. J. (2006). *How to Prove It: A Structured Approach* (second edition). Cambridge: Cambridge University Press.
- Walker, R. C. (1974). *The Stone-Čech Compactification*. New York: Springer-Verlag.
- Weir, A. (2016). Informal proof, formal proof, formalism. *The Review of Symbolic Logic*, **9**(1), 23–43.

DEPARTMENT OF PHILOSOPHY
UNIVERSITY OF BRISTOL
BRISTOL, UK
E-mail: otattonbrown@gmail.com