

GROUNDWORK FOR A FALLIBILIST ACCOUNT OF MATHEMATICS

BY SILVIA DE TOFFOLI

According to the received view, genuine mathematical justification derives from proofs. In this article, I challenge this view. First, I sketch a notion of proof that cannot be reduced to deduction from the axioms but rather is tailored to human agents. Secondly, I identify a tension between the received view and mathematical practice. In some cases, cognitively diligent, well-functioning mathematicians go wrong. In these cases, it is plausible to think that proof sets the bar for justification too high. I then propose a fallibilist account of mathematical justification. I show that the main function of mathematical justification is to guarantee that the mathematical community can correct the errors that inevitably arise from our fallible practices.

Keywords: mathematical justification, proof, a priori, fallibility, basing relation, reliability.

Is it necessary to add that mathematicians themselves are not infallible?

(Poincaré)

I. INTRODUCTION

Mathematics is hard, but the epistemology of mathematics appears to be easy. Proofs are often taken to be the uncontroversial and unique path to first-hand mathematical knowledge.¹ Mathematical justification is assumed to be infallible, such that when a subject is first-hand justified in believing theorem T, and she does believe it, then she knows it. This is different from other types of beliefs such as perceptual and testimonial beliefs, where there is usually thought to be a gap between justification and truth.

The picture is clear and straightforward, at least as long as we can distinguish proofs from non-proofs. But therein lies the rub—we are fallible humans,

¹ My main focus is mainstream mathematics, and I am setting aside the (hard) problem of knowing the axioms.

unfortunately, and even experts cannot always tell the difference. For an illustration, we need to look no further than the English mathematician Andrew Wiles.

In 1637 French mathematician Pierre de Fermat wrote in the margins of his copy of Diophantus' *Arithmetica* that he had a proof of the theorem that no three non-zero integers a , b and c satisfy the *Diophantine* equation $a^n + b^n = c^n$ for any n greater than 2. However, Fermat also noted that he could not write out the proof for reasons of space. (What can you do? Margins are only so big!) So Fermat had a proof, or so he claimed.

In 1986 Wiles embarked on the project of proving Fermat's Last Theorem. To be sure, the mathematical community had already made significant progress, but a key result, the Taniyama–Shimura–Weil conjecture, remained unproven. Wiles announced his success in proving the conjecture, and consequently *Fermat's Last Theorem*, in 1993. And Wiles had a proof, or so he claimed.

As it turned out, Fermat's proof was never found, and Wiles' initial 'proof' was fallacious. A few months after its public announcement, a significant gap was found. Wiles managed to fix his 'proof', but only after a year of intense work in collaboration with his former student Richard Taylor. He finally managed to produce a correct argument in 1995.²

This short vignette shows that distinguishing a proof from a non-proof is often no trivial matter. Only through public scrutiny did Wiles realise that his argument was not a genuine proof. Our clear and straightforward picture of mathematical justification has to give way to something much blurrier. At this point, I propose to shift our attention from proofs to what I call *simil-proofs* (SPs)—arguments that *look like proofs* to the relevant experts. In good cases SPs are indeed proofs; in bad cases they are not. The chance that an SP is a genuine proof increases with its simplicity and the type and degree of scrutiny to which it is subjected. By focusing on SPs, I propose to move towards an account of mathematical justification calibrated to broad features of our social nature and cognitive architecture—including our shortcomings.

My main goal is to individuate some of the norms of mathematical justification for beliefs *that we can and do comply with* in typical mathematical practice. As the case of Wiles suggests, it is the public availability of arguments that allows other mathematicians to perform quality control. This is essential to filtering out non-proofs. That is why I consider the in-principle ability to share one's mathematical arguments (in a way to be described) to be a necessary condition of mathematical justification.

In Section II, I sketch the standard view of mathematical justification, the focus of which is the justification of mathematical propositions rather than

² Still, the set-theoretic assumptions needed for Wiles' proof continue to be a matter of debate (see McLarty 2010).

beliefs and the level of idealisation is very high. I then put forward working definitions of *proof* and *simil-proofs*.

In Section III, I offer an account of what I call *mathematical justification*. This type of justification is required when we do mathematics and take personal responsibility for our results. That is, if our beliefs turn out to be wrong or unsupported, we are to be considered epistemically blameworthy, at least in the absence of attenuating factors.³ Mathematical justification involves having a good mathematical argument and being able to articulate and defend it. I link this type of justification with SPs. Crucially, not all justified beliefs about mathematics are *mathematically* justified.⁴ In many circumstances it is perfectly fine to rely exclusively on testimony. However, on pain of regress, some beliefs must be based on mathematical arguments. The cases requiring mathematical justification are those in which an agent vouches for a result and is expected to be able to share her argument with her peers. Such cases include not only mathematicians like Wiles proving new results, but also non-professional agents. For example, students learning elementary geometry are supposed to come up with their own arguments and be able to defend them. It is important to notice that for an agent to have a mathematical argument does not preclude her from relying on previously accepted results, provided that she would be able to point to appropriate references. I also show that, in the proposed framework, mere reliabilism is inadequate for mathematical justification because the latter always involves a mathematical argument.

The presence of norms deriving from the social nature of the mathematical enterprise shows that mathematics is not some sort of individual quest in the domain of perfect a priori justification, but rather a practice that is often on all fours with the empirical sciences (and some philosophy). Obviously, in a paper of this length, some of the details and examples will have to be left aside. The overall goal is to paint a picture of mathematical justification and (simil-)proofs that is, if not clear and straightforward like the traditional picture, at least more representative of how the discipline is actually practiced.

II. PROOFS AND SIMIL-PROOFS

II.1 *Highly idealised accounts*

From the end of the nineteenth century until recently, philosophy of mathematics did not pay much attention to the nature of justification in *individual*

³ This is in line with a deontological conception of epistemic justification (see Alston 1988). See Brown (2020) for a discussion of genuinely epistemic blame.

⁴ Similarly, not all our justified beliefs about perceptual items are perceptually justified. For example, I can hold the justified belief that there is a cup on your table if you tell me that there is, and I have no reason to distrust you.

subjects.⁵ With the development of modern logic and the rise of foundational issues for mathematics, philosophers and philosophically inclined mathematicians started to focus their efforts exclusively on the problem of the justification of *bodies of mathematical theories*, rather than on the justification of individual *mathematical beliefs*.⁶

At the beginning of the twentieth century, the notion of mathematical proof was first equated with ‘deduction from the axioms’:

It was simply assumed that mathematical knowledge would have to be a matter of proof, that is, deduction from the axioms; the only question, then, was how the axioms and inference rules of the relevant axiomatic systems could be justified. Thus, the epistemology of individual discovery simply dropped off the agenda. So did any concern with actual thinking in mathematics. (Giaquinto 2007: 5)

Epistemology of mathematics thus ignored actual mathematical practice, in which both the individual and the social dimension are crucial.⁷ Questions were often very general, encompassing mathematics as a whole, and the actual thinking process of the individual never came up. Examples of such general questions were: Is mathematics just a part of logic? Is the law of excluded middle universally valid in mathematics? Later, in the twentieth century, epistemological inquiries into mathematics remained oblivious to the individual and social levels and focused largely on how to justify whole axiomatic systems (e.g. Hilbert’s program, conventionalism and holistic empiricism *à la* Quine). This approach could not (and was not designed to) appreciate the richness of actual mathematical practices and their epistemological relevance. Philosophers tended to adopt monolithic approaches seeking to fit all of mathematics into ‘a single system or framework’:

A reductionistic approach that simplifies all this complexity, turning everything into developments of a single system or framework, will not only flatten the interesting complexity of mathematical methodology, but also make it impossible to understand the process of building up mathematical knowledge. (Ferreirós 2016: 37)

Contemporary mainstream philosophy of mathematics has taken a similar attitude towards individual reasoning. Even when considering the beliefs of individual mathematicians, the problem is often reduced to beliefs about axioms. For example, Hartry Field’s (1988) version of Paul Benacerraf’s (1973) problem of access to *abstract objects*, which is a problem for all our mathematical beliefs, is reduced to the problem of explaining the reliability of our beliefs in *axioms*. The idea that *if* we are justified in believing the axioms, and we accept a logic,

⁵ Fortunately, in the last two decades a new approach to philosophy of mathematics took shape: *the philosophy of mathematical practice*. See Mancosu (2008) and Carter (2019) for an overview.

⁶ Frege’s *anti-psychologism* is one motivation for such a focus.

⁷ Even if L. E. J. Brouwer, the father of intuitionism, took the individual subject to be central to his conception of mathematics, he ignored the social dimension of mathematics.

then we are justified in believing the theorems we generate is generally taken for granted in this context. In other words, the whole business of *reasoning through mathematical arguments* is often not considered or problematised, as though we are logically infallible.^{8,9} While the received view in the philosophy of mathematics bypasses the typically human problem of going through mathematical arguments, the one I offer here bypasses the problem of the justification of axioms and of the appropriate logic—which, to be sure, is a difficult problem. We can therefore see the two as *complementing* each other, rather than being in tension. More generally, according to Paolo Mancosu (2008), the efforts of philosophers of mathematical practice can be seen as complementary to those of scholars working on foundational issues.¹⁰

To be sure, fallibilism has been a working assumption of the philosophers of mathematical practice.¹¹ Imre Lakatos (1963) inspired a great deal of practice-based research. Taking the history of mathematics as starting point, he articulated a quasi-empirical account of the development of mathematics. Interdisciplinary research has also contributed to painting a picture of mathematics as practiced by human, fallible subjects.¹²

Not all fallibilisms are alike, however. Ian Dove (2003) identifies five different forms of fallibilism in mathematics and argues that they are less controversial than what is normally thought.¹³ It will be useful to briefly list them. The first form of fallibilism is inspired by Einstein's view of applied mathematics as 'riddled with uncertainty' (*ibid.*: 18). It is a mundane form of fallibilism that does not bear on pure mathematics. The second and the third forms of fallibilism concern, respectively, non-deductive methods and the justification of axioms and are therefore beyond the scope of this article.¹⁴ The fourth form of fallibilism is the only one that bears on what is at stake here. It arises from the gap between formal and informal arguments. Dove argues that this type of fallibilism in turn divides into two sub-types. One arises from informal inductive arguments; 'deductive informal arguments do not support mathematical fallibilism' (*ibid.*: 111). The other has to do with Lakatos's notion of 'heuristic counterexample' and involves the redefinitions of concepts. Neither of these sub-types of fallibilism are relevant for the current inquiry. The fifth has to do with Quine's naturalism, which is also beyond the scope of this article.

The type of fallibilism advocated here is at once more and less mundane than those surveyed by Dove. It is more mundane because it has to do with

⁸ As I discuss below, there are exceptions.

⁹ A further complication of this view is that *our* mathematical arguments are not formal arguments, with very few exceptions.

¹⁰ Not all philosophers of mathematical practice are so ecumenical, however. More polemical approaches can be found in Lakatos (1963) and Corfield (2003).

¹¹ Thanks to one of the anonymous referees for suggesting me to add this discussion.

¹² See van Kerkhove & van Bendegem (2005).

¹³ I wish to thank an anonymous referee for bringing this reference to my attention.

¹⁴ For a discussion of these types of fallibilism in mathematics see also Paseau (2015).

mathematical thinking in general and the threat of introducing mistakes in complex deductive arguments. It is less mundane because it questions the usual account of mathematical justification in terms of proofs. Phillip Davis and Reuben Hersh's (Davis and Hersh 1999) *The Mathematical Experience* deals with a similar type of fallibilism in mathematical reasoning.¹⁵ However, my discussion moves in a different direction. While I give plenty of space to the realm of social issues in mathematics, I do not reduce, as they seem at times to do, the notion of correctness to it; see also Hersh (2014: 149).¹⁶

Most importantly, although my account is wedded to mathematical practice, it remains normative insofar as it delineates an ideal toward which actual mathematical practice can be seen as striving but not always achieving.¹⁷ Taking actual mathematical practice as the starting point also does not threaten the objectivity of mathematics. I am trying to capture some of the norms that regulate *human mathematical practice*, and not to give a sociological analysis of a specific human activity. As John Burgess puts it:

Science [including mathematics] is the product of certain intelligent organisms, ourselves, in a certain environment, the universe. [...] it is to be expected that our science would be the way it is in part because the universe is the way it is and in part because we are the way we are. Intelligent organisms like us in a different universe, and intelligent organisms unlike us in this same universe, would both have a different science, one unlike ours. (Burgess 1990: 13)

The key (and liberating) observation is that this does not bring disarray and subjectivism (or anti-realism):

It would obviously be fallacious to infer from the premise that we contribute as much as the universe to science being as it is, the conclusion that we contribute as much as the universe to the objects of science is about being as they are [...]. (*ibid.*)

II.2 Proofs

My goal in this section is to sketch a notion of proof that cannot be reduced to 'deduction from the axioms' but is tailored to human agents.¹⁸ The definition of proof I am looking for is not the only one possible, but it is one that is useful

¹⁵ Other examples are Kitcher (1984), Aspray & Kitcher (1988), Avigad (2020) and Prawitz (2012). Dag Prawitz explains:

To avoid misunderstandings let me say that I am in no way excluding the possibility of errors about whether something constitutes a conclusive ground. By speaking of conclusive grounds obtained by deductive reasoning, I am thus not advocating any kind of infallibilism. (Prawitz 2012: 890)

¹⁶ See De Toffoli (2020) for a discussion of this issue.

¹⁷ Even in a fallibilistic framework, not everything goes. There are many cases in which mathematicians violate relevant norms. For example, when they rely on authority in circumstances in which a mathematical argument would be necessary.

¹⁸ I will not deal with the justification of axioms or conjectures. We can assume that the axioms are true, and we know them or alternatively endorse a if-thenist account of mathematics.

for tracking proofs *as they are deployed in practice*.¹⁹ I will use it to spell out what it takes to achieve mathematical justification, which is first-order justification for a mathematical claim.

I am approaching proofs with the aim of articulating what mathematicians need to justify their beliefs. Admittedly, there are many more dimensions of epistemic evaluation beyond the justification of individual beliefs. However, one must start somewhere. The kind of justification I aim to delineate is non-factive, but in good cases it is one that can underwrite mathematical knowledge. Thus, even if at this stage I move only at the level of justification, the type of knowledge in sight remains *individual propositional knowledge*. I thus exclude from the present inquiry both (1) knowledge-how and (2) social knowledge.²⁰ However, I hope that the ground of the present account will be a fertile one from which it will be possible to formulate an account of these other forms of knowledge as well.

Enough with the preliminaries. It is now time to move towards a working notion of proof. I propose to start by considering Thomas Tymoczko's (1979) classic characterisation of mathematical proofs as: (1) *convincing*, (2) *surveyable* and (3) *formalisable*. I share Tymoczko's motivating intuitions, but after 40 years of methodological innovations in mathematics, it is worth considering whether refreshing his conditions is in order.

(1) *Proofs are convincing. No, they are shareable.* A proof is certainly not 'what convinces'. Without appropriate training and the appropriate circumstances, a proof can both not convince or convince for the wrong reasons (such as testimony), and an argument that is not a proof may be found convincing. Tymoczko puts it like this:

It is because proofs are convincing to an arbitrary mathematician that they can play their role as arbiter of judgment in the mathematical community. (Tymoczko 1979: 59)

But who is this arbitrary mathematician? Perhaps it is better to work with an *idealised mathematician*, one with the relevant training, who operates without debilitating contingent factors such as tiredness, morning sickness or other cognitive disturbances. The important point is that it does not matter whether proofs actually convince flesh and blood mathematicians. This is because, for

¹⁹ My characterisation of proof is high level. Its details can be filled out in different ways. See, for example, Avigad (2020), Azzouni (2004), Burgess (2015), Hamami (2019), and Tennant (1986) for conceptions of proofs that emphasise (in different ways) the connection between regular proofs and formal proofs and Detlefsen (2008: 3–32), Leitgeb (2009: 263–99), Rav (1999), and Tanswell (2015) for accounts in which this connection is downplayed.

²⁰ For a discussion of socializing scientific knowledge, see Longino (2002: ch. 6), and for a focus on social knowledge in mathematics, see Ferreirós (2016). Here I focus on the justification of individual beliefs. My view, however, can be extended to the justification of group beliefs. This would be particularly interesting to explain cases in contemporary mathematics in which proofs are so long that no individual can individually grasp them. See Lackey (2016) for a general discussion of justification of group beliefs.

some proofs, there might not be an instantiation of our ideal mathematician, that is, an actual agent with the right training that could assess the correctness of the proof, except its author *bien entendu*. This could happen when a particularly brilliant and particularly isolated mathematician develops some substantial new mathematics. To be able to appreciate her results, other mathematicians would need to digest all the relevant preliminary materials, and this process could take years. Real life examples are William Thurston's work on the geometrisation conjecture (Thurston 1994: 175) and perhaps Shinichi Mochizuki's argument for the *abc* conjecture.²¹ I thus propose to use the following more precise notion:

An argument is *shareable* if its content and supposed correctness²² could be grasped by relevantly trained human minds from a (possibly enthymematic) perceptible instance of a presentation of it.

Being shareable, unlike being convincing, does not entail that a proof must be graspable by a single individual mind. For some proofs, a community of minds, each working on a part, may grasp a proof even if no individual grasps the whole. For instance, the proofs of the classification of finite simple groups are enormous and require multiple agents to check them in their entirety. The presence of very long and complicated proofs is an ever more common phenomenon in contemporary mathematics and thus can no longer be ignored by any viable analysis of proofs.

Shareability is naturally a graded notion that presents many orthogonal dimensions of evaluation. Trivially, a very short and simple argument will be more shareable than a complex and long one. Comparisons are not always so simple, however. Proofs can in fact be evaluated with respect to features that influence their shareability in competing ways and thus give rise to trade-offs. For example, a very long proof deploying only basic techniques will be graspable by a wider audience. A short proof of the same result deploying sophisticated mathematical machinery will have a much more restricted audience, but will be verified more reliably and more quickly by the experts. Shareability is influenced by features like length, conceptual prerequisites, perspicuity and ease of verification.

²¹ Although many mathematicians now think that Mochizuki's argument has serious gaps, others claim that his argument is not understood because there is no appropriately trained audience for it (Fesenko 2019). See also De Toffoli, 'Mathematical Justification', draft for a detailed analysis of this case. Thanks to one of the anonymous referees for prompting me to present these cases.

²² Correctness is only 'supposed' because the shareability criterion can also be satisfied by fallacious arguments. I will assume that a proof presentation can contain minor, easily fixable, errors—called *mathematical typos* in the jargon—and still be a presentation of a proof. Otherwise I would have had to discard as invalid too much of contemporary accepted mathematics. See Avigad (2020: 3) for a similar approach.

The distinction between deductive and non-deductive arguments is orthogonal to shareability. Indeed, this condition narrows the gap between empirical sciences and mathematics by requiring mathematicians, like scientists, to participate in an essentially common enterprise in which results are in principle within reach of other practitioners who can submit them to scrutiny.

(2) *Proofs are surveyable. No, they are transferable, or maybe they are a priori verifiable.* Here is how Tymoczko spells out the surveyability constraint:

A proof is a construction that can be looked over, reviewed, verified by a rational agent. [...] It is an exhibition, a derivation of the conclusion, and it needs nothing outside of itself to be convincing. (Tymoczko 1979: 59)

This condition arises from the challenge posed by the use of automatic computations in mathematics, and in particular, from the very first computer-assisted proof—the proof of the 4-color theorem proposed by Kenneth Appel and Wolfgang Haken in 1974. This proof generated a heated discussion on the nature of proofs and the a priori status of mathematics. Computer-assisted proofs remain problematic for the mathematical community to this day. Consider the following case. In 1998, Tom Hales proposed a computer-assisted proof for the Kepler conjecture—which is about how to minimise space when packing equally sized spheres. In mundane terms, it states that there are no better ways to arrange oranges than the pyramid often used in fruit-stands. Hales' argument was only accepted in the *Annals of Mathematics* 2005 after 12 referees spent a great deal of time and effort to check it (Hales 2005). However, they admitted that they were only 99 per cent sure of its validity. It was only in 2017 that, thanks to a full formalisation of it, Hales' argument was finally accepted without reservations by the mathematical community (Hales et al. 2017). Notwithstanding their controversial status, computer-assisted proofs cannot be completely ignored since they are now a fairly common phenomenon.

Not only are computer-assisted proofs becoming more common in the landscape of contemporary mathematics but also, as discussed above, proofs requiring a collaboration of multiple agents are common. These are genuine proofs, even if they are not be surveyable. An up to date notion in the vicinity is 'transferability', discussed by Kenny Easwaran (2009) to find a criterion separating *probabilistic proofs*²³ from deductive proofs. As in the case of *surveyability*, the central idea behind *transferability* is that proofs are self-sufficient:

A proof is transferable just in case the sequence of propositions itself constitutes the proof—nothing about the method by which the propositions were

²³ These are not computer-assisted (deductive) proofs, but inductive arguments for mathematical conclusions whose correctness can be checked only on the condition that certain propositions are chosen at random. It is thus misleading to call them 'proofs'. Easwaran aims at answering the challenge to find a genuinely epistemic criterion able to distinguish probabilistic from deductive proofs posed in Fallis (1997), and Paseau (2015).

generated is essential. That is, mere consideration of the proposition suffices for a relevant expert to become convinced of the conclusion [...] (Easwaran 2009: 354)

Transferability is a generalisation of *surveyability*. The significant difference is that there are transferable (but not surveyable) proofs that no single expert could grasp:

as long as all the relevant proofs are transferable, there is no particular step for which testimonial justification is essential. (Easwaran 2009: 355)

Although it accounts for long proofs that no individual mathematician could hope to survey, transferability only partially solves the problem that arises with computer-assisted proofs. While no particular step of a computer assisted proof requires testimonial justification, it would be, at least in significant cases, impossible for the relevant experts to become convinced of its conclusion without actually using computers.

Ultimately, like *surveyability*, this notion is trying to capture the epistemic property that distinguishes mathematical from empirical arguments. This is one intuition behind the a priori. And indeed, there is a reasonable conception of the a priori²⁴ according to which an argument is transferable iff it can be verified a priori.²⁵ Since I have a different plan for this paper, I will avoid the quagmire of the debate about the a priori. However, I still discuss the issue *en passant* since the integration of a condition that is generally associated with very traditional views of mathematics may make my account more palatable to some readers.

When proofs are so long that a human being could not even in principle survey them, doubts about their a priori status emerge. As a matter of fact, even less would be sufficient to raise questions. With any substantive proof, we need to memorise and write down the numerous steps. The traditional infallibility of a priori warrants is incompatible with the fallible nature of memory and perception.²⁶ This is considered by some to be enough to introduce an a posteriori element in proofs. Others accept that a priori warrants can be fallible and delineate a conception of the a priori that can accommodate long proofs and that can be compared to transferability.²⁷ For instance, Tyler Burge

²⁴ The a priori has been, and can be, conceived in different ways. As Paul Boghossian & Christopher Peacocke (2000: 2) explain ‘There are several variant notions of the a priori, of varying degrees of strength. Each variant notion is generated by a different construal of “experience”’.

²⁵ Easwaran considers this possibility in Footnote 19 (2009: 360), but he does not decide one way or the other. Don Fallis (1997) rejects the idea that complex deductive arguments could be verified a priori, but he considers strict versions of the a priori.

²⁶ René Descartes’s seventh rule in his *Regulae ad directionem ingenii* states that to gain true knowledge it is necessary to be able to survey an argument in an uninterrupted sweep of thought (Descartes 1628).

²⁷ For a discussion of fallibilist and infallibilist conceptions of a priori justification see Casullo (2003: ch. 2).

(Burge 1993) argues that the role of memory in proofs is a merely enabling role:

Reliance on memory does not even add to the justificational force of the deductive justification. (Burge 1993: 463)

Writing down the steps of a proof can be put in the same category as memorizing the step.²⁸ The external medium is an extension of our memory that, had we had more internal memory, we would not have needed. The case from which Tymoczko's discussion started, however, is by no means dissolved. Computer assisted proofs still challenge the idea that proofs can be verified a priori. While Burge (1998) clears the way for thinking that we can get a priori justification from computer-assisted proofs, this is a longer and more controversial discussion (which exploits the analogy between knowing by a computer and knowing by testimony). However, note that it is exactly the same problem that the notion of *transferability* faces. Easwaran acknowledges that computer-assisted proofs 'are still in the borderline area' (Easwaran 2009: 352).²⁹

(3) Proofs are formalisable. Yes, and thus they are correct deductive arguments. Tymoczko points out that the formalisation is a 'local' rather than 'global' criterion. Due to Gödel's Incompleteness theorems, we know that a single formal system will not suffice to formalise all proofs. Moreover, not all formal systems will do. We cannot include what we want to prove in the axioms, for example. It must be a formal system whose axioms and inference rules are accepted by the relevant agents. If this looks problematic, we can always think of mathematical results as conditional on the axioms.³⁰ Requiring that proofs are formalisable is supposed to exclude fallacious arguments and assure the correctness of (informal) proofs. As Jeremy Avigad puts it:

One justifies an informal claim by proving it, and if the proof is correct, with enough work it can be turned into a formal derivation. Conversely, a formal derivation suffices to justify the informal claim. So an informal mathematical statement is a theorem if and only if its formal counterpart has a formal derivation. Whether or not a mathematician reading a proof would characterize the state of affairs in these terms, *a judgement as to correctness is tantamount to a judgment as to the existence of a formal derivation*[.] (Avigad 2020: 3, *emphasis added*)

I agree with Avigad's observations. There are various ways, however, to articulate the relation between informal proofs and formal derivations. For

²⁸ In other work in progress, I argue that even diagrammatic proofs can be conceived to be a priori.

²⁹ Here Easwaran refers to the acceptability criteria in journals, but these are the criteria he aims at tracking with the notion of transferability.

³⁰ This is common among mathematicians.

instance, it is important to underline that there are many formal derivations that can be appropriately associated with the same informal proofs.

To sum up, we have arrived at the following tentative definition of proof:

A *mathematical proof* is a correct deductive argument for a mathematical conclusion from accepted premises that is shareable and verifiable a priori.

Since the issue of the a priori status of mathematics is controversial and not our main concern, we can content ourselves with the following, more permissive notion of proof:

(PROOF) A *mathematical proof* is a correct deductive argument for a mathematical conclusion from accepted premises that is shareable.

This is not necessarily the unique or best definition of proof, but it is one that works well in this context. The *shareability constraint* is the innovative element. It is what guarantees that proofs are the sort of things that can be shared among appropriately trained human subjects. This may seem obvious, but in fact it is not. If proofs are equated with valid deductive arguments in a formal system, then not all proofs would be shareable since proofs could be so long that they could never be grasped, even by a collection of human minds, and thus are not shareable. Moreover, if some extra-terrestrial creatures were so cognitively different and advanced that for them the inference from the relevant axioms to Fermat's Last Theorem was an easy step, they would probably manage with reason alone to get to many results that would not be shareable with us.

The notion of proof that emerges from this discussion is context-dependent. Proofs, in fact, are supposed to be shareable among *appropriately trained* human agents. Benedikt Löwe & Thomas Müller (2008) also endorse a context-dependent conception of proofs and use it to argue that mathematical knowledge is similarly context-dependent. Since they only deal with knowledge, their framework might suggest that we have justification for believing a mathematical proposition (axioms aside) only when we have a proof of it. But the matter is not so simple. This is because there are cases of justified beliefs that do not constitute knowledge. The problem I consider here is that the ability mathematicians have to tell whether an argument is a proof is not infallible. Accordingly, our conception of mathematical justification should not be infallibilist. As Henri Poincaré reminds us: 'Is it necessary to add that mathematicians themselves are not infallible?' (Poincaré 1914: 47). We thus need to articulate a fallible notion in the vicinity of proof.

II.3 *Simil-proofs*

Consider Gauss's original argument for the fundamental theorem of algebra; it contained minor mistakes. Strictly speaking it was not a proof as defined above. However, the corrected version is a proof and it is generally considered

as being ‘essentially another version’ of the original one. Benjamin Fine and Gerhard Rosenberger narrate the intriguing history of the arguments proposed in support of the fundamental theorem of algebra:³¹

The first *published proof* of the Fundamental Theorem of Algebra was then given by D’Alembert in 1746. However there were gaps in D’Alembert’s *proof* and the first fully *accepted proof* was that given by Gauss in 1797 in his Ph.D. thesis. This was published in 1799. Interestingly enough, in reviewing Gauss’s *original proof*, modern scholars tend to agree that there are as many holes in this *proof* as in D’Alembert’s *proof*. Gauss, however, published three other *proofs* with no such holes. He published second and third proofs in 1816, while his *final proof*, which was *essentially another version of the first*, was presented in 1849. (Fine and Rosenberger 2012: 3, *emphasis added*)

In this passage, the authors are using the term ‘proof’ for arguments that are not actually proofs in my sense, but only appear as such. This is a standard practice, but to avoid confusion, I introduce the notion of *simil-proof*:

(SIMIL-PROOF) An argument is a *Simil-Proof (SP)* when it is shareable, and some agents who have judged all its parts to be correct as a result of checking accept it as a proof.³² Moreover, the argument broadly satisfies the standards of acceptability of the mathematical community to which it is addressed.

Not all SPs are proofs, and not all proofs are SPs. The former is true because SPs can contain significant mistakes; the latter is true because a subject S can come up with an argument that is in fact a proof, but neither S nor any other mathematician recognises it as such. Although an argument can be an SP accepted by only a single subject, due to the acceptability constraint, not *every argument* can be an SP. Roughly, we can say that an argument is ‘acceptable’ if it is a serious candidate for publication in a respected international journal or if it is already included in the existing literature.³³ This is a context-dependent *external* constraint on SPs. A subject could wrongly believe that her argument is an SP, for example by missing a blatant mistake.³⁴

While D’Alembert’s ‘proof’ was an SP accepted by him and part of the community (starting with those who published it), Gauss’s SP (published in 1799) was at the time accepted by the whole community as an SP. However, the same argument would nowadays not be considered a proof at all; hence, it is no longer an SP. Clearly, an argument can be an SP at one time but later cease to be an SP. Moreover, the *acceptance* of an SP is graded and is indexed to a

³¹ Thanks are due to Marcus Giaquinto for this reference and for long discussions of this topic.

³² This checking can include explicitly defined and documented computer-assisted computations.

³³ See De Toffoli (2020) for a discussion of acceptability criteria in mathematics.

³⁴ Thanks to Thomas Kelly for pointing out this issue to me. The situation is actually more complicated because the criteria of acceptability change with time and with context. However, since here I mainly focus on contemporary mathematics, I will set this issue aside.

particular time. Even considering a single mathematician, the time-dimension matters. It is common knowledge that when one thinks one has a proof, one should ‘sleep on it’ and re-evaluate the situation the next day. Only after this step should one start testing the argument by presenting it to one’s peers.

This reflects the idea, expressed by John Burgess (2015: 4), that proofs, or better, SPs, are generally convincing, but do not justify absolute confidence. He quotes this very apt passage of Hume:

There is no. . . mathematician so expert in his science, as to place entire confidence in any truth immediately upon his discovery of it, or regard it as any thing, but a mere probability. Every time he runs over his proofs, his confidence increases; but still more by the approbation of his friends; and is rais’d to its utmost perfection by the universal assent and applauses of the learned world. [Hume (1739: part IV, section I) in Burgess (2015: 4)]

A rough gauge of the strength of the acceptability of an SP, and thus of our justified confidence in it, is given by how many subjects (or communities) accept the SP in a non-testimonial way, in addition to how long it has been accepted. Note that many members of the mathematical community will accept an SP without checking it themselves, just by relying on other experts’ judgment—their acceptance of it comes therefore from mere testimony and will not contribute to the strength of the acceptance of the SP.³⁵ The degree of acceptance of an SP is hard to estimate due to the presence of dependence relations between practitioners. In Andersen, Andersen & Sørensen (2020) it is argued that mathematicians generally require that a putative proof survives multiple checks by several mathematicians before relying on its conclusion without checking it.³⁶ This is because, the authors argue, the possibility of the argument being fallacious diminishes with different checks. While I agree with their diagnosis, following Goldman (2001), I suspect that the situation is not that simple. They write: ‘our account does not at all place emphasis on who has checked a given proof, but on how many have checked the proof’, but the assumption that such checks are independent is, in practice, far-fetched. That is why we cannot bracket questions about who checked the proof and the dependence relations between different agents.

Using the notion of SP instead of proof, we can make sense of situations in which a well-functioning agent makes mistakes, but it is nevertheless plausible to think she is justified. This is the case when certain results are accepted by the community and it is only after a long time that a mistake is spotted. Here is a case.

³⁵ This implies that even if the actual number of practitioners who become convinced increases, the strength of acceptance of the SP will remain invariant.

³⁶ Thanks to one of the anonymous referees for bringing this reference to my attention.

(DEHN) In 1910 Max Dehn published an argument for what is now called *Dehn's Lemma*.³⁷ In 1929 Hellmuth Kneser discovered that Dehn's argument contained a significant gap and thus that the lemma has not been proved (although it had been used as a premise to prove many other results). It was only in 1957 that Christos Papakyriakopoulos published a proof of it.

Before 1929, Dehn had an SP that was widely accepted by the community. It is plausible to think that he was justified in believing he had a proof. Even if he did not, in this phase his confidence should be raised, in Hume's words, 'to its utmost perfection'. Another similar case is the one of Alfred B. Kempe's first attempt at proving the 4-color conjecture in 1879. His argument was accepted for 11 years before it was revealed to be fallacious.³⁸ Note that these are cases of published proofs. They show that even publications in esteemed journals can include fallacious arguments.³⁹

Still, there is a stark difference between Dehn and Kempe's situation and that of a reckless mathematician who has a bogus proof that he did not submit to any external scrutiny. The notion of SP can differentiate these cases since it makes sense of cases of mathematicians who did not lack cognitive diligence, but nevertheless got things wrong. Let's look further at how.

III. MATHEMATICAL JUSTIFICATION

I have not denied the claim that genuine proofs are the gold standard for mathematical justification. I just want to draw attention to the fact that mathematicians can go wrong, and as in the case of Dehn and Kempe, to the fact that they can go wrong while being cognitively and socially diligent. This is why we have to resort to SPs to spell out a standard for mathematical justification calibrated to *human* agents.

(MATHEMATICAL JUSTIFICATION) A subject S's belief that mathematical claim C (in need of a proof) is *mathematically justified* if and only if S *has* an SP.

Mathematical justification should not be equated with the more general notion of *justification for a mathematical claim*. An agent relying merely on testimony can hold justified beliefs about mathematical claims (and even know them) without being mathematically justified.⁴⁰ Mathematical justification is the type

³⁷ Dehn's Lemma claims that, given a 3-dimensional manifold and a disk within it (with possible self-intersection) whose boundary is a simple closed curve embedded in the 3-manifold without singular points of the disk, there exists another disk that is embedded in the manifold and has the same boundary curve of the original disk.

³⁸ See Sipka (2002).

³⁹ (Andersen, Andersen and Sørensen 2020: 3).

⁴⁰ See Paseau (2015).

of first-hand justification involved in *doing* mathematics.⁴¹ This is required in professional contexts but also in contexts in which non-experts produce and evaluate mathematical arguments. It is important to notice that evaluating SPs requires previous training since not every logical step will be spelled out. Moreover, in most cases, mathematicians do not include in their SPs arguments for all results they appeal to. As a general rule, in most situations, they should, however, be able to point to reliable sources or to produce an argument themselves for those results. Moreover, those results should be supported by highly acceptable SPs.⁴²

This understanding of mathematical justification brings mathematics closer to the empirical sciences, where it is almost universally accepted that justification is not factive.⁴³ Mathematics remains special since its method, when correctly carried out, is that of deductive proof. Those sorts of proofs entail their conclusion, and are thus infallible. When we focus on epistemology, however, logical relations do not exhaust the whole story. As I spelled out above, mathematical justification is not just a matter of logic but is relative to the agent and to our human shortcomings.⁴⁴

A caveat is in order. Although I have been talking about full-on beliefs and will continue in the same fashion, talk of degrees of belief would be more accurate. For a given proposition and a given SP, only a certain range of credences will be mathematically justified. Such a degree of justification varies along many dimensions, among which are: its evaluation with respect to shareability, the agent's thoroughness in her first-order verification of it and the strength of its acceptance.

The notion of mathematical justification explains our practice of taking responsibility for our beliefs. If our belief turns out to be false or not adequately supported and if we have no attenuating factors, then we are blameworthy. It is therefore restricted to reflective subjects—it excludes cases of mathematical cognition in children and non-human animals.⁴⁵ This is not a problem since we are dealing with maximally reflective contexts.

⁴¹ Proofs require previous training since not every logical step will be spelled out. Moreover, in most cases, mathematicians do not include in their proof all the proofs of the result they appeal to. They should be, however, able to point to reliable sources or to produce an argument themselves for those results.

⁴² This is in line with actual mathematical practice, Andersen, Andersen & Sørensen (2020).

⁴³ An exception is Littlejohn (2012). For a discussion of fallibilism in epistemology generally see Cohen (1988) and Brown (2018).

⁴⁴ In a more general setting, Gilbert Harman argued that logical deductions are not a kind of *reasoning* and cannot satisfactorily describe reasoning. This implies that to understand human rationality and human reasoning, we cannot abstract away from human limitations with respect to time and computational power. See for example Harman (2002).

⁴⁵ This is not an instance of over-intellectualisation since I am not concerned with all types of first-order justification for mathematical beliefs—such as the ones that non-reflective children can have.

The last thing we have to specify is how to understand the expression ‘having an SP’. I propose the following:

(HAVING AN SP) S *has an SP* of C if and only if, when prompted to articulate a reason for her belief in C, in the appropriate context, S would (in good faith) share the SP. Moreover, S would be able to appropriately reply to challenges and hold related dispositions. For instance, if the validity of an inferential step of her SP is questioned, S would be disposed to abandon it if she cannot defend it.

The condition ‘in the appropriate context’ is there to exclude trivial counterexamples such as the ones arising from temporary amnesia, stress, time constraints or other factors that might impede success in providing a proof to others.

This is a sufficient condition. Being able to articulate the SP and holding specific dispositions is enough to have an SP, and thus to be justified. Having a proof can be understood along the lines of the *basing relation*—that is, the relation that holds between a justified belief and the reason on which it is based. I prefer, however, to avoid discussing the basing relation because that would introduce controversial claims that are not immediately relevant here. According to the received view, the basing relation includes a causal component but having an SP does not.⁴⁶ The correct psychological explanation of how the belief that C was formed is not relevant in the present context, especially when S is mistaken about it. For instance, if S thought he formed the belief that C by considering the good reason R, but in fact S formed the belief because he was struck by lightning, S can still be justified if he is able to articulate an SP and respond appropriately to challenges.

This is also a necessary condition. As mentioned above, it is not always necessary for beliefs about mathematics. Testimony as well as non-deductive arguments could suffice to gain justification, and even knowledge in mathematics.⁴⁷ It is, however, necessary for *mathematical justification*. Without having an SP, the subject would not have an appropriate mathematical argument which other people could check. If Wiles did not present an SP, but an intuition in support of Fermat’s Last Theorem, other people would not have been able to appreciate his justification. It is only with mathematical arguments that it is possible for the community to detect and correct the errors that are a natural outcome of our fallibility.

I have spelled out a specific norm for mathematical justification. In practice, it is not always easy to determine whether the norm is complied with. A particularly interesting example involves so-called *folk theorems*.⁴⁸ These are theorems that are included in the accepted repertoire of results of a particular sub-field but that lack a published proof. Exactly because of the lack of public documentation, it can be hard to evaluate whether the practitioners are

⁴⁶ See Korcz (2019). An exception is Leite (2004).

⁴⁷ See Paseau (2015).

⁴⁸ Thanks to one of the anonymous referees for suggesting me to discuss this case.

mathematically justified in believing them or are simply relying on testimony or even on hearsay. The issue is vexing because it raises issues of epistemic injustice. As pointed out in Rittenberg, Tanswell & Van Bendegem (2020), the publication of a proof for folk theorems is sometimes prevented, leading to an unfair distribution of merit and the difficulty of filling gaps in the public literature.

III.1 *Simil-proofs are hopeful*

According to the proposed picture, in order to be mathematically justified, one needs to be able to articulate, in reasonable time and in the appropriate circumstances, a plausibly good mathematical argument. Mere reliability will thus never be enough—a mathematical belief produced by a reliable process (and lacking any other warrant) might be justified but will not be *mathematically justified*.

The inadequacy of mere reliability has been discussed in ways that will be relevant here by experimental philosophers in the context of challenging the acceptability of appeals to intuitions in philosophical arguments. Jonathan Weinberg (2007) explains that the problem with intuitions as a source of evidence cannot be reduced to the fact that intuitions are fallible.⁴⁹ All (or at least most of) our sources of evidence are fallible. There is, however, an important difference:

My contention here will be that what many uses of philosophical intuition are guilty of, but which our other standard sources of evidence are not, is *unmitigated* fallibility—a fallibility uncompensated by a decent capacity for detecting and correcting the errors that it entails. (Weinberg 2007: 323)

The parallel in mathematics is evident. Suppose Gianna is a mathematician who has an extremely reliable mechanism to form mathematical beliefs.⁵⁰ This mechanism, however, is a black box that neither she nor anybody else can look into. Because of her track-record, Gianna has faith that her mathematical beliefs will be true, notwithstanding her inability to produce mathematical arguments. Gianna is special, but she is human and thus fallible. And even if she might beat everyone else in terms of pure reliability, her fallibility is unmitigated. Her errors are incorrigible. In Weinberg's terminology, her reliable mechanism is hopeless:

A source of evidence that is not practically infallible is hopeful to the extent that we have the capacity to detect and correct its errors. Parallelism (as well as malice aforethought) demands that we similarly stipulate the term hopeless for devices for which our practices lack an appropriate sensitivity to their errors, and capacity for correction when such

⁴⁹ Thanks to Jessica Brown for bringing this reference to my attention.

⁵⁰ The character of Gianna could for example share analogous features with the mathematician Srinivasa Ramanujan.

errors are found. In general, a source of evidence will be hopeful to a greater or lesser degree, but I will reserve “hopeless” for those that are pretty far down at the low end of the spectrum. (Weinberg 2007: 327)

SPs are hopeful, and their degree of helpfulness increases and decreases together with their degree of shareability. Leaps of faith or appeals to a putative mathematical intuition might be reliable generators of mathematical beliefs but will always remain hopeless.

The importance of adopting a hopeful source of evidence is particularly strong in contemporary mathematics, since its results are increasingly sophisticated and difficult to check. As the case of Wiles’ proof of Fermat’s Last Theorem exemplifies, mathematics can flourish only on the condition that practitioners can overcome their individual shortcomings by checking on each other.

IV. TAKING STOCK

The focus of this paper was individual mathematical justification. My concern was restricted to mathematical claims that are thought by experts to require a proof. First, I sketched a working definition of proof, starting from Tymoczko’s characterisation of proofs as convincing, surveyable and formalisable. I reached a new characterisation, involving the notion of shareability. That proofs are shareable guarantees that the mathematical community can perform quality checks on the results it produces. I also suggested that there is an appropriate notion of the a priori according to which proofs are also a priori.

I then introduced a related notion—*simil-proofs*. These are arguments that look like proofs but may not be. I tied justification to the notion of SPs. This move led me to a non-factive account of justification—in my view, it is possible to hold a justified belief in virtue of a fallacious mathematical argument. This does not mean that mathematical justification is cheap, since serious external constraints must be met. To hold a mathematically justified belief, one must have an SP. This condition boils down to being able to articulate an adequate mathematical argument and having appropriate dispositions.

Although I focused on how SPs provide justification, SPs play other roles as well—they promote the development of mathematics by introducing new methods, notations and concepts that can then be redeployed in other mathematical contexts.⁵¹ However, it is the way in which the justificatory function is achieved that makes these other functions possible.⁵²

⁵¹ These roles have been discussed in Rav (1999).

⁵² In case of theorems proved in many different ways, such as the Pythagorean theorem, we do not seek further justification with each new proof, but this does not imply that the justificatory role is not carried out each time.

One might worry that social norms for mathematical justification would lead to conceiving mathematics as ‘socially constructed’, thus threatening its objectivity. However, as I explained in the introduction, such a worry is misplaced. It is true that our arguments about the entities and facts of mathematics are communal and inter-subjective, but this does not imply that such entities and facts are merely a product of social agreement. More generally, the epistemology of mathematics, as conceived here, does not entail any ontology of mathematics.

In sum, I proposed an alternative to the standard position in the philosophy of mathematics according to which justification requires a genuine proof and is therefore infallible. Requiring SPs is a norm we fallible humans can in most cases comply with and, therefore, one that can be used to articulate a notion of justification that is in line with actual practice. A fallibilistic account of justification does not entail that mathematics is misleading—mathematics is, in fact, very successful, and this is because mathematicians overcome individual shortcomings by engaging in a cooperative activity regulated by specific norms. The picture of the epistemology of mathematics proposed here gives a central role to our particular ways of grasping mathematical arguments. Proofs are understood as calibrated to our human cognitive powers and can be evaluated in terms of how easily we can process and verify them. This opens the door to a human—albeit for traditional views all-too-human—epistemology of mathematics.

ACKNOWLEDGEMENT

I wish to thank Marcus Giaquinto for many conversations on this topic. Thanks are also due to Andrew Chignell, Thomas Kelly and Hilary Kornblith for useful discussions. Finally, I wish to thank the two anonymous referees for helping me improve this paper.

REFERENCES

- Alston, W. (1988) ‘The Deontological Conception of Epistemic Justification’, *Philosophical Perspectives*, 2: 257.
- Andersen, L. E., Andersen, H. and Sørensen, H. K. (2020) ‘The Role of Testimony in Mathematics’, *Synthese*, doi: 10.1007/s11229-020-02734-9.
- Aspray, W. and Kitcher, P. (1988) *History and Philosophy of Modern Mathematics*. Minneapolis: University of Minnesota Press.
- Avigad, J. (2020) ‘Reliability of Mathematical Inference’, *Synthese*, doi: 0.1007/s11229-019-02524-y.
- Azzouni, J. (2004) ‘The Derivation-Indicator View of Mathematical Practice’, *Philosophia Mathematica*, 12: 81–105.
- Benacerraf, P. (1973) ‘Mathematical Truth’, *The Journal of Philosophy*, 70: 661–79.
- Boghossian, P. and Peacocke, C. (2000) ‘Introduction’, in *New Essays on the A Priori*. Oxford: Clarendon Press.

- Brown, J. (2018) *Fallibilism: Evidence and Knowledge*. Oxford: OUP.
- (2020) ‘What is Epistemic Blame?’, *Noûs*, 54: 389–407.
- Burge, T. (1993) ‘Content Preservation’, *The Philosophical Review*, 102: 457–88.
- (1998) ‘Computer Proof, Apriori Knowledge, and Other Minds: The Sixth Philosophical Perspectives Lecture’, *Noûs*, 32: 1–37.
- Burgess, J. (1990) ‘Epistemology & Nominalism’, in A. D. Irvine (ed.) *Physicalism in Mathematics*. Springer Netherlands: Kluwer Academic Publishers.
- (2015) *Rigor and Structure*. Oxford: OUP.
- Carter, J. (2019) ‘Philosophy of Mathematical Practice—Motivation, Themes and Prospects’, *Philosophia Mathematica*, 27: 1–32.
- Casullo, A. (2003) *A Priori Justification*. Oxford: OUP.
- Cohen, S. (1988) ‘How to Be a Fallibilist’, *Philosophical Perspectives*, 2: 91.
- Corfield, D. (2003) *Towards a Philosophy of Real Mathematics*. Cambridge: CUP.
- Davis, P.J. and Hersh, R. (1999) *The Mathematical Experience. Illustrated Edition*. Boston, MA: Mariner Books.
- Descartes, R. (1628) ‘Rules for the Direction of the Mind’, in E. Anscombe and P. T. Geach (eds) *Descartes: Philosophical Writings*, 1954th edn. Nelson.
- Detlefsen, M. (2008) ‘Proof: Its Nature and Significance’, in B. Gold and R. A. Simons (eds) *Proofs and Other Dilemmas: Mathematics and Philosophy*. Mathematical Association of America.
- De Toffoli, S. (2020) ‘Reconciling Rigor and Intuition’, *Erkenntnis*, doi: 10.1007/s10670-020-00280-x.
- Dove, I.J. (2003) *Certainty and Error in Mathematics: Deductivism and the Claims of Mathematical Fallibilism*. Rice University.
- Easwaran, K. (2009) ‘Probabilistic Proofs and Transferability’, *Philosophia Mathematica*, 17: 341–62.
- Fallis, D. (1997) ‘The Epistemic Status of Probabilistic Proof’, *The Journal of Philosophy*, 94: 165–86.
- Ferreirós, J. (2016) *Mathematical Knowledge and the Interplay of Practices*. Princeton: Princeton University Press.
- Fesenko, I. (2019) ‘About Certain Aspects of the Study and Dissemination of Shinichi Mochizuki’s IUT Theory’, *Manuscript*. <https://www.maths.nottingham.ac.uk/plp/pmzibf/rapg.pdf>. (Date accessed: October 30, 2019).
- Field, H. (1988) ‘Realism, Mathematics and Modality’, *Philosophical Topics*, 16: 57–107.
- Fine, B. and Rosenberger, G. (2012) *The Fundamental Theorem of Algebra*. New York: Springer.
- Giaquinto, M. (2007) *Visual Thinking in Mathematics*. Oxford: OUP.
- Goldman, A. (2001) ‘Experts: Which Ones Should You Trust?’, *Philosophy and Phenomenological Research*, 63: 85–110.
- Hales, T. (2005) ‘A Proof of the Kepler Conjecture’, *Annals of Mathematics*, 162: 1065–185.
- et al. (2017) ‘A Formal Proof of the Kepler Conjecture’, *Forum of Mathematics, Pi*, doi: 10.1017/fmp.2017.1.
- Hamami, Y. (2019) ‘Mathematical Rigor and Proof’, *The Review of Symbolic Logic*, doi: 10.1017/S1755020319000443.
- Harman, G. (2002) ‘Internal Critique: A Logic is Not a Theory of Reasoning and a Theory of Reasoning is Not a Logic’, *Elsevier Science*, 1: 17.
- Hersh, R. (2014) *Experiencing Mathematics: What Do We Do, When We Do Mathematics?* American Mathematical Society.
- Hume, David (1739) *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*. London: John Noon.
- Kitcher, P. (1984) *The Nature of Mathematical Knowledge*. Oxford: OUP.
- Korcz, K. A. (2019) ‘The Epistemic Basing Relation’, in Edward N. Zalta (ed.) *Stanford Encyclopedia of Philosophy*. [The Epistemic Basing Relation](https://plato.stanford.edu/archives/fall2019/entries/basing-epistemic/). <https://plato.stanford.edu/archives/fall2019/entries/basing-epistemic/>. Accessed 10 June 2020.
- Lackey, J. (2016) ‘What is Justified Group Belief?’, *Philosophical Review*, 125: 341–96.
- Lakatos, I. (1963) *Proofs and Refutations*. Cambridge: CUP.
- Leite, A. (2004) ‘On Justifying and Being Justified’, *Philosophical Issues*, 14: 219–53.
- Leitgeb, H. (2009) ‘On Formal and Informal Provability’, in Ø. Linnebo and O. Bueno (eds) *New Waves in Philosophy of Mathematics*. UK: Palgrave MacMillan.
- Littlejohn, C. (2012) *Justification and the Truth-Connection*. Cambridge: CUP.

- Longino, H. (2002) *The Fate of Knowledge*. Princeton: Princeton University Press.
- Löwe, B. and Müller, T. (2008) 'Mathematical Knowledge Is Context Dependent', *Grazer Philosophische Studien*, 76: 91–107.
- Mancosu, P., ed. (2008) *The Philosophy of Mathematical Practice*. Oxford: OUP.
- McLarty, C. (2010) 'What Does it Take to Prove Fermat's Last Theorem? Grothendieck and the Logic of Number Theory', *Bulletin of Symbolic Logic*, 16: 359–77.
- Paseau, A. (2015) 'Knowledge of Mathematics Without Proof', *The British Journal for the Philosophy of Science*, 66: 775–99.
- Poincaré, H. (1914) *Science and Method*. London: Thomas Nelson and Sons.
- Prawitz, D. (2012) 'The Epistemic Significance of Valid Inference', *Synthese*, 187: 887–98.
- Rav, Y. (1999) 'Why Do We Prove Theorems?', *Philosophia Mathematica*, 7: 5–41.
- Rittenberg, C., Tanswell, F. and van Bendegem, J. P. (2020) 'Epistemic Injustice in Mathematics', *Synthese*, 197: 3875–904.
- Sipka, T. (2002) 'Alfred Bray Kempe's "Proof" of the Four Color Theorem', *Math Horizons*, 10: 21–23.
- Tanswell, F. (2015) 'A Problem With the Dependence of Informal Proofs on Formal Proofs', *Philosophia Mathematica*, 23: 295–310.
- Tennant, N. (1986) 'The Withering Away of Formal Semantics?', *Mind & Language*, 1: 302–18.
- Thurston, W. (1994) 'On Proof and Progress in Mathematics', *Bulletin of the American Mathematical Society*, 30: 161–77.
- Tymoczko, T. (1979) 'The Four-Color Problem and its Philosophical Significance', *The Journal of Philosophy*, 76: 57–83.
- van Kerkhove, B. and van Bendegem, J. P., eds (2005) *Perspectives on Mathematical Practices: Bringing Together Philosophy of Mathematics, Sociology of Mathematics, and Mathematics Education*. Springer.
- Weinberg, J. (2007) 'How to Challenge Intuitions Empirically Without Risking Skepticism', *Midwest Studies in Philosophy*, 31: 318–43.

Department of Philosophy, Princeton University, Princeton, New Jersey, USA