Brit. J. Phil. Sci. **66** (2015), 775–799

Knowledge of Mathematics without Proof

Alexander Paseau

ABSTRACT

Mathematicians do not claim to know a proposition unless they think they possess a proof of it. For all their confidence in the truth of a proposition with weighty non-deductive support (for example, the Riemann hypothesis), they maintain that, strictly speaking, the proposition remains unknown until such time as someone has proved it. This article challenges this conception of knowledge, which is quasi-universal within mathematics. We present four arguments to the effect that non-deductive evidence can yield knowledge of a mathematical proposition. We also show that some of what mathematicians take to be deductive knowledge is in fact non-deductive.

1 Introduction

Playing around with the odd integers, you notice the following pattern:

$$1 = 1 = 1^2$$
$$1 + 3 = 4 = 2^2$$
$$1 + 3 + 5 = 9 = 3^2$$
$$1 + 3 + 5 + 7 = 16 = 4^2$$
$$1 + 3 + 5 + 7 + 9 = 25 = 5^2.$$

doi:10.1093/bjps/axu012

Advance Access published on July 23, 2014

Having tested a few more cases, by hand or computer, and observed that the pattern continues, you form the conjecture that the sum of the first n odd numbers is n^2. The conjecture formed, suppose you can't for the moment see how to prove it. At this point, is your evidence a sufficient basis for you to know the fact—for it is one—that the sum of the first n odd numbers is n^2? No. No matter how many instances of

$$1 + 3 + 5 + \cdots + (2n - 1) = n^2,$$

you check, you cannot thereby attain knowledge of the general proposition. The reason seems simple: although you possess some evidence, you lack a proof.

Mathematicians more generally claim not to know a proposition unless they think a proof of it is available.[1] The Riemann hypothesis is probably the most famous example: most experts believe it is true, for a variety of reasons, yet since the hypothesis has not been proved, they would not claim to know it.[2] Indeed, in mathematics it is a sort of methodological sin to claim knowledge in the absence of proof. Knowledge of a mathematical proposition apparently entails that someone or other has come up with a proof of it.

Is this quasi-universal attitude within mathematics correct? I shall argue that it is not and that mathematicians' conception of knowledge is overly strict: one can know a mathematical proposition without possessing a proof of it.

2 Why It Might Matter

Given its ubiquity, mathematicians would be foolish to deny the existence of testimonial knowledge of mathematics, which is empirical and hence non-deductive.[3] What they assume is that a proof lies at the knowledge's source, i.e. that someone or other in the mathematical community has proved the result. (We elucidate what we mean by 'proof' in Section 3.) Our question, therefore, is whether there can be non-testimonial knowledge of mathematics without proof. Where relevant below, construe any proof that contains a testimonial component (for example, because it exploits a computer to verify cases) as if it were grasped by a single individual. This qualification understood, why might it matter whether non-deductive evidence is sufficient for knowledge of mathematics? We briefly set out why it matters to the

[1] 'Proposition' in this article usually abbreviates 'mathematical proposition'.

[2] The Riemann zeta function, $\zeta(s)$, is the analytic continuation of $\sum_{n=1}^{\infty} \frac{1}{n^s}$ to $\mathbb{C}\backslash\{1\}$. Its trivial zeros are $-2, -4, -6, \ldots$ The Riemann hypothesis states that the non-trivial zeros of $\zeta(s)$ all lie on the line with real part ½, the so-called critical line. See (Borwein *et al.* [2008]; Franklin [1987]) for some of the evidence.

[3] Burge ([1993]) argues against this common conception of testimonial knowledge. As the focus is on non-testimonial knowledge, our main points are independent of Burge's arguments.

philosophy of mathematics, and we conclude the article with a potential implication for mathematics itself.[4]

The conception of mathematics as starting from self-evident axioms and proceeding via certain preserving rules of inference to theorems which via this process come to be known with certainty might, following Lakatos ([1962]), be called 'the Euclidean programme'. It is of course based on (a philosophical gloss of) Euclid's *Elements*. Pascal's *Of the Geometrical Spirit* is perhaps the *locus classicus* of the programme's philosophical articulation. Although the Euclidean programme no longer holds sway as a paradigm of mathematical (or of human) knowledge, one or two of its tenets are still standing. In particular, deduction from axioms is thought to be the highest form of justification available for a mathematical proposition and, most relevant here, as being necessary for knowledge. To challenge this tenet is thus to challenge one of the last bastions of Euclideanism.

Second, there are two aspects to the deeply ingrained understanding of mathematics as distinct from the empirical sciences. The objects of mathematics are standardly thought to be abstract and its sole knowledge-generating method deductive. The empirical sciences, in contrast, are concerned with concrete entities (as well as abstract ones), and employ non-deductive knowledge-generating methods (as well as deduction). To deny that knowledge of mathematics must be deductive is thus to chip away at the epistemological dimension of the empirical science/mathematics divide. It is to deny mathematics its alleged special epistemic status by assimilating its epistemology to that of the empirical sciences.

Third, an important question of mathematical epistemology is whether every mathematical truth is knowable. Optimists believe that all mathematical truths are knowable;[5] pessimists that some are unknowable.[6] However we understand the modality involved here, all participants in this debate have hitherto presupposed that the relevant knowledge is deductive in nature. If this presupposition is mistaken, however, knowability optimists have further resources to draw on. Even if there are mathematical propositions that

4 Williamson ([2000]) offers an extensive and spirited defence of the privileged role of knowledge in epistemology. There is a subfield of epistemology that concerns itself with the value of knowledge over and above that of other epistemic goods, a question sometimes called 'the extra value of knowledge problem', as in (Davis and Jäger [2010]) and many other papers. There are in fact several extra-value-of-knowledge problems, depending on what other epistemic goods are already in place (true belief, justified true belief, reliable true belief, and so on). Epistemologists usually assume that the extra-value-of-knowledge problem is not domain specific.

5 Gödel seemed to be one of their rank, indeed he proposed that all mathematical truths are provable in some higher set theory (Gödel [1964]).

6 The more general form of epistemic optimism is that every truth is knowable. Many of the Fitch-based arguments against general epistemic optimism, for example, (Williamson [2000], Chapter 12), do not touch the mathematical case, since the Fitch counterexample, 'q but it's not known that q' is not a strictly mathematical proposition. I discuss Fitch's argument in (Paseau [2008]).

cannot—in some appropriately interesting sense—be known deductively, perhaps they can all be known non-deductively. Suppose for argument's sake that the strongest system consisting of axioms human beings can ever come to know deductively is *S*, with associated theorem set *Thm*(*S*). If non-deductive knowledge of mathematics is impossible, humanly knowable mathematical propositions are a subset of *Thm*(*S*). But if it is possible, we could come to know mathematical propositions not contained in *Thm*(*S*). The existence of non-deductive knowledge of mathematics thus potentially strengthens knowability optimism.[7]

Fourth, the hypothesis that knowledge of mathematics may be non-deductive is important for mathematical instrumentalists. Instrumentalists construe some parts of mathematics, typically the more sophisticated ones, as instruments for establishing statements in other parts of mathematics, typically the elementary ones. The best-known twentieth-century instrumentalist is Hilbert, who insisted that a proof of the consistency of instrumental mathematics or, as he called it, 'ideal' mathematics should be given within elementary or 'real' mathematics. More generally, the instrumentalist divides mathematics into two parts, real and ideal,[8] and sees real mathematics as in some manner superior to ideal mathematics. Ideal mathematics discharges its instrumental role successfully only if it delivers the 'right' sort of real statements.

Gödel's second incompleteness theorem apparently shows that real mathematics cannot even prove its own consistency, let alone that of the stronger ideal mathematics. The theorem is therefore generally thought to defeat Hilbert's instrumentalism. The 'inductivist' response to this argument is that, although it may sink Hilbert's instrumentalism, Gödel's second incompleteness theorem does not sink mathematical instrumentalism more generally (Paseau [2011]). Non-deductive evidence may be sufficient to warrant high rational credence in the appropriate instrumentalist constraint—for example, consistency—thereby meeting one of instrumentalism's key aims. But what about knowledge? Since a proof of consistency would presumably amount to knowledge, the inductivist case would be strengthened if it turns out that non-deductive evidence can also yield knowledge. That non-deductive knowledge of mathematics is possible therefore also potentially strengthens mathematical instrumentalism.

In sum, the possibility of non-deductive knowledge of mathematical propositions opens several doors in the epistemology of mathematics.

[7] This is not to evade the limitative Gödelian results, unless non-deductive methods give rise to a non-recursively enumerable class of new items of knowledge.

[8] We stick to the vivid labels 'real' and 'ideal' without commitment to the particulars of Hilbert's instrumentalism (Hilbert [1967]).

3 Two Further Examples and Preliminaries

Following the tradition of Russell ([1910]), let's call non-deductive evidence 'inductive'. By inductive evidence for p, then, we mean any kind of evidence for p other than a proof of p; by definition, deductive and inductive evidence for p partition the types of evidence for p. Two further examples help illustrate some of the varieties of inductive evidence in mathematics.

Goldbach's conjecture (GC) states that every even number greater than two is the sum of two primes. Number theorists are highly confident of GC's truth on inductive grounds. First, there is enumeratively inductive evidence for GC, of the same kind as in our opening arithmetical example, though much more of it: GC has been checked for every even number up to about 4×10^{18} and double-checked up to a number not much smaller.[9] Second, various slightly weaker claims, for example that every sufficiently large odd number is the sum of three primes (Vinogradov's theorem),[10] or that every sufficiently large number (in another sense) is the sum of a prime and either a prime or the product of two primes, or that every even number is the sum of no more than six primes, have been proved. (Note in passing that this illustrates our broad use of the phrase 'inductive evidence for p' to include deductive evidence for statements related to, though distinct from, p.) Third, let G(n), the Goldbach number of n, be the number of different ways in which n can be written as the sum of two primes. GC can then be expressed as the claim for all even n greater than two, $\mathrm{G}(n) \geq 1$. Computer evidence shows that the function G(n) broadly increases for even n as n increases (with oscillation, but with an increasing trend), so that for instance for even $n \approx 10^5$, $\mathrm{G}(n) \geq 500$. In light of this evidence, that G(n) will suddenly drop to zero appears highly unlikely. Fourth, the ratio $R_N = \frac{1}{N} \cdot$(number of $k \leq N$ such that $G(2k) = 0$) has been proved to tend to zero as N tends to infinity; in other words, the density of counterexamples to GC is zero. (If GC is true, then R_N is zero for all N.) Fifth, Hardy and Littlewood's formula for the asymptotic number of representations of $N = p_1 + \cdots + p_m$, where $p_1 \leq \ldots \leq p_m$ are m primes, has been proved for $m \geq 3$, and if true for $m = 2$ implies GC for sufficiently large even numbers. There are more reasons to believe in GC, but we will stop here since our intention is only to illustrate a few inductive methods in mathematics. On the basis of this and other evidence, mathematicians are virtually certain of GC's truth—although hardly anyone would say we know GC since it lacks a proof.[11]

[9] The latest results can be found on the web page maintained by Tomás Oliveira e Silva <sweet.ua.pt/tos/goldbach.html>, accessed 31 May, 2014.

[10] At the time of writing, Harald Helfgott claims to have proved the unbounded version of this theorem.

[11] References I have drawn on in this paragraph include (Echeverría [1996]; Kumchev and Tolev [2005]; Nathanson [1996]).

Rabin's ([1980]) primality test, which performs a simple test on a randomly chosen number k smaller than N, illustrates a different kind of inductive argument. The detailed nature of the test is not important here, only the following facts: the conditional probability of a positive test result for k given that N is prime is one; in contrast, the conditional probability of a positive result for k given that N is composite is a constant $c < 1/4$. The tests are independent for different k, so if we run the test for T times, i.e. for T numbers smaller than N, and each time the test result is positive, then the probability that N is composite decays like 4^{-T} (taking $c = \frac{1}{4}$) so long as the prior that N is prime is non-zero. For example, if one's prior credence in N's compositeness is $1/2$, then Bayes' Theorem shows that one's posterior credence in N's compositeness after T positive results (and no other information) should be $1/(4^T + 1)$, and one's credence in its primeness $1 - 1/(4^T + 1)$.[12]

In a typical application of the Rabin test, one's confidence that the witness numbers are picked at random will admittedly be less than one. If the random number-generator is a 'black box'—its workings are unknown to the mathematician applying the test—then clearly the Rabin test applied a thousand times should not deliver a degree of belief anything close to $1 - 1/(4^{1000} + 1)$. But randomness in this context merely requires independence from any property relevant to the test. Knowledge of the Rabin test (which is easy to understand and to program) and some understanding of the mechanics of the random number generator at the algorithmic level are usually sufficient to convince oneself that the numbers selected are random in the appropriate sense. Similarly, a modest amount of testing suffices for virtual certainty that a particular computer application of the test is free of both hardware and software failure. For all these reasons, one's credence in N being prime as a result of applying the Rabin test to 1000 numbers smaller than N will typically not be quite as close to one as the probabilistic argument suggests; but in the best cases it will be fairly close.

Our three examples so far—the sum of odd numbers, GC, and primality testing—offer, it seems to me, progressively stronger candidates for inductive

[12] More generally, suppose $d \in [0,1)$ is one's prior credence that N is composite, and let z_T be one's posterior credence in the same proposition after obtaining T positive test results after T applications of the test (so that $d = z_0$). Assuming that the conditional probability of a positive result for k given that N is composite is exactly $1/4$ gives, by Bayes' theorem:

$$z_{T+1} = \frac{1}{4}.z_T / \left(\frac{1}{4}.z_T + 1.(1 - z_T)\right) = \frac{z_T}{(4 - 3z_T)}.$$

One may check by induction that the solution to this recursion equation is

$$z_T = z_0/(4^T - (4^T - 1)z_0)),$$

and that the case in which $z_0 = \frac{1}{2}$ yields $z_T = \frac{1}{(4^T + 1)}$. Note that $z_T \leq \left(\frac{1}{4^T}\right).\frac{z_0}{(1 - z_0)}$.

knowledge of mathematics. But even the strongest of the three is controversial, as are other examples in the literature (for instance, the example in (Steiner [1975], pp. 102–8)). If one is to establish that inductive knowledge of mathematics is possible—and indeed actual—one apparently needs to go beyond such examples and produce suasive theoretical reasons. That will be the task of the next four sections.

Before that, it is worth addressing a 'naturalist' objection, which I have heard voiced by some mathematicians, to the effect that no philosophical considerations can overturn virtual consensus amongst mathematicians that knowledge is impossible without proof. One response might be to question the presumption that philosophy cannot challenge established mathematics. A second response is more effective in this context. Our challenge is not to what mathematicians say about mathematics but to what they say about its epistemology, which is not their professional specialty. It is neither especially implausible nor arrogant to suppose that they might be collectively mistaken so far as the epistemology of their subject is concerned, to suppose more precisely, that they are prey to a tendency to misapply the term 'knowledge' and err in related epistemic discriminations. If the arguments below are on the right track and mathematicians systematically misjudge what it takes for knowledge of a mathematical proposition, this is not a devastating indictment of the profession. All but the most ardent naturalist should accept the legitimacy of challenging mathematicians' epistemological, as opposed to mathematical, pronouncements.

A few words about the deductive/inductive knowledge distinction are also in order. For the arguments in Sections 4–7 to go through, we require no more than a decent handle on this distinction. We can all agree, for instance, that empirical knowledge of physical facts is not deductive (Section 7), or that the type of evidence we possess on behalf of some current axioms and new axiom candidates for set theory is not entirely deductive (Section 6). Although there is no need to analyse the notion of deductive knowledge for our purposes, a tentative and partial elucidation might nonetheless prove helpful. We may think of deductively known axioms as true statements that the subject believes on the basis of sufficient intrinsic justification. Anyone interested in analysis or a deeper account will of course want to know more about each of the terms 'sufficient', 'intrinsic', and 'justification'. For our purposes, suffice it to say that justification is intrinsic if it does not depend on other statements of the system or on the evidence for them. Note that our characterization allows that axioms may be known in this way without attaining the exacting standard of self-evidence, or that they may be thus known by virtue of constituting the subject matter at hand (an example may be set theory's axiom of extensionality). We may in turn think of a statement as known deductively just when it is known by deduction from deductively known axioms.[13] Our attitude towards what constitutes a legitimate deductive argument or proof is roughly that of the working mathematician: proofs are the types of things generally referred to by that name in mathematics journals and mathematics lectures. In particular, we do not insist that proofs need be formal or even formalizable, and we certainly do not insist that proofs must take the form of a strict deduction in some formal system; we allow that proofs may include pictures or diagrams, as they customarily do in geometry or algebraic topology say; and we allow as a proof something that a stricter usage might call a proof sketch, i.e. a fairly telegraphic set of instructions for generating a deductive argument. Our elucidation of deductive knowledge is thus rather liberal, as it should be if the conclusion that there can be non-deductive (non-testimonial) knowledge of mathematics is to be interesting.

With these preliminaries behind us, we turn at last to the arguments. The first two are plausibility considerations from general epistemology; the third is based on how we know the axioms of a literally-interpreted branch of mathematics (here: set theory); and the fourth attempts to establish that knowledge of mathematics is attainable without proof by proposing that such knowledge is sometimes derivable from empirical knowledge.

4 An Exclusive Epistemic Virtue of Proof?

Our first general-epistemological plausibility argument is that deduction does not seem to possess any special knowledge-conferring property that inductive arguments lack. As we cannot hope to survey all the possible candidate properties here, and since much work on this front has already been done, we shall be brief. Fallis ([1997], [2002]) has investigated the epistemic virtue probabilistic 'proofs' might be thought to lack that deductive arguments enjoy, and has assessed the merits of certainty, threshold degree of certainty, conditional certainty, and *a priori* warrant. His conclusion is that none of them fit the bill. This section is intended as something between a reminder of, and a short supplement to, Fallis' arguments generalized to the question of what epistemic virtue inductive arguments might lack that deductive arguments enjoy. We select four virtues with either *prima facie* plausibility or historical pedigree—to wit, certainty, unrevisability, reliability, and explanatoriness—and sketch how the arguments might go in each case.[14] The conclusion that apparently

[13] We allow one-line deductions so that deductively known axioms are deductively known *tout court*.

[14] Setting testimony aside, as explained.

emerges is that there is no (non-question-begging)[15] property had by deductions necessary for knowledge of mathematics that cannot be had by inductive arguments.

Certainty: According to the traditional picture, axioms are certain and inferential steps are certainty-preserving. Deductions thus deliver their conclusions with certainty, whereas inductive arguments only offer uncertain unknowledge.

This traditional picture is a long way off the mark. That we are in possession of a proof of p does not imply that we should be certain of p. We may not be certain of the proof's premises, and it would be irrational to be more certain of an argument's conclusion than its premises if the argument is our sole reason for believing the conclusion. We don't know the axioms of set theory with certainty, and in practice the starting points of contemporary research in a branch of mathematics are not known with certainty, even if they are universally assumed. It seems rationally permissible to give degree of belief less than one to many claims, in mathematics and elsewhere.[16] Given the normative (some would say constitutive) link between degrees of belief and betting behaviour, certainty in a statement implies that you should be willing to stake your life, risk eternal damnation in hell, and so on, for no return on its truth. Yet it is rationally permissible to be unwilling to go to the stake over the truth of the conclusion of any argument; it is rationally permissible for fallible inquirers to maintain a degree of doubt in virtually any claim.

Second, inference rules need not be certainty preserving nor indeed credence-preserving. Typical deductive steps in mathematics are relatively large. Indeed, at the research level many such leaps are transparent only to experts. Even 'atomic' inferential steps, of the type found in natural deduction systems, need not be credence-preserving. It is not irrational, for example, to give credence of less than one—say credence $1 - \varepsilon$ for positive (perhaps very small) ε—to the proposition that double negation elimination (not-not-p, therefore P) is truth-preserving. Double negation elimination is a classical rule questioned (for non-decidable statements) by intuitionists and intuitionist sympathizers, a camp that includes L. E. J. Brouwer, Arend Heyting, and Michael Dummett—formidable thinkers all. Even if you are a committed classical logician, given the history of this dispute it would not be unreasonable—and it would certainly be rationally permissible—to lower your credence in the rule's truth-preservation by some small positive ε. It seems rationally permissible for example to give positive credence, however slight,

15 A question-begging property would be a property such as being deductive or logically entailing the conclusion.

16 We interpret certainty in p as entailing degree of belief one in p; no commitment to the converse is made.

to the hypothesis that Dummett's ([1973]) arguments for intuitionism are sound.[17]

In sum, deductive steps do not necessarily preserve certainty for two reasons. One is that a mathematician might reasonably doubt a step (for example, double negation elimination). Another is that a mathematician might reasonably doubt that all the steps were applied correctly, as is illustrated by the fact that published proofs are occasionally retracted (see below).

Two final points are also worth noting. At least some empirically (and thus inductively) justified propositions about one's immediate environment, or perhaps of one's own mental states, are more justified than various theorems we believe. And in any case the certainty of knowledge is a side issue, since our interest is in whether inductive knowledge of mathematics is possible, not whether such knowledge would be certain. All these points could be developed further, of course, but it should be fairly clear that the delivery of certain knowledge is not the hallmark of deduction.

Unrevisability: A deductively valid argument is monotonic in the sense that new premisses will not convert it into an invalid one, in contrast to inductively valid arguments. The latter are systematically subject to one type of threat of non-permanence that deductive arguments are not. This, it might be said, is the preserve of deductive arguments.

It is indeed characteristic of inductive arguments that the addition of premises can invalidate them. This logical sense of permanence—monotonicity—however, is quite distinct from epistemic permanence. Proof almost never (if at all) achieves epistemological permanence in the sense of rational unrevisability of its conclusion. We may later come to reject some of the premises of any proof we currently accept, or we might reject some of its inferential steps, or we may for some other reason come to suspect the argument of invalidity (for example, because of its length). The history of mathematics throws up many examples of erroneous 'proofs', even by leading mathematicians,[18] and of centuries-long tenets now widely repudiated: that all numbers are positive or that they are all real, that space is Euclidean, that the whole is greater than the part, that all functions are linguistically expressible, and so on. If fallibilism about p is the doctrine that we should maintain belief in p in an open-minded spirit that allows that

[17] Paseau ([forthcoming]) develops these ideas further. Although deductively valid arguments shouldn't always preserve high credence, it is notoriously difficult to give a theory of rational credence that allows for this to fail. As many have noted, for example, Franklin ([1987], p. 13), a formal theory of these facts is not available should not lead us to cast doubt on their existence.

[18] See, for example, (Davis [1998]; Hersh [1979], p. 41; [1997] pp. 43–5; Crowe [1988], Section 4, pp. 263–5; Krantz [2011], pp. 215–7). The remarkable Lecat ([1935]) lists 500 examples of long-standing—as opposed to easily detected—mathematical errors, committed by 330 mathematicians, including Abel, Cauchy, Cayley, Chasles, Descartes, Euler, Fermat, Galileo, Gauss, Hermite, Jacobi, Lagrange, Laplace, Legendre, Leibniz, Newton, Poincaré, and Sylvester.

future evidence may give us reason to reject p, then we should be fallibilist about many deductively justified mathematical beliefs. In short, the traditional picture that proof entails unrevisability applies at best to small pockets of mathematics.[19]

Reliability: Although deduction is by and large more reliable than induction as a general method, in mathematics and elsewhere, and although the best deductive arguments may be more reliable than the best inductive ones, it is not true that knowledge-generating deductive arguments are always more reliable than inductive ones. The Rabin test exemplifies just how reliable inductive evidence can be, namely, extraordinarily reliable even by the high standards of the exact natural sciences. More generally, although most mathematicians quite reasonably regard deductive evidence as generally speaking more secure than inductive evidence, they would concede that a deductive argument is not always a better guarantee of its conclusion's truth than an inductive one. Indeed, keen 'experimentalists'—mathematicians who develop and make much use of such methods—point out that strong inductive arguments are more convincing than the specious deductive arguments one regularly encounters in mathematical journals.[20]

As a brief illustration—such examples are easily multiplied—at the time of writing many group theorists think that the best evidence that the classification of finite simple groups is complete is not the behemoth of a proof finally collated in the 1980s. They tend to value more highly the inductive evidence that no finite simple group has yet been found that fails to slot into any of the classifying families. Note in passing that these points suitably elaborated see off the claim that deduction is necessary for knowledge of mathematics, but not knowledge of other domains, because mathematicians have high standards.[21]

Explanatoriness: As for explanation, more briefly, a deductive proof might not offer an explanation either. Many deductive proofs are not explanatory. For example, proofs of generalizations that proceed disjunctively, with different instances having entirely different sub-proofs, are not explanatory. Famous examples include the Appel-Haken proof of the four-colour theorem in 1976, Hales' 1998 proof of Kepler's sphere-packing problem, and the

[19] For a more systematic discussion of the revisability of our knowledge of mathematics, see (Casullo [1988]).

[20] As a leading experimentalist puts it:

> I don't trust humans a lot. You know, people think that a written proof is the gold standard. I think many mathematical papers and arguments contain errors and gaps and the only reason we don't find them is because they don't get read. On the whole, the building of mathematics is sound, but if a mathematical statement works out in a computer test, then I believe it a lot more. (Sturmfels [2008], p. 5)

[21] Or to put it in contextualist terms, because mathematical contexts are high-standard contexts.

already-mentioned classification of finite simple groups. These examples are of course cases of testimonial knowledge, since they involve computer assistance or reliance on others' results. Bracketing the testimonial aspect, the point is that the resulting proofs in themselves are too disjunctive to be explanatory. If I had more space, I would also develop the point that if by explanation we mean something that conveys understanding,[22] heuristic arguments can be more explanatory than deductive ones.

In sum, none of our candidates differentiates inductive arguments from deductive ones. Combining these considerations with those advanced by Don Fallis, a challenge arises for those who would defend orthodoxy: find an epistemic property, *E*, enjoyed by (some) deductive arguments but no inductive ones, and show that *E* is necessary for knowledge-generation in mathematics. Although we have by no means exhausted the possible candidates for *E*, the absence of plausible contenders is highly suggestive.[23]

5 Analyses of Knowledge

Although it is virtually tautologous that knowledge of mathematics is one of the many forms of human knowledge, this fact is pregnant with implications for mathematical epistemology. General considerations about knowledge suggest that inductive knowledge can exist in mathematics just as it can elsewhere.

As is very familiar, Gettier ([1963]) demolished a long tradition of equating knowledge with justified true belief. Ever since, the search has been on for the solution to the equation

$$\text{Knowledge that } p = \text{True Belief that } p + X.$$

Notoriously, none of the myriad proposed conditions has achieved consensus, although many have been thought to be along the right lines, or at least to cover an important range of cases. For brevity, call any analysis that has gained at least some traction in the literature a 'right-track analysis'. Any right-track analysis falls into one of two categories, I claim: either it allows that knowledge of mathematics may be inductive; or it does not apply to knowledge of mathematics, and *a fortiori* does not privilege a deductive, as opposed to an inductive, route to knowledge of mathematics.

[22] Brown ([2012]) sees this as one of the two chief senses of explanation.

[23] Easwaran ([2009]) argues that proofs but not inductive arguments are transferable. Whether inductive arguments are truly untransferable remains unclear; see Jackson ([2009]) for critical discussion. Either way, by his own account, Easwaran is offering a descriptive account of what makes an argument mathematically acceptable. He is not suggesting that transferability is necessary for knowledge-generation independently of what mathematicians think.

Establishing this moral calls for a vast literature review. We cannot do that here, so we illustrate the point with a handful of analyses, starting with two classic conditions (supplementary to true belief):

- the belief that p is justified (see Plato's *Theatetus*);
- the belief that p was caused by the fact that p ([Goldman 1967]).

The Platonic condition allows for inductive knowledge of mathematics, assuming that inductive evidence (for example, of the type we have for GC, or that found in primality testing) can justify a mathematical proposition. This assumption follows on virtually all prevailing conceptions of justification;[24] Franklin ([1987]) is among the many who have argued that inductive evidence in mathematics is indeed a form of justification. On the other hand, on the usual platonist picture, Goldman's causal condition falls under the other category: it rules out all knowledge of mathematics, deductive or inductive, given the relatively uncontroversial assumption that facts about abstract objects do not cause our beliefs in them.[25] Precisely for this reason, Goldman did not intend his analysis to apply to mathematics.

The same style of argument can be run on other post-Gettier right-track analyses, some of the earlier ones being:[26]

- the belief that p is not inferred from a false lemma (Clark [1963]) or its justification must not essentially rest on a false assumption (Harman [1973]);
- there is a law-like connection between the fact that p and the belief that p (Armstrong [1973]);
- the belief that p is produced by a reliable process not undermined by the subject's cognitive state (Goldman [1986]);
- if it were the case that not-p then the subject wouldn't believe p (Nozick [1981]).

In the best inductive cases (think of judicious primality testing or the examples from Section 7 below), the belief that p is justified and its justification does not rest on a false lemma or a false assumption. Armstrong's ([1973]) 'reliable-indicator' analysis is in the other camp, as it was not intended to cover

[24] Some epistemologists have recently tried to narrow the gap between justification and knowledge by seeing the former as a type of 'would-be knowledge' (Bird [2007]; Smith [2010]); Sutton ([2007]) takes this tendency to an extreme by arguing that there is no gap between justification and knowledge. McGlynn ([2012]) criticises the former accounts, supporting instead a traditional position according to which one is justified in believing that p iff p is sufficiently likely given one's evidence. A conception of justification that assimilates it closely to knowledge may be compatible with—indeed, it may promote—the thesis that inductive knowledge of mathematics is possible.

[25] But not entirely uncontroversial. For instance, some Gödelian platonists may deny this claim.

[26] These are some of the proposed analyses' characteristic clauses. I have omitted qualifications, refinements, and extra clauses that do not affect our general point.

knowledge of mathematics. On most accounts of reliability, the best inductive methods in mathematics satisfy Goldman's ([1986]) 'process-reliabilist' conditions. Finally, Nozick's truth-tracking condition is not easily applied to mathematics because mathematical propositions are usually thought necessary, and we lack a decent account of counterfactuals with impossible antedents.[27] Countermathematical conditionals about which we do have intuitions, for what these are worth, tend to support our conclusion. For example, we are liable to think that if the Riemann hypothesis were false we would lack the evidence we in fact possess for it, and consequently that we wouldn't believe it. Similarly, if the perpendicular bisectors of Euclidean triangles did not meet in a point then the physical experiments described in Section 7 presumably wouldn't lead us to conclude that they do. If a particular number were composite rather than prime then the primality test evidence would be different. And so on. In sum, whether we take such intuitions seriously or not, the truth-tracking condition does not privilege a deductive route to knowledge of mathematics.

Notice in passing that the best inductive mathematical cases are quite unlike the much-discussed 'lottery cases', in which a subject has an arbitrarily high degree of justification for her belief that her ticket won't win ($1 - \frac{1}{N}$, where N is the number of tickets) and yet doesn't know this fact. In the lottery cases, the subject's evidence for p is insensitive to whether p is true: she has the same evidence and belief in the scenario in which she holds the winning ticket as in the $N - 1$ nearby scenarios in which she holds a losing ticket. Not so for the best inductive cases in mathematics, for the reasons just sketched.

Nozick's truth-tracking condition has come to be known as 'sensitivity'. Williamson ([2000]) considers the related 'safety' condition on knowledge, that the subject could not easily have falsely believed that p. The best inductive cases satisfy this condition too, since mathematical beliefs are safe if true.

Of course this is only a selection from a vast number of analyses and criteria in the literature, and we have not considered accounts of knowledge other than analyses or criteria. But broadening the range would not alter the moral. General epistemology has not converged on any condition that excludes an inductive (non-testimonial) route to knowledge of mathematical propositions. No right-track analyses allow deductive routes to mathematics whilst ruling out inductive ones. As a prominent field of inquiry, mathematics is, and should be, a test case for general epistemology.[28] If the only (non-testimonial) route to knowledge of a mathematical proposition were deductive, you would expect some prominent general accounts of knowledge to have this

[27] Many philosophers claim that counterfactuals with impossible antecedents are true, but they usually do so for shallow technical reasons. Soames ([2010], pp. 56–8) offers a succinct overview.

[28] That said, in practice several epistemologists follow Goldman ([1967]) in explicitly restricting themselves to empirical knowledge or knowledge of contingent truths.

implication. That none does is not a vindication of the possibility of inductive knowledge of mathematics, but it does lend it plausibility.

6 The Inductive Basis of (Some) Deduction

It is familiar that some mathematical axioms gain their justification from their mathematically desirable properties. Russell ([1907]) was the first modern philosopher to emphasize this form of justification.[29] He called it the 'regressive method' and used it to justify his axiom of reducibility (Russell [1910], pp. 250–1). Gödel famously held a variant of this view ([1964], p. 261). The more recent work of Penelope Maddy has stressed the various extrinsic reasons for believing the axioms of set theory ([1988], [1990], [2011])—that the axioms are mathematically effective and fruitful—as opposed to intrinsic reasons, for example that the axioms are self-evident or part of the concept of set ([2011], p. 47). In this terminology, at least part of the reason for accepting ZFC's axiom of choice, the (infinitely many instances of the) axiom scheme of replacement, or the axiom of infinity is inductive. The extrinsic argument for these principles is based on the assessment of their mathematical power; for example, the axiom of choice yields a workable theory of transfinite arithmetic. Indeed, both the 'Z' and the 'F' in 'ZFC'—Zermelo and Fraenkel—believed that the system's axioms were in part extrinsically justified.[30]

Let's consider the situation schematically. Suppose a set of axioms, O (for 'old'), is given. Here we are excluding 'structuring' axioms—about general structures such as groups, rings, fields, topological, or metric spaces—which define their subject matter. Our focus instead is on axioms about a particular structure: think of set-theoretic axioms, aiming to capture facts about the universe of sets. It does not matter here just how the axiom set O is known. A new axiom candidate, N, is proposed at some point. N is quickly seen to have many extrinsic virtues; as a consequence it is adopted as an axiom. In the course of one mathematical generation, say, the set of axioms expands from O to $O \cup \{N\}$.

Suppose p is a consequence of N in conjunction with axioms in O, as comes to be appreciated, and that N is necessary for p's derivation. Thus:

$$O \nvdash p,$$
$$O \cup \{N\} \vdash p.$$

[29] Irvine ([1989]) offers a pertinent discussion.

[30] See (Zermelo [1967], p. 189), which contains his dictum 'Actually, principles must be judged from the point of view of science, and not science from the point of view of principles fixed once and for all'. Fraenkel: 'the intuitive or logical self-evidence of the principles chosen as axioms [of set theory] naturally plays a certain but not decisive role; some axioms receive their full weight rather from the self-evidence of the consequences which could not be derived without them' (cited in (Lakatos [1976], p. 250), which contains more relevant quotes).

What is the evidence for p, assuming a proof from $O \cup \{N\}$ is our route to it?[31] The evidence for logic, the evidence for O,[32] and the evidence for N. But the evidence for N, we said, was inductive. Hence the evidence for p is in part inductive; it is not strictly deductive. If one knows that p via logic and knowledge of N and O, then if one's knowledge of N is inductive then in the typical case one's knowledge of p will also be inductive.

It is worth bringing into the open an epistemological rule of thumb implicit in this argument. To use a metaphor from genetics, in the normal run of cases inductive evidence is 'dominant' and deductive evidence 'recessive', meaning that a combination of the two is usually inductive (compare the 'dominance' of *a posteriori* evidence and the 'recessiveness' of *a priori* evidence). For example, suppose that I justify p_1 by deducing it from some premises that include the fact that the British winter of 2010 was cold, for which I have inductive justification (I experienced it); that I justify p_2 by deducing it from some premises that include the fact that $2^{61} - 1$ is prime, my justification for which is inductive (it consists of computational evidence);[33] and that I justify p_3 by inferring it using a non-deductive rule of inference (enumerative induction, say). My justification for each of p_1, p_2, and p_3 is inductive overall: if I know any one of these propositions, I know it inductively.[34]

The schematic argument just given is an argument for the possibility of inductive knowledge of mathematics (more precisely, for knowledge of any branch of mathematics with an intended subject matter). That such knowledge actually exists follows from the supplementary fact that our knowledge of some axioms is in fact inductive. The argument for this last fact—for instance, that the mentioned axioms of ZFC are known inductively—can be found in the works cited above and in many others, and will not be rehearsed here. Perhaps a twenty-first century example of an inductively supported axiom will be $AD^{L(\mathbb{R})}$, i.e. the axiom of determinacy restricted to the smallest inner model for ZF that contains the reals, or more plausibly perhaps, a principle that implies $AD^{L(\mathbb{R})}$.[35] Should the considerable inductive evidence for $AD^{L(\mathbb{R})}$ be deemed strong enough for it to be adopted as an axiom, then propositions

[31] Of course, p could come to be known via another epistemic route. If p is otherwise deductively knowable and $O \vdash N \leftrightarrow p$, then N could come to be known deductively via deductive knowledge of p and of O; but this is not the case under consideration.

[32] More precisely, the evidence for the principles in O used in the derivation.

[33] As it happens, a proof that this number is a Mersenne prime also exists.

[34] In some cases, inductive evidence can be combined to form deductive evidence. For example, a proof Π_1 of p_1 may be inductive evidence for $p_1 \wedge p_2$, and similarly a proof Π_2 of p_2 may be inductive evidence for $p_1 \wedge p_2$. Taken together, however, the two proofs Π_1 and Π_2 are deductive evidence for $p_1 \wedge p_2$. Similarly, checking odd cases provides inductive evidence for a number-theoretic conjecture, as does checking even cases; putting the two together yields a proof. See (Baker [2008], p. 338) for a version of this point.

[35] Maddy ([2011], pp. 47–51, 126–31) summarizes the case for the extrinsic justification of $AD^{L(\mathbb{R})}$. She stresses the primacy of extrinsic over intrinsic justification more generally ([2011], pp. 134–7).

deduced from it will be known on inductive grounds, even if they are deduced logically from an axiom. Another way of putting the point is that it is insufficient for a principle to be known deductively that it should be both known and an axiom. ZFC's infinity, replacement, and choice axioms are cases in point; and if $AD^{L(\mathbb{R})}$ were adopted as an axiom tomorrow because of the considerable inductive evidence behind it, our justification for it would not turn from inductive to deductive overnight.

It does not follow from this argument that the justification for axiom N will forevermore be inductive. Perhaps its justification changes over time. We must also distinguish between accepting N and elevating N to the status of axiom. If N is already known, it does not follow from the fact that there are extrinsic reasons for elevating N to axiomhood that the proposition p, which follows from O and N, is known inductively. Our claim is that we know some statements that are axioms inductively, not merely that axiom status is bestowed upon some already-known statements for inductive reasons.

7 Physical to Mathematical Linkages

We can derive knowledge of mathematical propositions from knowledge of related physical ones. Let us examine an elementary geometric case in some detail. Suppose you draw a triangle on a piece of paper and construct perpendicular bisectors on its sides. You then notice that the three constructed lines meet at a point (Figure 1).

You repeat the exercise with many different triangles and observe the same fact. You try triangles of different sizes, of different angles (in particular, triangles with an obtuse angle, right-angled triangles, and triangles in which all the angles are acute), different orientations, and so on. You use a very fine pencil to draw the various line segments; you might even exploit technology to make sure the lines are accurately drawn. On the basis of these 'experiments', you come to believe with great confidence that the perpendicular bisectors of any planar triangle meet at a point.

Let us imagine that you did not know the related mathematical fact before performing the paper-and-pencil experiments. Imagine further that you are sensitive to the distinction between physical and mathematical space. You appreciate that your drawings are not part of the Euclidean plane. What you believe, correctly, is that they are approximately isomorphic to figures in this plane, a relationship you symbolize as:

$$\text{Triangles drawn on paper} \cong_a \text{Triangles in Euclidean plane } (*).$$

(*) is supposed to capture a two-way relationship (since approximate isomorphism is symmetric): any physical triangle drawn on your piece of paper is approximately isomorphic (with respect to a certain class of geometric

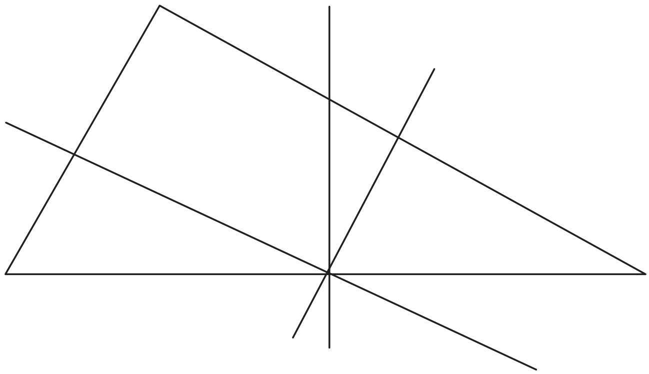

Figure 1. A triangle's perpendicular bisectors meet at a point.

properties) to a Euclidean triangle; conversely, any Euclidean triangle is approximately isomorphic (with respect to the same properties) to a physical triangle.[36] It follows from (*) that if all drawn triangles have their perpendiculars bisectors meet in a point, then all Euclidean triangles do too. Suppose that you know (*); that you know the facts about your drawings empirically; and that you combine these two facts to infer the corresponding mathematical fact about triangles in the Euclidean plane. You have thereby gained inductive knowledge of a mathematical fact.

It would be hard to deny that you can come to know the generalization about physical triangles in this way, assuming that your diagrams are drawn with sufficient care and that you have taken all the relevant variables into account. If you cannot come by such knowledge under such propitious circumstances, that does not bode well for knowledge of other physical generalizations, for example, that if you drop an object it will fall. Note that the physical generalization in question is about all physical triangles drawable on your sheet(s) of paper rather than the more open-ended and vulnerable one about all physical triangles in spacetime. We could in any case avoid generalization if we wish, basing the argument on the fact that a given physical triangle of specified side lengths has the stated property.

You could argue that since the physical generalization is exact rather than approximate, you cannot come to know it via experimentation because experiments cannot give you more than approximate knowledge. But by parity of reasoning that would imply that you cannot know an exact relationship between variables in a physical equation, which would be a devastating consequence for physics and for the possibility of exact scientific knowledge more

[36] What the relevant geometric properties are exactly need not detain us. If precision is sought, take them as the properties invariant under translations, rotations, reflections, and (non-zero) scalings.

generally. Anyway, we could change the example to an approximate statement rather than a precise one. For instance, if the two medians, AA' and BB', of triangle ABC (where A' and B' are the midpoints of BC and AC) meet in a point, G, then the length of AG is roughly twice the length of GA'. (In the Euclidean plane the ratio is exactly two.) In sum, inductive knowledge of physical facts about triangular drawings does not seem particularly problematic, *modulo* fairly strong scepticism about the possibility of empirical knowledge.

It would also be hard to deny knowledge of the approximate isomorphism claim. After all, we exploit mathematical–physical approximate isomorphisms all the time to gain knowledge of the physical world from mathematical models. The first step in such uses is to argue that a physical structure is exactly or approximately isomorphic to a mathematical one; the second is to reason about the mathematical structure to derive a mathematical conclusion; the third is to infer the analogous conclusion about the physical structure, this time exploiting the exact or approximate isomorphism in the opposite direction than was used in our example, that is, in the mathematical-to-physical direction at this last stage. (I am not saying that all applications of mathematics to the physical world are of this simple form; only that some are.) It would be a devastating blow to the use of mathematics to gain knowledge of the physical world, surely a successful enterprise, if approximate isomorphisms such as (*) were not knowable.

A sceptic could of course grant that every physical triangle is approximately isomorphic to a Euclidean one, thus allowing for applications, but question the fact that every Euclidean triangle is approximately isomorphic to a physical one, which is what we are relying on in our example. One response to this sceptic would be to marshal our evidence for (*), which supports both directions of the approximate isomorphism. If the class of relevant properties is invariant with respect to (non-zero) scaling, it is hard to see what could be the basis of such scepticism. Another response is to modify the conclusion: the mathematical proposition the physical experiments allow us to know inductively is that some Euclidean triangles—not necessarily all—have (or approximately have) the perpendicular bisector property.

The particular example of a triangle's perpendicular bisectors meeting at a point is elementary in the extreme and admits of a trivial proof.[37] One might question whether you could realistically use (*) to come to know this fact, since, it might be said, to know (*) you must have tried out the approximate

[37] Let X be the point where the perpendicular bisector of AB and AC meets. Then by the property of perpendicular bisectors, X is equidistant from A and B as well as equidistant from A and C. Hence X is equidistant from B and C and therefore also lies on the perpendicular bisector of B and C.

isomorphism on this kind of case in the first place. I doubt that knowledge of the perpendicular bisector property is a prerequisite for knowing (*). In any case, one can pick a less obvious example, say the existence of the Feuerbach nine-point circle.[38]

You might argue that perhaps space is not Euclidean, in which case the isomorphism would not obtain. But (*) is an approximate isomorphism, and only requires that physical space be (approximately) locally Euclidean, which is both true and something we have good reason to believe. Moreover, that the three perpendicular bisectors of a triangle meet in a point is a property of spherical and hyperbolic, as well as Euclidean, geometry.

Finally, knowledge of the truth-preservation of the logical steps used to derive the mathematical belief from belief in the physical generalization and (*) is no more problematic than any other form of logical knowledge. We conclude that the described experiment does provide an inductive route to knowledge of a geometric proposition.

This example and others like it do not prejudge the issue regarding any conjecture of research interest. Kepler's 1611 sphere-packing problem was only resolved in 1998 by Thomas Hales, who confirmed that 'shot-pile packing' (in which the next layer of spheres is placed in the gaps between spheres in the preceding layer) is the most efficient method. Although this fact is broadly familiar to anyone who has ever repeatedly stacked oranges or cannonballs, I am not claiming that grocers and artillerists have known for a long time what Hales only recently proved.[39]

Exploiting physical knowledge to gain knowledge of mathematics is not limited to geometry. Similar examples can plausibly be found in many other areas of mathematics. One might count physical objects to obtain arithmetical knowledge, exploiting both the existence of an isomorphism between some set of physical objects and a finite initial segment of the natural numbers (the N physical objects and the numbers $\{1, 2, \ldots, N\}$) and empirical knowledge of the physical objects. For example, one might empirically determine of some physical objects: (1) that they are arranged in R rows and C columns, and (2) that there are N of them. Exploiting the physical-to-arithmetical isomorphism

[38] That is, the circle passing through the three feet of the perpendiculars from a vertex to the opposite side, the three midpoints of the sides and the three midpoints of the line from a vertex to the orthocentre.

[39] Many physicists would take the view that some important mathematical facts are known for physical reasons, and that proving them would be little more than a cleaning-up exercise. Here's Witten on Yang-Mills theory and the mass gap hypothesis:

> A mathematical proof that quantum Yang-Mills theory exists in four dimensions would be a milestone in coming to grips mathematically with twentieth century theoretical physics. The reaction of physicists, however, would be that with the renormalization group and asymptotic freedom, one already understands why this theory exists, and that mathematicians have merely succeeded in supplying the ε's and δ's. (Witten [2002], p. 25)

just mentioned would constitute an empirical route to the pure-mathematical fact that $R \times C = N$. Although this example gives us a quick argument for this article's main thesis, the rest of our discussion shows that this is not an isolated instance, that we can have inductive knowledge of significant mathematical generalizations, and that a good deal of knowledge of mathematics is already inductive. Another empirical method might be to use coin tosses, die throws, and Monte Carlo experiments or other simulation methods to obtain probabilistic knowledge (about means and higher-order moments of various distributions say). At least in some such cases, knowledge of the appropriate exact or approximate isomorphism combined with knowledge of the appropriate physical facts yields knowledge of the relevant mathematical proposition.

8 Conclusion

Four arguments were presented for thinking that non-deductive knowledge of mathematics is both attainable and attained. Moreover, we saw in Section 6 that some of what passes as deductive knowledge in set theory is in fact inductive. To finish, some disclaimers are in order.

First, the possibility and actuality of inductive knowledge of mathematical propositions is consistent with the claim that in the great majority of cases in mathematics we only come to know that *p* when we find a proof of *p*. Second, inductive knowledge may be more easily attained in some branches of mathematics than others. A task for a sequel is to examine arguments against the possibility of inductive knowledge in specific areas or obtained using certain techniques. Third, nothing we have said settles cases of current research interest. That inductive knowledge of mathematics is possible creates the conceptual space for an argument to the effect that we already know GC or the Riemann hypothesis on the basis of current evidence—but no such argument has been mooted here. A natural next step would be to consider, for various inductively supported mathematical conjectures *p*, whether *p* is in fact already known on the basis of the available evidence, whether more evidence is required before *p* can be inductively known, or whether for some reason *p* is in principle inductively unknowable. Fourth, we are not denying the truism that one of the fundamental aims of mathematics is proof. To be sure, mathematicians prefer deductive proof and actively look for it even in the presence of overwhelming inductive evidence. As Frege ([1953], Section 2) put it, 'It is in the nature of mathematics always to prefer proof, where proof is possible, to any confirmation by induction'. The ends of mathematical activity are not just to increase knowledge, but to achieve it with (a certain kind of) proof.[40]

[40] See (Rav [1999]) for more on the role of proof in mathematics.

Nothing here directly challenges mathematics' elevation of proof as an end in itself, even if we recognize that proof sometimes starts from inductively supported principles. The aim of the mathematical game, so to speak, may be to prove conjectures—and to do so in ways that suggest further interesting developments (though that is another story)—even if knowledge of some conjectures is non-deductively attainable.[41]

Finally, what are the implications for the practice of mathematics? One natural reaction is that the possibility and indeed actuality of inductive knowledge of mathematics suggests that mathematics should be significantly revised.[42] To draw a parallel with physics, physics until the late nineteenth century was a unified discipline. By the early twentieth century, the physics community had split in two: those who speculated about theory (theoretical physicists) and those who verified the speculations (experimental physicists). Jaffe and Quinn ([1993]) urge that the bifurcation of mathematics into those who offer rigorous proofs and others who offer less rigorous arguments should continue, subject to significant provisos. Our conclusion could be used as grist to this mill, to support further bifurcation in mathematics.

To flesh this out a little further, consider how widespread acceptance of my argument might affect the type of propositions mathematicians try to prove. Suppose that X and Y are propositions with inductive support but that X is inductively known whereas Y isn't. Other things being equal, mathematicians might then focus on proving X or proving statements of the form 'if X then p' rather than on proving Y or statements conditional on Y. Mathematicians need good open problems and granting bodies support good projects. The difference between a proposition being known but not yet proved and its being unknown (and *a fortiori* unproved) might in some cases be the difference between a better and worse research project.

A more conservative reaction is also possible, however. Inductive knowledge of mathematics is attainable and indeed attained, as we have seen. But perhaps the aims of collective mathematical epistemology are best served if the community carries on thinking of proof as the only means of acquiring knowledge of mathematics. Perhaps more theorems are proved under the status quo, better mathematicians trained, and qualitatively better work produced overall than would be under an alternative epistemic regime. Perhaps the health of mathematics might even depend on mathematicians wrongly and dogmatically insisting that proof is the only means to knowledge of a mathematical proposition.

In sum, that knowledge of mathematics is achievable without proof does not necessarily imply that mathematicians should place less emphasis on proof.

[41] It is compatible with our argument that inductive knowledge of mathematics should not be called *mathematical* knowledge, i.e. the sort of knowledge one aspires to in mathematics.

[42] Echoing Fallis ([1997], p. 166).

Whether or not our conclusion can be used to argue for a greater bifurcation between deductive and so-called experimental mathematics depends on broader features of collective mathematical epistemology, features merely alluded to here. What should be clear at any rate is that much is at stake.

Acknowledgements

Thanks to Alan Baker, Dani Rabinowitz, Tim Williamson, and several of this journal's referees for comments on a written version; to Dan Isaacson, James Studd, and Peter Smith for discussion; and to various participants at a seminar in Oxford, at the Philosophy of Science conference at the IUC in Dubrovnik, at a Philosophy of Science Workshop in Bergen organized by Sorin Bangu, and at the Foundations of Mathematics conference in Cambridge. Thanks above all to the Mind Association, whose award of a Research Fellowship in the first half of 2012 enabled me to write this article.

Wadham College
Oxford University
Oxford OX1 3PN, UK
alexander.paseau@philosophy.ox.ac.uk

References

Armstrong, D. M. [1973]: *Belief, Truth, and Knowledge*, Cambridge: Cambridge University Press.

Baker, A. [2008]: 'Experimental Mathematics', *Erkenntnis*, **68**, pp. 331–44.

Bird, A. [2007]: 'Justified Judging', *Philosophy and Phenomenological Research*, **74**, pp. 81–110.

Borwein, P., Choi, S., Rooney, B. and Weirathmueller, A. [2008]: *The Riemann Hypothesis: A Resource for the Afficionado and Virtuoso Alike*, New York: Springer.

Brown, J. R. [2012]: *Platonism, Naturalism, and Mathematical Knowledge*, Oxford: Routledge.

Burge, T. [1993]: 'Content Preservation', *Philosophical Review*, **102**, pp. 457–88.

Casullo, A. [1988]: 'Revisability, Reliabilism, and *A Priori* Knowledge', *Philosophy and Phenomenological Research*, **49**, pp. 187–213.

Clark, M. [1963]: 'Knowledge and Grounds: A Comment on Mr. Gettier's Paper', *Analysis*, **24**, pp. 46–8.

Crowe, M. [1988]: 'Ten Misconceptions about Mathematics', in W. Aspray and P. Kitcher (*eds*), *Minnesota Studies in the Philosophy of Science, 11: History and Philosophy of Modern Mathematics*, Minneapolis: University of Minnesota Press, pp. 260–77.

Davis, P. [1998]: 'Fidelity in Mathematical Discourse: Is One and One Really Two?', in T. Tymoczko (*ed.*), *New Directions in the Philosophy of Mathematics*, Boston: Birkhäuser, pp. 163–75.

Davis, W. and Jäger, C. [2010]: 'Reliabilism and the Extra Value of Knowledge', *Philosophical Studies*, **157**, pp. 93–105.

Dummett, M. [1973]: 'The Philosophical Basis of Intuitionistic Logic', in his *Truth and Other Enigmas*, London: Duckworth, pp. 215–47.

Easwaran, K. [2009]: 'Probabilistic Proofs and Transferability', *Philosophia Mathematica*, **17**, pp. 341–62.

Echeverría, J. [1996]: 'Empirical Methods in Mathematics. A Case-Study: Goldbach's Conjecture', in G. Munévar (*ed.*), *Spanish Studies in the Philosophy of Science*, Dordrecht: Kluwer, pp. 19–55.

Fallis, D. [1997]: 'The Epistemic Status of Probabilistic Proof', *The Journal of Philosophy*, **94**, pp. 165–86.

Fallis, D. [2002]: 'What Do Mathematicians Want? Probabilistic Proofs and the Epistemic Goals of Mathematicians', *Logique et Analyse*, **45**, pp. 373–88.

Franklin, J. [1987]: 'Non-deductive Logic in Mathematics', *The British Journal for the Philosophy of Science*, **38**, pp. 1–18.

Frege, G. [1953]: *The Foundations of Arithmetic*, *trans. by* J. L. Austin, London: Basil Blackwell.

Gettier, E. [1963]: 'Is Justified True Belief Knowledge?', *Analysis*, **23**, pp. 121–3.

Gödel, K. [1964]: 'What is Cantor's Continuum Problem?', in S. Feferman, J. W. Dawson, Jr., S. C. Kleene, G. Moore, R. Solovay and J. van Heijenoort (*eds*), *Kurt Gödel: Collected Works*, Volume 2, New York: Oxford University Press, pp. 254–70.

Goldman, A. [1967]: 'A Causal Theory of Knowing', *The Journal of Philosophy*, **64**, pp. 357–72.

Goldman, A. [1986]: *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.

Harman, G. [1973]: *Thought*, Princeton: Princeton University Press.

Hersh, R. [1979]: 'Some Proposals for Reviving the Philosophy of Mathematics', *Advances in Mathematics*, **31**, pp. 31–50.

Hilbert, D. [1967]: 'On the Infinite', in J. van Heijenoort (*ed.*), *From Frege to Gödel*, Cambridge, MA: Harvard University Press, pp. 367–92.

Hilbert, D. [1967]: 'The Foundations of Mathematics', in J. van Heijenoort (*ed.*), *From Frege to Gödel*, Cambridge, MA: Harvard University Press, pp, 464–79.

Irvine, A. D. [1989]: 'Epistemic Logicism and Russell's Regressive Method', *Philosophical Studies*, **55**, pp. 303–27.

Jackson, J. C. [2009]: 'Randomized Arguments are Transferable', *Philosophia Mathematica*, **17**, pp. 363–8.

Jaffe, A. and Quinn, F. [1993]: '"Theoretical Mathematics": Toward a Cultural Synthesis of Mathematics and Theoretical Physics', *Bulletin of the American Mathematical Society*, **29**, pp. 1–13.

Krantz, S. G. [2011]: *The Proof is in the Pudding*, New York: Springer.

Kumchev, A. V. and Tolev, D. I. [2005]: 'An Invitation to Additive Prime Number Theory', *Serdica Mathematical Journal*, **31**, pp. 1–74.

Lakatos, I. [1962]: 'Infinite Regress and the Foundations of Mathematics', *Aristotelian Society Supplementary Volume*, **36**, pp. 155–84.

Lakatos, I. [1976]: 'A Renaissance of Empiricism in the Recent Philosophy of Mathematics?', in J. Worrall and G. Currie (*eds*), *Mathematics, Science, and Epistemology: Philosophical Papers, Volume 2*, Cambridge: Cambridge University Press, pp. 24–42.

Lecat, M. [1935]: *Erreurs de Mathématiciens des origines à nos jours*, Brussels and Louvain: Castaigne and Desbarax.

Maddy, P. [1988]: 'Believing the Axioms', *Journal of Symbolic Logic*, **53**, pp. 481–511, 736–74.

Maddy, P. [1990]: *Realism in Mathematics*, Oxford: Oxford University Press.

Maddy, P. [2011]: *Defending the Axioms: On the Philosophical Foundations of Set Theory*, Oxford: Oxford University Press.

McGlynn, A. [2012]: 'Justification as "Would-Be" Knowledge', *Episteme*, **9**, pp. 361–76.

Nathanson, M. B. [1996]: *Additive Number Theory: The Classical Bases*, New York: Springer.

Nozick, R. [1981]: *Philosophical Explanations*, Cambridge, MA: Harvard University Press.

Paseau, A. [2008]: 'Fitch's Argument and Typing Knowledge', *Notre Dame Journal of Formal Logic*, **49**, pp. 153–76.

Paseau, A. [2011]: 'Mathematical Instrumentalism, Gödel's Theorem, and Inductive Evidence', *Studies in the History and Philosophy of Science*, **42**, pp. 140–9.

Paseau, A. [forthcoming]: 'What's the Point of Complete Rigour?'.

Rabin, M. [1980]: 'Probabilistic Algorithm for Testing Primality', *Journal of Number Theory*, **22**, pp. 128–38.

Rav, Y. [1999]: 'Why Do We Prove Theorems?', *Philosophia Mathematica*, **7**, pp. 5–41.

Russell, B. [1907]: 'The Regressive Method of Discovering the Premises of Mathematics', in D. Lackey (*ed.*), *Essays in Analysis*, New York: George Braziller, pp. 272–83.

Russell, B. [1910]: 'The Theory of Logical Types', *Essays in Analysis*, New York: George Braziller, pp. 215–55.

Soames, S. [2010]: *Philosophy of Language*, Princeton: Princeton University Press.

Smith, M. [2010]: 'What Else Justification Could Be', *Noûs*, **44**, pp. 10–31.

Steiner, M. [1975]: *Mathematical Knowledge*, Princeton: Princeton University Press.

Sturmfels, B. [2008]: '"Mathematicians Have a Different Perspective"—An Interview with Bernd Sturmfels', *Mathematical Association Focus*, (January 2008), pp. 4–7.

Sutton, J. [2007]: *Without Justification*, Cambridge, MA: MIT Press.

Williamson, T. [2000]: *Knowledge and Its Limits*, Oxford: Oxford University Press.

Witten, E. [2002]: 'Physical Law and the Quest for Mathematical Understanding', *Bulletin of the American Mathematical Society*, **40**, pp. 21–9.

Zermelo, E. [1967]: 'Investigations in the Foundations of Set Theory I', in J. van Heijenoort (*ed.*), *From Frege to Gödel*, Cambridge, MA: Harvard University Press, pp. 199–215.