Mathematical knowledge is context-dependent

Löwe, B.; Müller, T.

Publication date
2008

Published in
Grazer Philosophische Studien

MATHEMATICAL KNOWLEDGE IS CONTEXT DEPENDENT

BENEDIKT LÖWE AND THOMAS MÜLLER

> 'But a proof is sometimes a fuzzy concept,
> subject to whim and personality.'
>
> Kenneth Chang, **New York Times**
> *(April 6, 2004)*

1. Introduction

Mathematical knowledge appears to be of a special, privileged form. When somebody knows a mathematical fact, we say that she knows 'with mathematical certainty', and it is commonly assumed that nothing can be more firmly grounded than that. Not surprisingly, in philosophical contexts, mathematics is often used as an epistemological rôle model. Mathematical knowledge is assumed to be absolute and undeniably firm. The main reason for that special status lies in the fact that mathematicians *prove* their theorems: Mathematical knowledge is proven knowledge. What has been proven is established beyond all doubt. Thus, mathematical knowledge stands out as a uniform, privileged form of knowledge.

Or so it would seem. The main thesis of our paper is that like other forms of knowledge, mathematical knowledge needs to be understood in a context dependent way. Whether a given epistemic subject knows that P (for some mathematical proposition P) depends on specific requirements set by context. These requirements will be linked to a context dependent notion of proof and to the subject's mathematical skills.

Our thesis has a number of important consequences. For one thing, showing even mathematical knowledge to be context dependent can be seen as yet another, strong argument in favour of contextualism in general. Another consequence is that epistemologically, mathematics turns out to be much closer to the other sciences than many philosophers (or mathematicians, for that matter) would think.

A previous version of this paper was given at the conference 'Degrees of Belief', Konstanz, 24 July 2004. We like to thank the participants for the fruitful discussion following our presentation.

Before we begin our argumentation, we consider a few methodological issues. We will proceed from a broadly naturalistic assumption: In philosophising about mathematics, mathematical practice must be taken seriously. If certain expressions, such as 'knowledge', or '... knows that...', are used in the mathematical community, then that usage cannot be dismissed without good arguments. This is not to say that mathematical practice has the last word—but it certainly has to have the first word.

Thus, we will not be satisfied with an epistemology for mathematics according to which there is no (or hardly any) mathematical knowledge in the world—mathematical practice asserts that on the contrary, there is a lot of mathematical knowledge. As we will point out below, a number of seemingly innocuous assumptions about mathematical knowledge and proof quite easily lead to such a conclusion, and it will therefore be important to avoid these assumptions. On the other hand, we will also not be willing to accept an epistemology that identifies all true mathematical statements (as the necessary proposition $2+2=4$ in disguise), and that consequently grants that every epistemic subject knows all mathematical truths. There is certainly less mathematical knowledge than that! Someone supporting such an epistemology would claim that the difference between knowing that $2+2=4$ and knowing that if for some n, $x^n+y^n=z^n$ for all natural numbers x, y, z, then $n \leq 2$ (Fermat's Last Theorem) is a mere difference of *recognition*, not of *knowledge*. However, this view is deeply at odds with the way we talk about knowledge of mathematical statements. We normally say 'I didn't know that P', but we can hardly say 'Oh, I hadn't recognized that P is just $2+2=4$ in a different form', or 'I hadn't seen that P means that $2+2=4$'.

With respect to contextualism (for details, cf. Section 2), our methodology has another important consequence. A number of philosophers, *e.g.*, [Scha04], concede the context dependence of knowledge *attributions* (quite an obvious linguistic fact), but dispute the context dependence of knowledge itself. In order to support such a view, pragmatic rules are cited according to which in some contexts it is just inappropriate to say that S knows that P even though S does know that P, or in which it is appropriate to say that S knows that P even though S does not know that P. Effects of mismatch between semantics and pragmatics are indeed common. *E.g.*, it is true, but not normally appropriate to say, that the earth has less than 19 moons. These pragmatic effects may also play a rôle in some epistemic situations—*e.g.*, it is not normally appropriate to say that somebody knows that the earth has less than 19 moons, either. However, with respect to the question of knowledge attributions in mathematics, a non-contextualist reading has to specify a rather narrow criterion for knowledge that applies to much fewer cases than actual usage suggests. The pragmatic explanation will then have to claim that these are all cases of S not knowing that P in which it is still pragmatically adequate to say that S knows that P. Such cases are possible, but most examples of that kind of pragmatic/semantic mismatch are so-called white lies. *E.g.*, it may be appropriate to say to a nervous speaker after his talk that this was a nice talk even if one doesn't really think so. However, this is surely an odd account of what goes on in mathematical knowledge attributions generally. Our methodology urges us to take such statements at face value unless there are good reasons to the contrary—and such reasons had better not rest on a distaste for contextualism. In mathematical practice (*e.g.*, in exams), people argue about what a certain subject *knows*, not what *it would be appropriate to say* about a certain subject. (Try telling an examinor 'Of course I don't know, but it would be appropriate to say that I do, so I deserve an A'.) Thus context dependence of knowledge attribution in mathematics, given our methodology, immediately supports our thesis of context dependence of mathematical knowledge itself.

2. Contextualist epistemology: Lewis's analysis

Contextualism is a fairly recent attempt at answering one of the long-standing problems of epistemology, viz., the problem of skepticism. In spelling out contextualism, we follow David Lewis's general analysis, given in his 1996 classic, 'Elusive knowledge'. Lewis tries to find a way of holding on to the at least *prima facie* reasonable claim that knowledge must be infallible while avoiding the skeptical challenge. Infallibility means that a purported knower must have eliminated all possibilities that put her knowledge claim in doubt. However, the skeptic can always point to some far-fetched possibilities which the purported knower has overlooked: CIA plots, brains-in-a-vat scenarios, a deceiving demon. Thus, the purported knower cannot sustain her infallibility claim—and accordingly, there is no knowledge. As arguments of that type apply to all knowledge claims, it would seem that there is no knowledge at all. But this is absurd; we know a lot. Lewis counters the skeptic's moves by pointing out that these moves amount to changing the *context* of the knowledge claim. His own, context dependent analysis of 'S knows that P' is as follows:

(#) S *knows* that P iff S's evidence eliminates every possibility in which not-P—Psst!—except for those possibilities that we are properly ignoring [Lew96, 554].

The option of 'properly ignoring possibilities' allows for a spectrum of knowledge contexts from the loose standards of every-day usage (in which, *e.g.*, I know that my cat Possum is not in the study without checking the closed drawers; cf. [Lew96, 562]) to the demanding standards of epistemology (Cartesian Doubt), in which (almost) all knowledge claims are defeated. Consequently, a switch of context may destroy knowledge. This, so Lewis, both explains the force of skeptical arguments and points a way to a cure.

Lewis's paper and a number of other related works gave rise to a huge debate about details and technicalities of his version of contextualism, dealing with important questions about the specification of 'properly ignoring possibilities' and the context changes in communicative acts. This paper is not intended to be a direct contribution to that debate.[1] We rather assume that contextualism has a stable core, and we will try to employ that core notion to mathematics. Our emphasis will be on determining how a context for knowledge claims can be specified in mathematics, and how such contexts can be related to the mathematically central notion of proof. We will follow Lewis's outline of 'eliminating adverse possibilities' and 'properly ignoring possibilities' to fix the terminology. As to details of the debate, to put our cards on the table, we favour the approach to contextualism put forward by MacFarlane [MacF05], who points to a number of important phenomena connected with nesting knowledge claims and other modal operators. As we will treat stand-alone sentences only, we will not need to discuss the intricacies of the distinction between context of use, context of evaluation, and context of assessment here.

Contextualism has not been employed in the epistemology of mathematics so far. There is certainly a number of reasons why this is so. For Lewis, the main reason seems to be that he treats all true mathematical statements as the necessary proposition in disguise, thus blocking any way of distinguishing among them epistemologically. This is a consequence of Lewis's modal epistemology: A semantics for knowledge claims for Lewis must be based on possible worlds. As all mathematical statements are true in all possible worlds, modal semantics must treat all mathematical statements as the necessary proposition,

[1]Important recent contributions include [MacF05], [Scha04], and [DeR02]. Cf. also the April 2005 issue of *The Philosophical Quarterly*.

modelled as the set of all worlds. As we pointed out above, given our methodology, we cannot follow Lewis here. (Incidentally, in his [Lew93, 218], Lewis supports something very close to our methodology, so there may be a slight tension in his position.) It seems obvious to us that Lewis's modal approach to epistemology can be separated from his contextualist stance, and thus we will employ a contextualist analysis of knowledge along Lewisian lines, but with respect to finely individuated mathematical statements.

How finely? This is a good question, and we do not have a definite answer. However, from mathematical practice we learn that while we must not individuate too coarsely ($2+2=4$ is different from Fermat's Last Theorem), we must not individuate too finely either: One and the same mathematical statement can be expressed in various ways, *e.g.*, in different languages and with or without the explicit use of formal symbolism, and one and the same statement can have a number of different proofs. Any way of individuating mathematical statements that honours these constraints will be fine with us. We are not after a formal theory here.

3. Standard Mathematical Epistemology

In this section we give an outline of a standard view of knowledge in mathematics that, we claim, both plays an important rôle in philosophical debates and is held by the educated public:

($\star$) S knows that P iff S has available a proof of P.

Our final view, (‡), in some way still honours many of the intuitions behind ($\star$). However, as it stands, ($\star$) is open to a number of criticisms. Firstly, it is vague with respect the two key notions of the *explanans* ('proof' and 'having available'). Secondly, it does not cover all cases of knowledge attributions in mathematical practice: as we will argue in section 4.3, there are cases of mathematical knowledge in which availability of proof is out of the question. Thirdly and most importantly, attempts at removing the vagueness of ($\star$) usually proceed to fix the notions of 'proof' and of 'having available' in a uniform way, leading to an invariantist reading of ($\star$):

($\star^\sharp$) S knows that P iff S has available$^\sharp$ a proof$^\sharp$ of P.

Here, 'has available$^\sharp$' and 'proof$^\sharp$' are sharp, fixed notions. As we will show, such a reading leads to theories that run counter to our methodological presuppositions. The view sketched here thus provides the basis for our more detailed argumentation in section 4 below, through which we will establish and substantiate our thesis of context dependence of mathematical knowledge as a viable alternative to $(\star^\sharp)$.

As mentioned in the introduction, it has often been observed that compared to the other sciences, mathematics is an 'epistemic exception':[2] Whereas empirical claims have relative support through other empirical claims, mathematical claims admit of proof, and as we learn early on in our education, a mathematical proof is either correct or incorrect. A proven mathematical statement is beyond doubt in a way that even a well-supported empirical claim is not: The notion of proof sets an absolute standard of support for a mathematical claim. Furthermore, proofs do not admit degrees of correctness. To use Keith Devlin's polemic words: 'Surely, any math teacher can tell in ten minutes whether a solution to a math problem is right or wrong! [...] Come on folks, it's a simple enough question. Is his math right or wrong?' [Dev03][3]

This widespread belief in the objectivity of a notion of mathematical proof gives intuitive support to a context independent reading à la $(\star^\sharp)$. Such a position concerning mathematical knowledge is implicitly or explicitly shared by a large number of philosophers ancient and modern. *E.g.*, Plato in the famous *paîs* example (*Meno*, 82b–84a) shows how the slave, guided by Socrates, has available a proof of the Pythagorean theorem and thus, mathematical knowledge. Locke, in discussing knowledge in the *Essay* [LE], uses mathematics as a prime example. While he makes a number of observations based on actual mathematical practice, *e.g.*, on difficulties of long deductions (esp. in Bk. IV, Chap. ii, which is entitled 'Degrees of Knowledge'), he also subscribes to $(\star)$ in general. Kant, who holds that mathematical truths are synthetic *a priori*, limits the use of 'knowledge' generally to proven certainties and even claims that mere belief has no place in mathematics at all (*Critique of Pure Reason* [KCPR], A 823/B 851). Frege's project of logicism, i.e., grounding all of mathematics on pure logic alone, makes explicit formal proof the hallmark of secure knowledge (*Begriffsschrift* [FBS] IX f.). And to bring in a contemporary contextualist, Lewis also seems to endorse an invariantist view like $(\star^\sharp)$ when he remarks in passing that in 'the mathematics department, [...] they are in confident agreement about what's true and how to tell, and they disagree only about what's fruitful and interesting' [Lew00, 187f.].

[2]This has been an important topic in the sociology of science, discussed, *e.g.*, by Bloor [Blo76] and Livingston [Liv86]. Cf. also [Hei00, Chap. 1].

[3]Just for the record: Of course, Devlin is playing the *advocatus* here, arguing that even checking proofs is not as trivial as is often believed.

Arguably even more important is the fact that an invariantist reading of the standard view is deeply entrenched in the image of mathematics in the educated public.

It is no coincidence that the standard view is so widespread. In a way it expresses, in a form suitable for mathematics, just the intuitions that stand behind the classical tri-partite analysis of knowledge as justified true belief:

(§) S knows that P iff S has the justified true belief that P.

Taking for granted the context independence of ($\star^\sharp$), we get that proof is the only source of knowledge in mathematics, and whether someone has available a proof or not is an objective, indisputable fact. This is a crisp explication of the word 'justified' in (§) in that it describes a unique possibility of justification for mathematical belief: having available a proof.

Now Gettier-style problems[4] show that the classical analysis (§) merits critical discussion. Instead of a context independent notion of justification, we seem to need a context dependent notion of 'appropriate justification'. The received view about mathematics seems to imply that mathematics is immune to these problems, as there is only one 'appropriate justification' in mathematics and that is the objective notion of proof. We shall show that this view is as naïve as the invariantist reading of (§), leading to conflicts with our methodological assumptions. Instead, a context dependent reading of ($\star$) is needed.

4. Mathematical knowledge

In this section we have two main aims. The first is negative: we will criticize the context independent reading of the standard view of mathematical knowledge, ($\star^\sharp$), outlined in the previous section. Our second aim is positive: we will show how contextualism à la Lewis can be spelled out for the case of knowledge in mathematics. To this end, we first give an overview of how the two crucial notions of ($\star$), 'proof' and 'having available', are employed in mathematical practice. This overview strongly supports a contextualist reading of ($\star$).

4.1. **Proof in mathematical practice: derivation and informal proof.** Proof is perhaps the central notion in mathematics. Despite that, it is not exactly clear what a proof is: a wide range of texts is called 'proof' in actual mathematical practice. The guiding idea of

[4]Cf. Gettier's classic [Get63]; for the ensuing discussion, cf., *e.g.*, [Pap78] and [DeR99].

proving something is to arrive at the result through a number of secure steps, but one needs to specify which steps may be used. Frege in his *Begriffsschrift* proposed that the steps should be so small that a mechanical procedure was available for checking each step. This led to a mathematically precise definition of formal proof which was then available for metamathematical investigations leading, *e.g.*, to Gödel's completeness and incompleteness results. We will use the term *derivation* to stand for formal proof in a mathematically well-defined system.[5] Outside metamathematical investigations and a few very specialised areas,[6] one will not find derivations in mathematical publications. Mathematical journals and textbooks (as well as lectures, research notes and conference talks) instead contain *informal proofs*.[7] The notion of informal proof does not have a mathematically precise definition—if it did, it would be just another version of derivation.

From the point of view of derivation, informal proofs contain gaps and appear to be unfinished. It is therefore tempting to see an informal proof just as an imperfect stand-in for a derivation. However, mathematical practice strongly supports the view that the important notion of proof in mathematics is not derivation, but informal proof. One reason for this is communication: 'The point of publishing a proof [...] is to communicate that proof to other mathematicians. [...] [T]he most efficient way [...] is not by laying out the entire sequence of propositions in excruciating detail' [Fal03, 55]. Instead, mathematicians publish informal proofs. However, there is more to informal proof than ease of communication. It just isn't the case that mathematicians have a derivation in mind and transform it into an informal proof for publication in order to reach a wider public—the entire procedure of doing research mathematics rests on doing informal proofs. The proofs in mathematical research papers are so far removed from derivations that only a few experts could produce a derivation from them even if they wanted to, and only a minority considers that a worthwhile enterprise. We need to take seriously the fact that derivations are hardly ever used. Subscribing to the tempting image of the derivations as the real objects

[5]There are various competing notions of derivation, but their differences do not matter for our purposes. For first-order logic, the competing formal systems are equivalent in allowing one to prove exactly the same theorems.

[6]*E.g.*, the *Journal of Formalized Mathematics*, which focuses on derivations in the specific proof system `MIZAR`, *cf.* `http://www.mizar.org/JFM/`; or the publications of the Coq group discussed in section 4.2.

[7]Cf., *e.g.*, [Rav99] for discussion of this distinction.

of mathematical study to which informal proofs are imperfect approximations would be a violation of our maxim of taking mathematical practice seriously.

Informal proofs come in many flavours. One can, *e.g.*, distinguish semi-formal textbook proofs for beginning students, graduate-level textbook proofs, journal proofs, informal research notes, and proof sketches. Each of these types is pragmatically fairly well delineated—try submitting a textbook-style proof to a mathematical journal, or presenting research note-style proofs to beginning students, and you will feel the force of the boundaries.

It is often possible to compare proofs for one and the same P with respect to the level of detail they exhibit. One proof may give more details than another, even though both are valid and complete proofs of certain types. *E.g.*, a textbook proof may contain a whole page of detail for a certain inductive construction where a research note would just say 'by induction'. Thus it makes sense to say of a proof that it contains gaps relative to another. However, we do not subscribe to an absolute notion of 'gaps in proofs', because that would presuppose an absolute standard of a 'gap-less' proof.[8]

If we wish to analyse mathematical knowledge by ($\star$) taking into account many different forms of proofs, we thus have two options. Either we accept that the word 'proof' will be analysed by different types of proofs in different contexts, thus leading to a context dependent reading of ($\star$), or we fix a level of detail and declare the corresponding type of proof$^\sharp$ as constitutive for mathematical knowledge via ($\star^\sharp$). For both options it is necessary to spell out what it means to have available a proof.

4.2. **Having available a proof.** What does it mean to have available a proof? A literal reading in terms of having access to a material copy of the proof is inappropriate. It is too narrow, because there just aren't enough copies of proofs to back even a fraction of true mathematical knowledge claims (especially if one demands derivations, of which there

[8]Note that there is a different notion of gap in proof, which Fallis [Fal03] calls 'untraversed gaps' in contrast to the 'enthymematic gaps' that we just discussed: If in proving one fails to note a certain special case, the proof will be incomplete—it won't even belong to the intended class of informal proof. Here the gap terminology is appropriate in an absolute sense.

are hardly any around).[9] But it is also too wide: A mathematical illiterate on the first floor of UC Berkeley's Evans Hall (the math library) has available lots and lots of proofs, but it would be odd to say that the mere location could affect any change in mathematical knowledge (*genius loci* nonwithstanding).

Thus, 'having available' cannot be spelled out in terms of actual physical access; it needs to be given a modalised reading in which the epistemic subject S plays an active rôle. A reformulation of ($\star$) that makes that modalisation explicit is the following:

(†) S knows that P iff
S could in principle generate a proof of P.

That model, with 'proof' mostly fixed to mean 'derivation', has been advocated by many writers, and we would venture to say that something like (†), quite possibly with the 'derivation' reading of 'proof', is what most people will come up with when pressed to explicate ($\star$).[10]

[9] *E.g.*, almost certainly no living mathematician has seen a derivation of the Feit-Thompson Theorem, yet there are (many) mathematicians who know that every group of odd order is solvable. The original paper, [FeiTho63], has over 250 pages. Only specialists in finite group theory will know even an informal proof, let alone a derivation. On the other hand, the theorem is rather well known.

[10] Derivation is often called the gold standard of mathematical proof. That metaphor is quite telling. First, a few historical facts. Implementing a gold standard means making a fixed weight of gold the standard economic unit of account. This can, *e.g.*, be established by using coins made of gold. More practically, gold is stored in some central reservoir, and paper money is issued as certificates entitling the holder to a fixed amount of gold. Such systems were established in the late 19th century in many Western countries, and there were earlier, similar systems in many places. A positive aspect of an international gold standard is free convertibility of currencies, which was important in boosting international trade. A negative aspect of such a system is that even though gold is nice stuff, what people actually need isn't gold (except in some cases related to dentistry), but other goods, and the scarcity or otherwise of gold dictates in effect the price of other goods. The successor of the early international gold standard, the Bretton Woods system, collapsed in the early 1970ies. Since then, many countries have sold off much of their gold. Other mechanisms of establishing trust between trading parties have proved to be more practical and more efficient.

We would like to draw a rather strict analogy between the rôle of gold for the exchange of goods and the rôle of derivation for the exchange of mathematical knowledge. Historically, of course there never was a period in the development of mathematics during which derivation was the generally accepted currency, but the logicist movement of the early 20th century surely was an attempt at establishing that currency. Just like gold *vs.* goods, derivation is neither the only store of value for mathematics, nor the most useful. If anything, trading in derivations is more impractical than trading in gold. (Given the scarcity of gold and the expanded international trade today, a return to an international gold standard would mean

The notion of 'could in principle generate' implies some type of modal idealisation. Such idealisations are often invoked when it comes to explaining what human beings do when they do mathematics. A classical example is Brouwer's *idealised mathematician*, the creative subject who creates his choice sequences.[11] Kitcher [Ki84, Chap. 6.III], in a similar vein, employs the notion of an 'ideal agent' to account for the fact that actual operations of actual agents do not suffice to establish the truths of arithmetic as he conceives it.[12] Steiner [Ste75, Chap. 3] explicates the modal idealisation of (†) via the following thought experiment: In order to check whether a mathematician has available a proof of P and thus, knows that P, she is asked to transform her (informal) proof into a derivation with the aid of a logician who as a Socratean 'midwife' works out the formal details, but is not otherwise mathematically creative. "If the two can bang out a formal proof, then the mathematician is said to have known the proof all along, on the basis of the informal argument" [Ste75, p. 100]. Thus:

($\dagger_1$) S knows that P iff with the help of a logician, S can generate a derivation of P.

The three cited authors give quite specific readings of modal aspects of mathematics, and Steiner gives an explicit test for 'could in principle generate a proof'. This is what one needs to do if one is after an invariantist version of (†). However, mathematical practice provides counterexamples against any fixed notion—there is even knowledge without proof (cf. section 4.3 below). We will now explore the modal dimension of (†) in three steps, starting with a critique of Steiner's approach.

increasing the current price of gold more than ten-fold. But given the scarcity of derivations, establishing derivations as the sole vehicle of mathematical justification would at present completely stop the development of mathematics.)
Other mechanisms for establishing trust in the mathematical community are well established, and they are working. Of course that does not mean that derivations are worthless. Quite on the contrary—the belief in the possibility to generate, at least in principle, a derivation corresponding to any given informal proof, may well be one of the strongest sources of mutual trust in the mathematics community. It is just that actual derivations aren't really needed—except, if you allow, for exercises in mathematical dentistry.

[11] *Cf.* [Brow29]; for an historical overview of the notion, *cf.* [Troe82].

[12] *Cf.* Chihara [Chi90, Chap. 11.2] for decisive criticism of Kitcher's approach.

(i) Steiner's model ($\dagger_1$) is open to a number of criticisms, some already voiced in the original publication.[13] The envisaged test for knowledge only replaces one form of modalisation ('has available' or 'could in principle generate') with another, not much clearer one ('can produce, with the aid of a logician, ...')—and the kind of logician that is needed may well turn out not to exist. The logician's powers play a crucial rôle. Steiner rightly stresses that "we cannot envision a superhuman, because such a being would discover a completed proof despite the ignorance of the mathematician" [Ste75, pp. 101f.], rendering the test useless. In practice, even if two persons cooperate in producing a derivation, the rôles will never be as clearly delineated as the test suggests. It may be fine to say that *the pair* who succeeded in writing down a derivation had available a proof (and thus, knew that P), but that is of course no good as a test of the *mathematician's* knowledge.

Let us now consider two variants of Steiner's modalisation. In both variants, the dubious logician is replaced by a direct appeal to the subject's capabilities. The first variant is based on derivation, the second, on informal blackboard proof.

(ii) Suppose that we want to read ($\dagger$) without context dependence, and let us fix 'proof' to mean derivation. The task then is to try to find a good explication of 'could in principle generate'. The successes of formalized mathematics have shown that it is possible to provide derivations for many important mathematical statements, however doing so requires a long time: *e.g.*, the Coq community worked for over ten years before Geuvers, Wiedijk, and Zwanenburg were able to formalize the fundamental theorem of algebra [GWZ01]. Now, this suggests reading 'could in principle generate' as follows:

($\dagger_2$) S knows that P iff, given ten years,
she could write a formal derivation in the language Coq.

But compared to these ten years, the time we need to learn mathematical facts is short: many mathematicians could be in the situation that they don't know anything about P, but are able to learn within ten years both the mathematics needed to understand why P is true and then formalize it in Coq. These mathematicians would satisfy our fixed reading of ($\dagger_2$), but by assumption do not know P. For particularly bright beginning students, the time of ten years might be enough to study mathematics, enter graduate school, finish a PhD, and learn

[13] It must be said in fairness to Steiner that he does not it in the end subscribe to the model sketched. Rather, he gives an example of mathematical knowledge without proof and then argues for a Platonist understanding of mathematical knowledge.

Coq. Thus, the reading ($\dagger_2$) would grant almost indefinite mathematical knowledge to everyone who has the intellectual capacity to finish a mathematics degree. Clearly, not an intended reading.

The invariantist readings ($\dagger_1$) and ($\dagger_2$) face another difficulty. As soon as derivations or a system like Coq play a rôle, we need to concede that there was no mathematical knowledge prior to a certain point in time: *e.g.*, before the *Begriffsschrift*, nobody could give a derivation of anything, because the concept of derivation had not yet been invented.[14] But mathematics is commonly taken to be the prime example of historically stable knowledge—the ancient Greeks already *knew* the Pythagorean theorem.

(iii) At the other end of the spectrum let us read (†) without context dependence by fixing 'proof' to mean 'informal proof on the blackboard'. For many research situations in mathematics, the relevant notion of 'could in principle generate' is something like the following:

($\dagger_3$) S knows that P iff, given a blackboard and a piece of chalk, she is able to produce an acceptable blackboard proof within an hour.

The time frame of one hour is important here, as many research mathematicians do not have all of the proofs they need for their work at their immediate cognitive disposal. They need to try one or two standard approaches to tackle the problem, remember the important details, and then are able to provide an acceptable proof. If one keeps this time interval too short, then one arrives at too strict a criterion for knowledge. But the time frame is not uniform across all situations even for fixed P. Consider a student in an oral exam asked whether P or non-P is true. Suppose that the student erroneously believes that non-P is true but given a blackboard, a piece of chalk and one hour of time, this particular student might be able to create a blackboard proof of P, first trying to prove non-P, failing, getting some ideas from the failed attempts, then remembering some facts and ideas from lectures, and finally proving P. However, in the oral exam, the examiner will not wait for an hour, the student has to rely on his belief, says 'non-P' and fails. Does this mean that the oral exam is not testing knowledge? In view of our methodological maxim, that would be absurd.

These examples show that the temporal component in 'having available' is immensely important, and that it seems hopeless to try to fix a

[14] If you are not satisfied with taking the *Begriffsschrift* as the beginning of derivation, supply your favourite reading instead. The consequences are practically the same.

single reading of 'having available' for all contexts. *E.g.*, on a temporal reading of the modalisation, if one gives the subjects too much time to generate a proof, then one ends up with knowledge assertions that shouldn't be true, but if one gives them too little time, then some true knowledge assertions dissolve.

4.3. **Contextualism in mathematics.** Our discussion of the key notions of the standard analysis of mathematical knowledge, ($\star$), has revealed that both 'proof' and 'having available' admit of gradations. For the former notion, there are various notions of proof available, ranging from formal derivation to informal proof sketch. For the notion of 'having available', we have argued that a modalised reading is necessary, as actual availability is far too strict a criterion. We have shown that there are various such modalisations available, *e.g.*, a family of modalisations distinguished by a (counterfactual) time constraint.

If ($\star$) is to be given a context independent reading as in ($\star^\sharp$), then it must be possible to select, at least relative to any mathematical proposition P and/or a subject S, *a single* proof type and *a single* reading of the modalisation (*e.g.*, a single time frame) that could fix the standards for *all cases* of "S knows that P". It is not difficult to see that this is a hopeless enterprise. One will always be able to come up with a scenario in which S fulfills the conditions, and yet doesn't know, or knows despite failing the conditions. *No fixed notion of 'having available a proof' yields an adequate analysis for all cases of mathematical knowledge.*

Actually, things are even worse for the standard view: There are many examples of proper knowledge attributions where (†) is false—we commonly attribute mathematical knowledge to people who are currently in no position to produce a derivation or any other type of proof of mathematical facts that they know.

A good historical example is reported in Polya's study of Euler [Pól54, Chap. 2.6]: It is certainly true to say that Euler *knew* that

$$1 + 1/4 + 1/9 + 1/16 + 1/25 + \ldots = \pi^2/6,$$

but Euler didn't have available (and knew that he didn't have available, nor could in principle generate) a proof of that fact—he had established it via generally shaky generalisations from finite to infinite sums, and his evidence was to a large part inductive (*i.e.*, the first 20 or so decimal places coincided). Still, it would be ahistorical to say that Euler had just guessed.[15]

[15]Cf. [Ste75, Chap. 3.IV] for a similar assessment.

Cases of knowledge without proof are not rare at all, nor are they a thing of the past. In industrialised countries the majority of the public has mathematical knowledge of some kind, *e.g.* elementary algorithms of arithmetic, the Rule of Three, etc., but of course only a tiny fraction of the public would satisfy any reasonable reading of (†).[16] Furthermore, there is mathematical knowledge via testimony, for which proof plays no rôle at all—and yet, in many mathematical contexts it is fine to base a knowledge claim on testimony. That is obvious enough for claims to mathematical knowledge in the general public: Most people haven't actually seen any mathematical proofs at all. For beginning math students, a similar observation holds: While we certainly urge them to try to learn and understand the proofs, we also concede that the students do acquire knowledge (though not a very deep kind of knowledge) by just learning theorems by heart, and that may be enough to pass a first exam. And even in the context of research mathematics, some knowledge is just based on trust. If one works in cooperation with others, it will not normally be possible, nor required, to learn and check all proofs.

No context independent reading of (†) can cover all true knowledge attributions in mathematical practice. Thus, contextualism wins the day. But how? Our task now is to link the general contextualist analysis of knowledge to the specific case of mathematics. Recall Lewis's analysis (#): S knows that P iff S's evidence eliminates all possibilities in which non-P—except for those possibilities that we are properly ignoring. In mathematics, S's evidence and the ignored possibilities must be linked to the proof or other justification that S has available according to (†). As we saw, a context generally specifies a type of proof (or other justification) as appropriate. Very few contexts in mathematics demand derivations. Blackboard proofs are typical of research mathematics, and mathematical knowledge claims in the general public typically do not need to be backed by any form of proof at all. Similarly, S's evidence may be interpreted as the dispositional state of mind of S with respect to the required form of proof of P. Above we gave one explication by linking that disposition to a time frame that would be required to generate the proof in question. Thus one way of writing

[16]Behind these commonsensical examples, a vexing question in the philosophy of mathematics is hidden: Is it possible to have (a high degree of) knowledge of P by pure intuition without any formal proof in mind (the Ramanujan phenomenon)? Cf. [Thu94] for discussion of this point.

out (#) for the case of mathematics is the following:

($\#'$) *S knows* that P iff S's dispositional state of mind allows her to produce the required form of proof or justification for P within the time frame allowed by the context.

This analysis may be all that is needed, but it also comes with a certain problem: There does not seem to be an independent standard from which to assess the necessary time. Thus, ($\#'$) might be accused of being empirically void. We suggest that the notion of mathematical skill can help to improve the analysis.

4.4. **Mathematical knowledge and mathematical skills.** The notion of mathematical skill links the 'dispositional state of mind' of ($\#'$) with actual performance: Skill is both a modal notion (what somebody is able to do even while not doing it) and has an empirical side (skills can be tested). Our motivation for bringing skills into the picture is that through the Dreyfus-Dreyfus model of skill acquisition there is available a semi-formal theory of skill levels that has been fruitfully applied, *e.g.*, to chess skills and nursing skills [Ben84]. In the Dreyfus-Dreyfus model [DreDre86], there are five levels of skill ranging from novice to expert. These levels are delineated by their different relation to explicitly formulated rules. While a novice needs to stick to explicit rules in a step-by-step fashion, experts have internalised and transgressed such rules and are able to proceed intuitively.

Certainly the link between mathematical knowledge and mathematical skills merits further investigation, which will need to be left for a future occasion. Here we merely wish to argue that a skill-based analysis is plausible.

Using the notion of skill, we can reformulate ($\#'$), our preliminary synthesis of contextualism (#) and mathematical knowledge (†), as follows:

(‡) *S knows* that P iff S's current mathematical skills are sufficient to produce the form of proof or justification for P required by the actual context.

This analysis, we claim, is adequate as a general explication of mathematical knowledge. It refers to the actual context and is thus flexible with respect to both crucial aspects of mathematical knowledge: Context determines the required form of proof or other justification, and context also sets the standard for the modal component in terms of a required skill level. Skill levels provide the link of our analysis with independent constraints that was lacking in the case of ($\#'$)—unlike counterfactual time constraints, skill levels can be (and, more importantly: are) characterised independently of any theory of ours. Mathematical practice affirms that the concept of mathematical skill is well entrenched. It is customary to comment on students' or researchers' skills, and it is often possible to rank people with respect to their skills.[17] Skills are tested in exams and job talks, and it may well be that the aim of mathematics education is best characterised not as instilling mathematical knowledge, but as teaching mathematical skills.

5. Conclusion

In this paper we argued that contrary to first appearances, mathematical knowledge is not a fixed, context independent notion. Rather, we showed by appeal to mathematical practice that unless one disregards actual practice—which in our view would be just plain bad methodology—, one is forced to admit that mathematical knowledge is context dependent.

Many accounts of mathematical knowledge refer to the need to have available a proof. We concede that proof plays a crucial role in mathematics and in mathematical knowledge, but there is also mathematical knowledge without proof. Nor is proof a fixed notion: There are various forms of proof, and context determines which type of proof, if proof at all, is required. Furthermore, availability of proof is a modal notion that we suggested is best explained by reference to mathematical skills.

What then of formal derivation? The concept of derivation and its universal acceptance as a formalization of the intuitive notion of proof is important for the foundations of mathematics, but contrary to folklore, it hardly plays any rôle in determining the truth of 'S knows that P'—Psst!—unless the context explicitly demands it.

References

[Ben84] Patricia Benner, *From Novice to Expert: Excellence and Power in Clinical Nursing Practice*. Addison-Wesley 1984.

[Brow29] Luitzen E.J. Brouwer, Mathematik, Wissenschaft und Sprache. *Monatshefte für Mathematik und Physik* 36 (1929), pp. 153–164.

[Blo76] David Bloor, *Knowledge and Social Imagery*. University of Chicago Press 1976

[Chi90] Charles Chihara, *Constructibility and Mathematical Existence*, Oxford UP 1990.

[DeR99] Keith DeRose and Ted A. Warfield (eds.), *Skepticism. A Contemporary Reader*, Oxford UP 1999.

[17] An interesting question which again merits further investigation is the following: How finely do we need to individuate mathematical skills? Will it be enough to ascribe to persons a single 'mathematical skill level', or will we need to be more topic-specific, speaking, *e.g.*, of algebraic *vs.* geometrical skills?

[DeR02] Keith DeRose, Assertion, Knowledge and Context, *Philosophical Review* 111 (2002), pp. 167–203.

[Dev03] Keith Devlin, The shame of it, MAA Online: Devlin's Angle, May 2003.
URL: `http://www.maa.org/devlin/devlin_05_03.html`.

[DreDre86] Hubert L. Dreyfus and Stuart E. Dreyfus, *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*, Free Press 1986.

[Fal03] Don Fallis, Intentional gaps in mathematical proofs, *Synthese* 134 (2003), pp. 45–69.

[FeiTho63] Walter Feit and John G. Thompson, Solvability of groups of odd order. *Pacific Journal of Mathematics* 13 (1963), pp. 775–1029.

[FBS] Gottlob Frege, *Begriffsschrift*. Halle 1879.

[Get63] Edmund L. Gettier, Is justified true belief knowledge? *Analysis* 23 (1963), pp. 121–123.

[GWZ01] Herman Geuvers, Freek Wiedijk, and Jan Zwanenburg, A Constructive Proof of the Fundamental Theorem of Algebra without using the Rationals. In: Paul Callaghan, Zhaohui Luo, James McKinna and Robert Pollack (eds.), *Types for Proofs and Programs, Proceedings of the International Workshop, TYPES 2000, Durham*, Springer LNCS 2277, pp. 96–111, 2001

[Hei00] Bettina Heintz, *Die Innenwelt der Mathematik. Zur Kultur und Praxis einer beweisenden Disziplin*. Heidelberg: Springer 2000.

[KCPR] Immanuel Kant, *Kritik der reinen Vernunft*. 1781 (A), ^{2}1787 (B). Ed. by J. Timmermann. Hamburg: Meiner 1998.

[Ki84] Philip Kitcher, *The Nature of Mathematical Knowledge*. New York 1984.

[Lew93] David Lewis, Mathematics is megethology. *Philosophia Mathematica* 3 (1993), pp. 3–23. Reprinted in his *Papers in Philosophical Logic*, Cambridge UP 1998, pp. 203–229. Page references are to the reprint.

[Lew96] David Lewis, Elusive knowledge. *Australasian Journal of Philosophy* 74 (1996), pp. 549–567. Reprinted in his *Papers in Metaphysics and Epistemology*, Cambridge UP 1999. Page references are to the original.

[Lew00] David Lewis, Academic appointments: Why ignore the advantage of being right? In his *Papers in Ethics and Social Philosophy*, Cambridge UP 2000, pp. 187–200.

[Liv86] Eric Livingston, *The Ethnomathematical Foundations of Mathematics*. London: Routledge 1986.

[LE] John Locke, *An Essay Concerning Human Understanding*, 1689. Ed. by P.H. Nidditch. Oxford UP 1975.

[Mad90] Penelope Maddy, *Realism in Mathematics*. Oxford 1990.

[MacF05] John MacFarlane, The assessement sensitivity of knowledge attributions. In: T.S. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Oxford 2005.

[Pap78] George Pappas and Marshall Swain (eds.), *Essays on Knowledge and Justification*, Ithaca, NY: Cornell UP 1978.

[Pól54] George Pólya, *Induction and Analogy in Mathematics*, Princeton University Press 1954.

[Rav99] Yehuda Rav, Why do we prove theorems? *Philosophia Mathematica (3)* 7 (1999), pp. 5–41.

[Scha04] Jonathan Schaffer, Skepticism, contextualism, and discrimination. *Philosophy and Phenomenological Research* 69 (2004), 138–155.

[Ste75] Mark Steiner, *Mathematical Knowledge*, Cornell UP 1975.

[Troe82] Anne S. Troelstra, On the origin and development of Brouwer's concept of choice sequence. In: A. Troelstra and D. van Dalen (eds.), *The L.E.J. Brouwer Centenary Symposium*, Amsterdam: North Holland 1982, pp. 465–486.

[Thu94] William P. Thurston, On proof and progress in mathematics. *Bulletin of the American Mathematical Society* 30 (1994), pp. 161–177.

Institute for Logic, Language and Computation, Universiteit van Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands

E-mail address: `bloewe@science.uva.nl`

Philosophisches Seminar, Lehr- und Forschungsbereich III, Rheinische Friedrich-Wilhelms-Universität Bonn, Lennéstrasse 39, 53113 Bonn, Germany

E-mail address: `Thomas.Mueller@uni-bonn.de`