

ESSAY 1

Mathematical Knowledge as a Case Study in Empirical Philosophy of Mathematics

Benedikt Löwe
Thomas Müller
Eva Müller-Hill

Empirical research in philosophy of mathematics is rare. Possibly, this can be traced back to the (Kantian) conviction that mathematics does not depend on human experience, but is accessible to the *reine Vernunft*. The truth of mathematical statements and the status of mathematical knowledge according to this conviction should not depend on contingent facts about humans or human society. This view has not only left traces in philosophy of mathematics, but can also be seen in the practice of

sociology: Compared to sociological research dealing with other sciences, sociology of mathematics is severely underrepresented.¹

On the other hand, some of the central questions of philosophy of mathematics have an empirical core, and some of the statements that one finds in philosophical texts about mathematics are empirical claims.

For instance, consider the question of the connection between philosophical views and mathematical achievements. According to mainstream philosophical belief, many working mathematicians are Platonists believing in the actual existence of mathematical objects. It is part of this mainstream belief that these mathematicians interpret their own work as mental manipulation of abstract objects or of mental representations of abstract objects. This last claim is without any doubt an empirical claim about mathematicians. It is an interesting question whether there is a correlation between philosophical attitudes and mathematical success, or—to put it more bluntly—whether you have to be a Platonist in order to be a great mathematician. This statement can be true or false; the reasons for its truth or falsity need not be empirical; but the statement itself is an empirical statement, and the most natural way to find out whether it is true or false is to test it. Very few philosophers of mathematics take this last step, and, as mentioned before, it is not an easy step to take, as data on these questions is not abundant.

Philosophy of mathematics, like other areas of philosophy, relates phenomena (in this case, mathematics) to a philosophical theory. Whether the philosophical theory is correct/adequate or not is not independent of the phenomena. In analytic philosophy and in particular in philosophical logic, the analysis of phenomena is often done by a technique that one could call *conceptual modelling*, *philosophical modelling*, or *logical modelling*, in analogy to the well-known applied mathematics technique of mathematical modelling. This technique consists of a number of natural steps, one of which is to confront the philosophical model with the phenomena. We claim that in many areas of philosophy, especially in the case of philosophy of mathematics, this step is underdeveloped, and we propose to consider collecting data that allow us to identify stable

¹Cf. Heintz 2000: 9, “Die Soziologie [begegnet] der Mathematik mit einer eigentümlichen Mischung aus Devotion und Desinteresse” (sociology meets mathematics with a queer mixture of devotion and lack of interest). Her study thus reconfirms the earlier assessment of Latour (1987: 245f).

philosophical phenomena in mathematical practice.

This paper is a case study that should be seen as a part of a much larger research agenda of the first two authors: the development of an empirical study of philosophy of mathematics. Such an approach could be called “naturalistic” (as in Maddy 1997) or a “Second Philosophy of Mathematics” (as in Maddy 2007). We will use the label “*Empirical Philosophy of Mathematics*” in order to stress the fact that there is actual empirical work to be done in this field. The project *Empirical Philosophy of Mathematics* consists of a theoretical foundation and an unlimited number of questions and practical projects. Some first steps towards *Empirical Philosophy of Mathematics* have been documented in Löwe and Müller 2008, Müller-Hill 2009, Löwe 2007. The theoretical foundation should contain a sustained argument for the methodology of conceptual modelling as pointed to above, especially an argument for the necessity of empirical checks on the philosophical theories established via this method. This argument is not meant to operate on the (probably hopeless) level of generality of “giving a methodological basis for philosophy,” but rather should address the specific methodological issues of one particular argumentation scheme in philosophy, in our case, exemplified in philosophy of mathematics. Practical projects in *Empirical Philosophy of Mathematics* might, e.g., address a single philosophically interesting question about mathematics and employ empirical methods from psychology, cognitive science, linguistics, textual analysis, and quantitative or qualitative sociology, to provide an empirically well-founded answer to the given question.

In accord with what has been said above, this paper has two parts: a discussion of the method of *Empirical Philosophy of Mathematics*, and a case study addressing one specific aspect of the analysis of knowledge ascriptions in mathematics. In the first part, we shall mostly proceed descriptively, with few examples and no arguments.² Section 1 gives some background on the practice of mathematical modelling, and Section 2 draws parallels to the case of philosophical, or conceptual, modelling. However, the proper development of the main tenets of the theoretic-

²For a further discussion of the motivation for a new approach to philosophy of mathematics, we refer the reader to the introductory paper Buldt et al. 2008 in the special issue of the journal *Erkenntnis* entitled “Towards a New Epistemology of Mathematics.”

cal foundations for *Empirical Philosophy of Mathematics* has to be left for future work. In the second part (Section 3), we use data procured by the third author, described in more detail in the paper Müller-Hill 2009, to answer a specific question in the epistemology of mathematics. This part is meant to exemplify how the practical side of *Empirical Philosophy of Mathematics* would be carried out. We trust that this will provide enough detail to get an impression of the interplay between philosophical work and empirical studies.

1 *Mathematical Modelling*

The notion of a model has acquired a prominent place in contemporary philosophy of science. A great variety of uses of the term “model” has been studied (cf. Frigg and Hartmann 2006 for an overview). There is widespread agreement that models play a crucial role in scientific practice, and that a fair amount of that practice consists in modelling. In Section 2 we will suggest that the modelling paradigm also provides a good framework for certain philosophical investigations, and we will propose a modelling paradigm, which could be called *philosophical modelling* or *conceptual modelling*, as an approach to philosophy of mathematics.

In order to have some basis for drawing an analogy, in this section we will describe the practice of mathematical modelling, as exemplified in the sciences. The aim of this exposition is not to give an in-depth account of mathematical modelling, but to show its features relevant for the discussion in Section 2.

Following the Galilean conviction that the book of nature *è scritto in lingua matematica*, in science it is customary to employ mathematical methods to describe, explain, or predict phenomena wherever possible. Kepler’s work on planetary motion illustrates important aspects of such mathematical modelling. Kepler first proposed a model of the solar system in which a number of spheres containing the planets were arranged in an intricate manner determined by the five Platonic solids. This model was inspired by the conviction that God’s universe must exhibit a high degree of mathematical harmony, and it did fit the then available data reasonably well. Later on, when Tycho Brahe’s more accurate

data became available to him, Kepler proposed different mathematical models containing what we now call Kepler's three laws of planetary motion, describing elliptical instead of spherical orbits of the planets.

With hindsight, one might describe this episode in the history of science as follows:

- (1) Kepler started with a view of what a model of the solar system should look like. His view was that God must have arranged the (then known) six planets in such a way that their spatial layout is determined by some order of the five Platonic solids.
- (2) Guided by the available (Copernican) data, he found the best model from the mentioned class, i.e., a particular ordering of the solids.
- (3) A confrontation with Brahe's data afterwards showed that the model was inadequate. Kepler began to look for a new class of models, still believing that there must be mathematical harmony in the motion of the planets.
- (4) Thus Kepler went through a new cycle of modelling. He tried out a number of algebraic descriptions of planetary motion and finally arrived at his celebrated laws.

Abstracting from the historical setting, one can formulate an iterative method of mathematical modelling exhibited in this, but also in many other episodes from science:

Step 1 One starts with a class of models that appear as reasonable candidates. This class may be determined by pre-theoretical insight, or by earlier steps in the iteration.

Step 2 One collects data and tries to achieve a best fit within the available class of models.

Step 3 One determines the goodness of fit and will either be satisfied or revert to Step 1, having chosen a different class of models.

This basic scheme is at work in many areas of science, both historically and at present. Statistical tools have been developed for assessing the

“goodness of fit” of models and data. There is usually an additional layer of modelling for the data themselves (in order to handle measurement errors), and the determination of best-fit models is usually carried out on a computer (cf., e.g., Gershenfeld 1999). But the underlying structure appears to be preserved.

In mathematical modelling, the step of confronting the selected model with the observed phenomena is absolutely crucial. As every scientist will be proud to say, honesty with respect to that step is the mark of good science: A beautiful mathematical theory is worth nothing if it doesn’t fit the phenomena in question. Whether this ideal is always achieved in practice is, of course, another question, but the overall methodology is clear: Modelling is an iterative procedure, and unless one has achieved a good fit, one cannot leave the circle.

2 *Conceptual Modelling*

Viewed abstractly, the aim of establishing a “philosophy of X ” is quite similar to finding a “model for Y ” in the sciences: One wishes to gain theoretical insight into (some) aspects of a certain phenomenon by representing them in a specific way.³ In epistemology, to give just one example, one of the key questions is what knowledge consists in. Various models of knowledge are on the market; a time-honored conception, dating back to Plato’s *Meno*, takes knowledge to be justified true belief.

In an episode resembling Kepler’s confrontation of his Platonic model of the heavens with Brahe’s observational data, the Platonic conception of knowledge was also challenged by data taking the form of counterexamples: Gettier (1963) constructed plausible scenarios in which persons have justified true belief, but not knowledge. (E.g., someone might believe in a true disjunction, but her justification supports the false disjunct, while the other is, unbeknownst to her, in fact true.) The ensuing debate led to a repertoire of test cases that serves as a benchmark for theories of knowledge: One way or another, any respectable analytic model of knowledge has to answer the challenges posed by Gettier- and post-Gettier-examples.

³The question whether that is what’s behind the demand for an “explanation” of the phenomenon must be left open at this point.

As in the scientific case described in the previous section, the actual details of the story are of course more complicated. Surely everybody agrees that if Gettier cases are cases of justified true belief without knowledge, then knowledge cannot be justified true belief. But how do we know? It may be reasonably simple to determine whether a person has a justified true belief (depending on one's account of justification)—but how to judge the presence or absence of knowledge in a given scenario? People certainly have their intuitions and gut feelings, and in many cases that may be sufficient. But in the epistemological case at hand, intuition confronts intuition.

It is for such cases that we propose to mirror the method of mathematical modelling in the sciences more closely in a philosophical context.⁴ Thus, conceptual modelling of X also takes the form of an iterative process:

Step 1 Theory formation Guided by either a pretheoretic understanding of X or the earlier steps in the iteration, one develops a structural philosophical account of X , including, e.g., considerations of ontology and epistemology.

Step 2 Phenomenology With a view towards Step 3, one collects data about X and extracts stable phenomena from them that are able either to corroborate or to question the current theory.

Step 3 Reflection In a circle between the philosophical theory, the philosophical theory formation process and the phenomenology, one assesses the adequacy of the theory and potentially revises the theory by reverting to Step 1.

In many debates of contemporary epistemology, Step 2 consists of a presentation of the author's intuitions about the case at hand, possibly supported by anecdotal evidence. While this may be enough if there is widespread consensus about the analysis, a different solution needs to be found if one has to decide between competing models. The obvious solution, in view of the scientific modelling practice, is to supply more data from a more varied range of sources, including data established via accepted empirical methods. On this view, the key to successful

⁴For more details, cf. Löwe and Müller (to appear).

conceptual modelling lies in strengthening Step 2 of the above iterative scheme.⁵ In epistemology, the necessary data might, e.g., be supplied from empirical linguistics or from cognitive science. In the recent epistemological debate the need for such an empirical component is felt rather strongly, and a number of authors have discussed the philosophical relevance of “going empirical.”⁶

Similarly, philosophers of mind might in their investigations refer to results from cognitive science or neuroscience, while philosophers of language might refer to corpora from linguistics. Stich (2001) has argued that quite generally, it is time for philosophy to stop relying on intuitions, i.e., translated into our terminology: to look for an empirical basis for Step 2. In the philosophy of a science X , the relevant data might come from a description of the scientific practice, e.g., as provided by an empirically founded sociology of X . Since the 1980s, such an approach to philosophy of physics and biology has proved fruitful, even though there is no universal consensus about all of the results (Hacking 1999; Harding 1986; Knorr Cetina 1999; Latour and Woolgar 1979; and Pickering 1984).

In a similar vein, many authors emphasize the social embedding of mathematics as an important factor for philosophy of mathematics, for instance, the late Wittgenstein, Philip Kitcher (1984), Paul Ernest (1998), or David Bloor (1996 and 2004). Yet, there has hardly been done any *socio-empirical* research about actual mathematical research practice. As mentioned in the introduction, sociology of science has, with few exceptions, shunned away from taking mathematics as an object of study—mainly just because preconceived philosophical convictions made such studies appear senseless or impossible (cf. Footnote 1). The

⁵In this vein it is easier to reach a negative verdict (stop at Step 3) than to establish a positive result (start afresh at Step 1 and complete the cycle successfully). This may explain the affinity of analytic philosophy to producing more negative than positive results. Cf. also (Löwe and Müller to appear: §5.3). This affinity is connected to Gian-Carlo Rota’s (1991) sharp criticism of what he calls “mathematical philosophy.”

⁶For a still rather traditional example of employing linguistic data in epistemology, cf. Stanley 2004. More decidedly experimental views are discussed in Weinberg et al. 2001, Alexander and Weinberg 2006, Nagel 2007, Swain et al. 2008. For a detailed discussion of this debate in the framework of conceptual modelling, cf. (Löwe and Müller to appear: §4.3).

first large-scale socio-empirical study published was Bettina Heintz's (2000) work about the culture and practice of mathematics as a scientific discipline mentioned in the introduction.⁷

In philosophy of mathematics, conceptual modelling according to the above scheme is therefore hampered by the fact that there is little empirical data available; accordingly, one has to work on Step 2 and Step 3 of the modelling cycle: Acquire data and develop instruments to interpret these data. The objective of *Empirical Philosophy of Mathematics* is to integrate these two objectives, working hand in hand with sociologists, cognitive scientists, psychologists, and experts in mathematics education in order to gain access to the necessary empirical data. This is not an easy task, and it is only too tempting to oversteer: While trying to emphasize the importance of Step 2, one could forget the original philosophical agenda and transform *Empirical Philosophy of Mathematics* into pure sociology of mathematics. The philosopher has to ask herself what the philosophical content of empirical facts such as “the majority of mathematicians accepts that X knows that p ” is. How do we transcend the mere description of what mathematicians are doing and enter the study of understanding what mathematics is? In some cases, the empirical findings, or even the experience of collecting the empirical data, will be philosophically relevant for the argumentation for or against the particular methodology employed. This emphasizes the close relationship of the modelling cycle of conceptual modelling as described above and the hermeneutic circle.

Since giving an exhaustive description of and argument for the methodology of conceptual modelling is impossible, we will illustrate this approach in Section 3 via a case study that employed methods from quantitative sociology to generate the data. We think that this explains the methodology (and its difficulties) more than any attempt of a theoretical description. In our case study, the aim is to incorporate empirical data about actual mathematical practice in order to elucidate the knowledge concept used in mathematical practice. This should provide input for a philosophical theory of mathematical knowledge and the epistemological role of formalizability in mathematics. In terms of our modelling

⁷Before Heintz (2000), Markowitsch used qualitative sociological studies (interviews with mathematicians) in his 1997.

scheme, the study proceeds from Step 1 to Step 3, reaching a negative conclusion about the model chosen in Step 1.

3 *Case Study: The Role of Formal Proofs in Mathematical Epistemology*

3.1 THE PHILOSOPHICAL ISSUE AND THE EMPIRICAL TEST QUESTIONS

Mathematical knowledge is generally 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 (“*more geometrico demonstrata*”). Thus, mathematical knowledge stands out as knowledge with a uniform witness, the notion of proof, which since the end of the 19th century has been connected to the ideal of formal proof. Mathematical certainty seems to come from the fact that we believe that it is possible to formalize all informal proofs.

Even though this standard view of knowledge in mathematics is central in philosophical debates, neither formal proof nor the formalizations of informal proofs play a big role in actual mathematical practice. The philosopher of mathematics, at least if she wishes to develop a philosophy of mathematics that reflects to a certain extent mathematical practice, should try to explain this discrepancy and answer the question: What is essential about formalizability for a philosophical understanding of *mathematical knowledge*?

The traditional view could for example be presented as follows: For a mathematician X and a mathematical statement φ , “ X knows that φ ” holds if and only if X has access to an informal proof of φ that can be rendered into a formal derivation.

Note that such a characterization of mathematical knowledge is not a precise philosophical position since various notions (“has access to,” “can be rendered”) are vague and need explication.⁸ However, even from this imprecise characterization, one can deduce a number of features of mathematical knowledge (under very mild assumptions about the

⁸This is discussed in more detail in Löwe and Müller 2008: §3.

vague terms in the characterization). For instance, unless there is an inconsistency in our basic axioms of mathematics, it would be impossible for X to know φ at time t and to know $\neg\varphi$ at time t' , even if $t \neq t'$: If X had access to an informal proof at time t that can be transformed into a derivation D of φ and had access to an informal proof at time t' that can be transformed into a derivation D' of $\neg\varphi$, then the derivations D and D' witness that there must be an inconsistency in our underlying base axiom system.⁹

In the naturalistic spirit of *Empirical Philosophy of Mathematics*, we can test the above traditional view of mathematical knowledge: If the usage of mathematicians allows for conflicting knowledge ascriptions, then the traditional view must be abandoned.

It should be emphasized that it is an interesting question whether the results of our investigation are specific for mathematics at all. Our report here focuses on one particular aspect of the questionnaire study that was conducted, and could give the impression that the only thing that we are interested in is the question “Is the actual usage of mathematical knowledge attributions compatible with facticity of knowledge?” (Or, more specifically, “Can you adequately/truthfully say ‘I knew it this morning, but now I don’t know it anymore?’”). Reduced to this question, you might wonder whether our investigation would not yield exactly the same results if you told a story about biologists, physicists, or just ordinary people talking about red cars or clocks. And it actually might. We think that this is an interesting empirical research question for linguists with relevance for epistemology. If anything, such an empir-

⁹As one referee pointed out, this is assuming that there are no “different acceptable axiomatizations at different times.” This description can refer to several phenomena in mathematics, related to the subtle gradation of the use of axiomatics: On one end of the spectrum, axioms serve as part of the definition of a concept, on the other end, they are self-evident truths about mathematics. In the first case, “different acceptable axiomatizations” might refer to historical examples of changes in the meaning of basic concepts (cf. Buldt et al. 2008: §2 for a discussion of the notion of continuity); in the second case, we would be talking about something that Maddy (1997: 208) would call a failure of the maxim UNIFY. Both phenomena are interesting and worthy of study within our paradigm of *Empirical Philosophy of Mathematics*, but for the current investigation we will disregard the possibility of such a situation—in the example employed in our case study, a change in axiomatization is not at issue.

ical finding would only corroborate our belief that the idea of a special epistemic status of mathematics is highly problematic.

3.2 OUR TEST

In order to test the traditional view that mathematical knowledge is special because it is based on formalizable proof, we set up a quantitative experiment. It is difficult to get access to large enough numbers of test subjects for questions like this¹⁰; for this study, we used an online questionnaire that was advertised via scientific newsgroups. This guaranteed a relatively large number of test subjects, but limited the study to mathematicians who read newsgroups. Yet, we do not regard this as a serious limitation of the representativity of the sample, as a significant correlation between the habit of reading newsgroups and a certain attitude towards formalization is not very likely.

Our questionnaire contained mostly multiple choice questions, but also left some space for qualitative comments. It was posted on a webpage from August 2006 to October 2006 and announced via postings in three different scientific newsgroups, with about 100 valid responses out of 250 responses in total. A response was considered as valid if the personal data part and at least one question of the questionnaire was completed.¹¹

The personal data part served to decide whether a participant belonged to the target group of participants who have been involved in research or university-level teaching. We got 76 valid responses from the target group. Most of these responses came from the US, Germany, and the Netherlands. 46.1% of the target group participants hold a Ph.D. (or an equivalent degree), 19.7% an M.Sc. (or an equivalent degree), and 13.2% a B.Sc. (or an equivalent degree) in mathematics.

¹⁰Together with Francesco Lanzillotti and Henrik Nordmark, the first author has attempted to access research mathematicians via the organizers of international conferences with extremely limited success.

¹¹In the presentation of the results in Section 3.3 we will give the total count of valid answers for each reported multiple choice question.

3.3 SELECTED RESULTS

The questionnaire had 74 questions in total, but we will only present some questions and results that are significant for the interpretation regarding the test question.¹²

In one part of the questionnaire, participants were led through four scenarios. Each screen of the online questionnaire contained a piece of the scenario, and at the end of each screen the participants were asked whether they would ascribe knowledge to the protagonist of the scenario or not. In the following, we will give excerpts from one of these scenarios.

John is a graduate student, and Jane Jones, a world famous expert on holomorphic functions, is his supervisor. One evening, John is working on the Jones conjecture and seems to have made a break-through. He produces scribbled notes on yellow sheets of paper and convinces himself that these notes constitute a proof of his theorem.

Does John know that the Jones conjecture is true?

- Yes
- Almost surely yes
- Almost surely no
- No
- Can't tell

Screenshot 1: First screen of the “Jones Conjecture” scenario

On the following screens, the story continues: John presents the proof sketch to his professor, she discovers a gap and fixes it, they jointly write a paper that they submit to a mathematical journal of high reputation. After a thorough refereeing process of 18 months, the paper is accepted on the basis of a positive referee report. At the end of each screen the participants are asked the same question, “Does John know that the Jones conjecture is true?” Finally, the scenario ends with Screenshot 2.

¹²Cf. Müller-Hill 2009 for a more detailed exposition of the results of the survey.

After his Ph.D., John continues his mathematical career. Five years after the paper was published, he listens to a talk on anti-Jones functions. That evening, he discovers that based on these functions, one can construct a counterexample to the Jones conjecture. He is shocked, and so is professor Jones.

Does John know that the Jones conjecture is false?

- Yes
- Almost surely yes
- Almost surely no
- No
- Can't tell

Did John know that the Jones conjecture was true on the morning before the talk?

- Yes
- Almost surely yes
- Almost surely no
- No
- Can't tell

Screenshot 2: Final screen of the “Jones Conjecture” scenario

Table 1 contains the results for three of the questions for the Jones scenario. Q1 is the question whether John knows that the Jones conjecture is *true* after the paper was published, Q2 is the question whether John knows that the Jones conjecture is *false* after he discovers the counterexample, and Q3 is the question whether he knew that the conjecture was *true* on the morning before the talk. Note that Q2 and Q3 are on the same screen (see above), so participants are aware of their answer to Q2 when answering Q3 and vice versa. After finishing the scenario, the participants were given the opportunity to write down comments on the scenario in a free-text field. Here are some selected quotes, emphasizing different factors that influenced the answers in the Jones scenario:

- “How important is the Jones conjecture? How large is the com-

	Yes	Almost surely yes	Almost surely no	No	Can't tell
Q1	Count $\Sigma = 66$ Frequency 19 28.8%	37 56.1%	2 3.0%	3 4.5%	5 7.6%
Q2	Count $\Sigma = 62$ Frequency 9 14.5%	29 46.8%	4 6.4%	5 8.1%	15 24.2%
Q3	Count $\Sigma = 62$ Frequency 15 24.2%	29 46.8%	3 4.8%	9 14.5%	6 9.7%

Table 1: Results for the Jones scenario

- munity?”
- “My answers would have been very different with different time frames mentioned.”
 - “I don’t know John . . . , so I don’t have a good feel for how rigorously [he] . . . work[s].”

3.4 INTERPRETATION

The results from the Jones scenario are to a certain extent as one would expect. After thorough refereeing and publication of the result in a respected journal, John is taken to *know* that the Jones conjecture is true—the great majority (84.9%) of the participants gave a positive answer (“yes” or “almost surely yes”) to Q1. After discovery of the counterexample, a majority ascribes to John knowledge that the Jones conjecture is *false*—61.3% of the participants gave a positive answer to Q2, and only 14.6% deny his knowledge (“no” or “almost surely no”). The answers to Q3, however, might be surprising: In a situation in which the storyline has revealed that there is a counterexample to the conjecture, still a majority of 71% gave a positive answer to that question, i.e., would ascribe knowledge of the conjecture on the morning before he learned about the counterexample.

In order to appreciate this phenomenon, let us look at the correlation figures. Of the 38 participants who agreed to Q2, 27 also agreed to Q3 (71.1%). With this empirical result, we can now move to Step 3 of the modelling cycle: As argued in Section 3.1, a philosopher endorsing the traditional view of epistemology of mathematics is not able to accept the statement “John knows φ at time t and John knows $\neg\varphi$ at time t' .” However, a large number of our test subjects did exactly that, and so we either have to accept a notion of mathematical knowledge that ignores the usage of this large portion of the community, or give up the traditional view. Based on our methodological position, we discard the former option and choose the latter. The free text comments given in the survey further support the view that mathematical knowledge is highly context-sensitive.

Our socio-empirical analysis gives a negative result; it is now our task

as philosophers of mathematics to turn it into a philosophical insight: In the modelling cycle, we are now returning to Step 1. Based on the quantitative data, the free text comments, and further results from a qualitative interview study with mathematicians (planned by the third author), we should now develop a new understanding of knowledge that replaces the one that we called the “traditional view.”

4 References

- Joshua Alexander and Jonathan M. Weinberg. Analytic epistemology and experimental philosophy. *Philosophy Compass*, 1:56–80, 2006.
- David Bloor. *Scientific Knowledge: A Sociological Analysis*. University of Chicago Press, Chicago, IL, 1996.
- David Bloor. Sociology of scientific knowledge. In Ilkka Niiniluoto, Matti Sintonen, and Jan Woleński, editors, *Handbook of Epistemology*, pages 919–962. Kluwer, Dordrecht, 2004.
- Bernd Buldt, Benedikt Löwe, and Thomas Müller. Towards a new epistemology of mathematics. *Erkenntnis*, 68(3):309–329, 2008.
- Paul Ernest. *Social Constructivism as a Philosophy of Mathematics*. SUNY Press, Albany, NY, 1998.
- Roman Frigg and Stephan Hartmann. Models in science. In E.N. Zalta, editor, *Stanford Encyclopedia of Philosophy*. CSLI, 2006. Spring 2006 edition.
- Neil Gershenfeld. *The Nature of Mathematical Modelling*. Cambridge University Press, Cambridge, 1999.
- Edmund Gettier. Is justified true belief knowledge? *Analysis*, 23:121–123, 1963.
- Ian Hacking. *The Social Construction of What?* Harvard University Press, Cambridge, MA, 1999.
- Sandra Harding. *The Science Question in Feminism*. Cornell University Press, Ithaca, NY, 1986.

- Bettina Heintz. *Die Innenwelt der Mathematik. Zur Kultur und Praxis einer beweisenden Disziplin*. Springer, Wien, 2000.
- Philip Kitcher. *The Nature of Mathematical Knowledge*. Oxford University Press, Oxford, 1984.
- Karin Knorr Cetina. *Epistemic Cultures*. Harvard University Press, Cambridge, MA, 1999.
- Bruno Latour. *Science in Action. How to follow scientists and engineers through society*. Harvard University Press, Cambridge, MA, 1987.
- Bruno Latour and Steve Woolgar. *Laboratory Life. The Social Construction of Scientific Facts*. Sage, London, 1979.
- Benedikt Löwe. Visualization of ordinals. In Thomas Müller and Albert Newen, editors, *Logik, Begriffe, Prinzipien des Handelns*, pages 64–80. Mentis, Paderborn, 2007.
- Benedikt Löwe and Thomas Müller. Mathematical knowledge is context-dependent. *Grazer Philosophische Studien*, 76:91–107, 2008.
- Benedikt Löwe and Thomas Müller. Data and phenomena in conceptual modelling. *Synthese*, to appear.
- Benedikt Löwe, Volker Peckhaus, and Thoralf Räscher, editors. *Foundations of the Formal Sciences IV. The History of the Concept of the Formal Sciences*, volume 3 of *Studies in Logic*. College Publications, London, 2006.
- Penelope Maddy. *Naturalism in Mathematics*. Oxford University Press, Oxford, 1997.
- Penelope Maddy. *Second Philosophy: A Naturalistic Method*. Oxford University Press, Oxford, 2007.
- Jörg Markowitsch. *Metaphysik und Mathematik. Über implizites Wissen, Verstehen und die Praxis in der Mathematik*. PhD thesis, Universität Wien, 1997.

- Eva Müller-Hill. Formalizability and knowledge ascriptions in mathematical practice. *Philosophia Scientiae*, 13(2), 2009. To appear.
- Jennifer Nagel. Epistemic intuitions. *Philosophy Compass*, 2:792–819, 2007.
- Andrew Pickering. *Constructing Quarks*. University of Chicago Press, Chicago, 1984.
- Gian-Carlo Rota. The pernicious influence of mathematics upon philosophy. *Synthese*, 88:165–178, 1991.
- Dirk Schlimm. Axiomatics and progress in the light of 20th century philosophy of science and mathematics. In Benedikt Löwe, Volker Peckhaus, and Thoralf Räscher, editors, *Foundations of the Formal Sciences IV. The History of the Concept of the Formal Sciences*, volume 3 of *Studies in Logic*, pages 233–253, London, 2006. College Publications.
- Jason Stanley. On the linguistic basis for contextualism. *Philosophical Studies*, 119:119–146, 2004.
- Steven P. Stich. Plato’s method meets cognitive science. *Free Inquiry*, 21:36–38, 2001.
- Stacey Swain, Joshua Alexander, and Jonathan M. Weinberg. The instability of philosophical intuitions: Running hot and cold on truetemp. *Philosophy and Phenomenological Research*, 76(1):138–155, 2008.
- Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich. Normativity and epistemic intuitions. *Philosophical Topics*, 29(1&2):429–460, 2001.