

DOI: 10.1017/psa.2025.10134

This is a manuscript accepted for publication in *Philosophy of Science*.

This version may be subject to change during the production process.

Toward a Methodology for the Philosophy of Mathematical Practice

William D'Alessandro

Department of Philosophy, William & Mary

wbdalessandro@wm.edu


Abstract

Practice-based approaches to philosophy of mathematics have gone mainstream over the past several decades. As the paradigm has grown in popularity, however, there's been little sustained meditation—and still less any explicit consensus—on what precisely it means for philosophy to take practice seriously. The field's lack of a clear common methodology has begun to make itself felt in slowed and uncertain progress on core problems. Here I review the methodological situation and propose five canons to guide future research. I focus throughout on the study of explanation in mathematics as a guiding example.

Philosophy of mathematical practice (PMP) is an orientation or research program based on attentiveness to the ways of working mathematics. It emerged in its present form in the second half of the twentieth century—in harmony with the Kuhnian turn toward historical and sociological perspectives on science—though it had other distinctive causes, including Quine’s naturalistic doctrines (the epistemic primacy of scientific method, the denial of analyticity, the continuity of empirical science and mathematics), on display especially in Maddy’s pioneering work in philosophy of set theory (Maddy 1988a, 1988b, 1993). PMP entered a growth period in the first decade of the 2000s, seeing major new research programs on explanation in mathematics, informal proofs and the criteria governing their acceptability, methodological desiderata such as purity, the demonstrative and cognitive roles of diagrams, and so on.

It’s now 2025, and PMP’s mission to reinvent philosophy of mathematics seems largely complete. Works on PMP topics now appear often in philosophy’s most highly regarded journals. The Association for the Philosophy of Mathematical Practice boasts almost 300 members and runs a large biennial conference entering its eighth iteration. And justificatory appeals to working mathematics have become commonplace in the literature. Fig. 1 below shows occurrences of ‘mathematical practice’ in Google’s book corpus over time, with an exponential growth trend appearing to start around the middle of the 20th century:

Fig. 1: Trend in appearances of ‘mathematical practice’ in Google’s book corpus, from <https://books.google.com/ngrams/>.



As a card-carrying practice-oriented philosopher myself, I'm largely glad of these developments. As PMP's popularity grows, though, a couple of things appear increasingly clear. First, it's not always obvious what it means for philosophy to be faithful to mathematical practice; and insofar as there's agreement about the nature of the goal, controversy remains about the proper methods for pursuing it. (Rittberg 2019 makes this case with clarity and insight, pleading for more open methodological debate in PMP. Though he wouldn't endorse many specifics of my methodological vision, I believe the project is in line with his call.) Second, until some consensus is reached on these issues, it will grow harder to make progress on PMP topics.

As an illustration of this second point, consider the explanation-within-mathematics program, one of the best-developed and most influential lines of PMP research. The major initial job of philosophers in this area was to justify the program's existence, by documenting that mathematicians talk often about explanation and take explanatory concerns seriously. This was a nontrivial task—the subject met with skepticism in its early days (Resnik & Kushner 1987; Zelcer 2013), and amassing enough evidence to mollify critics took time—but it mostly required diligence rather than sophisticated methodological doctrine.¹

Having finished this preliminary job, we find ourselves in the position of determining just *what* lessons about explanation we should learn from mathematical practice. The new situation is less straightforward than the old. There are many prospective sources of evidence to consult, each offering a limited window into a particular practice in a particular context. Some of these sources disagree (or seem to). Others make apparently similar claims in different language. It isn't always clear which sources of information about a given practice are most informative or reliable, or even which practices we should aim to gather evidence about in the first place. Nor is it obvious how our prior philosophical commitments and general acquaintance with scientific practice should shape our interpretation of the mathematical data.

Different PMP works on explanation seem inclined to give different answers to such questions. A look at the recent literature reveals conflicts over the role of data-driven empirical methods as opposed to traditional case studies, the evidential value of one-off quotations from mathematical texts, the role of practice-based evidence as primary and indispensable as opposed to merely confirmatory, and so on.

¹ Much of this valuable work was done by Mancosu and collaborators in the 2000s: Mancosu (1999, 2000, 2001, 2008a) and Mancosu & Hafner (2008) are a few examples.

Such troubles are on particularly clear display in the explanation debate, but their nature is inherently general. As a result of these methodological misalignments, some cracks in the PMP edifice have begun to show. Progress beyond the earliest stages on explanation and other PMP topics is likely to be slow until the rifts are mended.

My goals in this paper are as follows. §2 describes what I take to be the epistemic motivation for practice-oriented philosophy and extracts a general methodological outlook. On the basis of these considerations, §3 offers five guiding canons for PMP research. I illustrate what's at stake with particular reference to the explanation literature. §4 briefly concludes.

2 Why Philosophy of Mathematical Practice?

In order to make progress on PMP issues, researchers must decide how to identify and weight the relevant mathematical practices. I suggest we start by asking why we aspire to do this at all. What does PMP hope to achieve by taking working mathematics as its guiding star?

A standard (and I think basically correct) answer runs along the following naturalistic lines. Philosophy should respect mathematical practice because, and to the extent that, mathematics is a highly successful epistemic enterprise. That is, we should accept experts' methods as *prima facie* appropriate and their judgments as *prima facie* correct because, and to the extent that, those methods and judgments reliably produce valuable epistemic goods. These goods include, for instance, knowledge, true belief, understanding, evidence, well-calibrated credence, accurate heuristics, fruitful concepts and correct explanations. It's common for naturalists to insist on the importance of both the internal viewpoint of pure mathematics and the external viewpoint of scientific applications in assessing mathematics' epistemic success.

Mathematics consists of a multitude of practices, many of which may be worth studying to gain insight into interesting patterns of human intellectual, professional and pedagogical behavior. But the naturalistic considerations above suggest a narrower focus for PMP. Practice-oriented philosophers of explanation and the like should principally aim to understand and respect those practices which are responsible for math's epistemic success. Accordingly, we should most strongly weight that evidence which provides the most informative signal about such practices.

I adopt a broad notion of "practice" here. This notion is inclusive of both enduring general methods (e.g., community-wide standards of proof) and smaller-scale phenomena (e.g., a

certain school of topologists' judgment that a given proof is explanatory). Both types of activity can constitute parts of successful mathematics, so both are potentially relevant to PMP.

From this viewpoint, the problem of methodology largely reduces to a set of questions about the quality of various kinds of evidence. PMP's ultimate goal is to fathom the norms, habits, methods and patterns of judgment with which well-functioning mathematics gets made. Since we can't peer directly into mathematicians' minds, we're obliged to rely on indirect proxies for these phenomena. A given such proxy is useful to PMP inasmuch as it tells us about the latter, the more reliably and informatively the better.

On this picture, the lack of, and need for, good PMP methods is of a piece with methodological problems elsewhere in science. Consider the ongoing replication crisis in psychology, medicine and social science, which has seen classic results repeatedly cast into doubt by subsequent reproduction failures (Shrout & Rodgers 2018). Broadly speaking, the replication crisis occurred because the prevailing standards for conducting and evaluating research were only weakly related to objective evidential strength. For instance, many researchers engaged in *p*-hacking and other manipulations to create a false appearance of strong evidence. Scientific journals' publishing decisions were systematically influenced by non-epistemic factors, such as the headline value of a paper's claimed results. And weak norms around data sharing and study design transparency made it difficult for third parties to spot evidential inadequacies.

Similarly, it's incumbent on PMP to ask which methods are strongly conducive to learning truths about the world—in this case, truths about relevant mathematical practices—as opposed to merely pleasing referees, aligning with a trend du jour or providing superficial cover against charges of armchair philosophizing. If (as I think) the field currently lacks consensus on such methods, we should worry, and take action. As the sciences have shown, good intentions and the trappings of academic rigor aren't enough to prevent crises. (Or, rather, to prevent philosophy's more typical failure mode, stagnation followed by neglect and irrelevance. Crises are a luxury for fields with dominant paradigms.)

In the framework of Carter (2019), the outlook described in this section corresponds with the “epistemological” strand of PMP—as distinct from the “agent-based” strand (which focuses on human and social dimensions of practice) and the “historical” strand (which examines the development of practices as a contingent evolutionary process). These latter traditions are, of course, worthwhile and necessary (if perhaps best understood as belonging to sociology, anthropology or history in some cases). I focus on the epistemological project in

part because it's the strand with which I'm most familiar, and in part because the methodological situation is somewhat different where naturalistic success criteria play a less central role. I believe many of the questions raised in this paper arise in some form for any practice-oriented project, however. And I hope the answers I suggest will prove useful for a broad range of PMP studies.

3 Five Methodological Canons for PMP

3.1 Evidence over Intuition

The first canon I'll propose may seem too obvious to bear mentioning. It enjoins PMP practitioners to *base philosophical claims on explicit practice-based evidence*, rather than (say) their own intuitions, impressions or unwarranted overinterpretations of mathematical texts. While this injunction is in some sense the constitutive tenet of PMP, it's heeded with surprising irregularity.

The justification for this canon needs little explaining. The goal of (the relevant strand of) PMP is to learn about successful mathematical practices. We look outside our own heads to inform ourselves about these practices because we mostly aren't working mathematicians, and thus we're in no position to opine authoritatively about how good math is done. To fall back on our own judgments as primary evidence, then, is to build our houses on sand.

This is a longstanding problem in the literature, even in work that professes to be doing PMP. Steiner (1978), the founding work on explanation in mathematics, begins by declaring that "Mathematicians routinely distinguish proofs that merely demonstrate from proofs which explain" (135). One might expect to see many examples of this alleged activity in support of Steiner's (then-controversial) thesis. The opening page does indeed offer two such, citing a Feferman chapter on predicative analysis and a Chang and Keisler model theory textbook. But the rest of the many explanation claims in the paper, including all the key theory-building examples and counterexamples, are represented as Steiner's own judgments, unsupported by references to mathematical practice.

It's perhaps unfair to judge Steiner too harshly; he wrote in the 1970s when PMP hadn't yet come into its own as a self-conscious program. But similar moves can be found even in the newest work. For instance, Antos & Colyvan (2024) explicitly aligns itself with PMP, declaring that "we need to draw on the judgments of mathematicians about which proofs are explanatory, if we are to respect mathematical practice... we see our task to be that of taking the judgments of mathematicians and trying to make philosophical sense of these" (290). In a

work full of suggestions about the explanatory value of proofs in descriptive set theory, though, I count just one citation of an explanation claim from a mathematical source (300). Elsewhere in the paper, a pair of mathematicians describe one proof of the Cantor-Bendixson theorem as “more informative” than another. Antos and Colyvan write that they “understand this use of ‘more informative’ to at least partly to mean ‘more explanatory’” (298). But the justification for this interpretative move isn’t obvious. A priori, informativeness has little to do with explanatory power, and there’s nothing in the relevant sources that strongly suggests such a reading.

To be clear, my claim here isn’t that non-experts’ intuitions about mathematical matters are always completely unreliable. Most philosophers understand at least some mathematics, and share at least some of the cognitive machinery responsible for mathematicians’ judgments, after all; it would be out of step with PMP’s naturalistic motivations to posit a radical discontinuity between infallible experts and incompetent outsiders. In particular, appeals to one’s own judgment may be fine in simple cases which don’t demand specialist knowledge or highly trained sensibilities. Since these sorts of cases aren’t of primary interest for most PMP projects, though, clear practice-based evidence ought to remain the default.

3.2 Practices over Particularities

My second canon enjoins PMP practitioners to *aim at evidence that indicates a widespread trend in a successful practice*. Individual judgments about particular cases will often be relatively noisy proxies for patterns, norms and methods of interest. This is true even when such judgments come from accomplished experts. Mathematicians, like most of us, possess idiosyncrasies which are imperfectly aligned with the main stream of successful practice. They sometimes speak imprecisely or unthoughtfully. They sometimes get things wrong. On the other hand, third parties may lack the context to interpret a remark in the intended way. While some practice-based evidence is certainly better than none, then, we should resist the temptation to consider the PMP box duly checked once we’ve managed to adduce a relevant expert quotation.

Again, an analogy with scientific epistemology is appropriate here. The results from a single study should almost never be taken as conclusive, or even as particularly strong evidence. This is true even when the study in question was well designed and carefully performed, because empirical research is an inherently noisy undertaking. A properly conducted meta-analysis can be much more informative, since the invisible biases, mistakes and randomness in individual studies will tend to cluster around the ground-level truth.

Of course, formal meta-analysis is rarely an option in the PMP context. But practice-oriented philosophers can adopt methods which are similar in spirit. Most importantly, we should strive to collect many relevant data points, present these in full, and try in good faith to discern the shape of the resulting pattern. In doing so, we shouldn't be put off by apparent disagreements or disparities; these are to be expected in any noisy data set. (But we can try to understand their causes and consider whether they can be reconciled, or alternatively disregarded as problematic outliers.) All else being equal, a larger number of convergent judgments from independent sources makes for a stronger evidential signal. So it's prudent to examine a variety of materials (e.g. textbooks, blog posts and explainer-style pieces, not just journal articles), and to seek out heterogeneous voices (e.g. temporally and ideologically diverse experts, not just a handful of name-brand famous mathematicians).

Hanna & Larvor (2020) touches on this theme in criticizing the practice of using one-off quotations to provide a veneer of practice-based justification for philosophical claims. As Hanna and Larvor point out, this sort of evidence is rarely probative in the absence of information about general trends and the existence of disagreement:

[W]e need the reflective testimony of mathematicians to help us to understand what they are doing. On the other hand, the usual reservations about practitioner-testimony apply to mathematics. Adepts in any practice can fail to understand what they are doing, how they are doing it and what conditions make it possible. ...[Given the inevitability of variation and disagreement], it makes little sense to use isolated quotations as indicators of how mathematics is uniformly or usually done. (1137)

Attempting to weigh and synthesize multiple sources of information also requires sensitivity to the origin and quality of each item. It's often possible to tell a throwaway remark from a deeply considered judgment, a simplification intended for students from a meticulous appraisal, an idiosyncratic opinion from a reflection of general sentiment, and an elder statesman's remembrances from the testimony of one immersed in the work. "Rigour, in the humanities disciplines that rely on textual quotation, means taking proper account of the context from which the quotations is taken" (Hanna & Larvor 2020, 1145).

It's sometimes mistakenly thought that naturalism requires philosophers to mindlessly rubber-stamp every line from a scientist's pen. On the contrary, making appropriate distinctions between expert sources is a crucial part of taking practice seriously.

3.3 Embracing Informality

My third canon exhorts PMP practitioners to *make good use of relevant informal sources*. It would be a mistake to assume that interesting mathematical practices are only on display in traditionally published scholarly literature. On the contrary, there are pressures in such media to limit discussions to minimal statements of theorems and proofs.

Mathematicians often discuss methods and motivation more freely in informal online venues: blogs, research-oriented question-and-answer fora like MathOverflow, unpublished notes on personal sites, substantive interviews with science magazines like *Quanta*, and so on. Philosophers should familiarize themselves with these sources and not shy away from citing them where appropriate.

Another benefit of engaging with informal sources is tied to the previous canon. Many online venues are participatory, often attracting mathematicians of varying backgrounds from around the world. They're therefore useful places to look for agreement and disagreement, and in general to form a more nuanced picture of experts' thinking.

Here, for instance, is a random example, discovered in a few minutes' search. The MathOverflow thread "Why are the sporadic simple groups HUGE?"², started by a mathematics graduate student, seeks explanatory understanding on several related questions:

[T]he order of the monster group is over 8×10^{53} , yet it is simple, so it has no normal subgroups... how? What is so special about the prime factorization of its order? Why is it 2^{46} and not 2^{47} ? Why is it not possible to extend it to obtain that additional power of 2 without creating a normal subgroup? Some of these properties seem really arbitrary, and yet must be very fundamental to the algebra of groups.

This question currently has five answers, four from professional mathematicians, presenting a variety of explanatory viewpoints. Most of the answers have further comments in a refining or contesting spirit. For instance, the representation theorist Scott Carnahan notes that the primes appearing in the prime factorization of 10^{53} are exactly the supersingular primes, and observes that this fact belongs to the more general "monstrous moonshine" phenomenon, for which "the question of a general conceptual explanation is still open". Noah Snyder, a topologist, suggests that in fact the sporadic finite simple groups aren't particularly large;

² <https://mathoverflow.net/questions/17617/why-are-the-sporadic-simple-groups-huge>.

several commentators debate this point, generating a small discussion about notions of structural complexity and their relationship to size. The algebraist Simon Lentner details the origins of the sporadic groups in a certain type of inductive construction process, which can be carried unusually far in the case of the Monster. And so on.

It's very difficult to get this sort of window into practice—especially the dynamic social practice that drives much progress in contemporary mathematics—by consulting only textbooks and journal articles. Nor is there anything special about the case just mentioned. MathOverflow is rife with discussions on explanation and other philosophically interesting topics, many of which feature more numerous, sophisticated and revealing answers.

My impression is that practice-oriented philosophers currently make little use of informal sources in their research. I believe that doing so would help with the quantity-of-evidence issues discussed above. It can after all be burdensome to locate relevant references in traditional scholarly literature, much of which is gated behind journal subscriptions, segmented across numerous repositories, undigitized or otherwise not readily searchable, and so on. If scholars limit themselves to these relatively research-unfriendly sources, it's no wonder they're tempted to lower their evidential standards: few projects would otherwise ever get finished. By contrast, blogs and other online fora are readily searchable by nature. Availing ourselves of these sources can therefore kill two methodological birds with one stone, improving both the amount of evidence in circulation and the richness of our picture of successful practice.

3.4 Care with Corpora

My fourth proposed canon urges PMP practitioners to *engage carefully with current empirical methods*, especially those relying on coarse-grained text analysis and similarly blunt tools.

In light of the quantity-of-evidence problems addressed by canons 1 and 2, the appeal of data-driven empirical methods is easy to appreciate: they hold out hope for an objective, systematic view of mathematical practice, immune to nagging worries about cherry-picked examples and file-drawer problems. It's plausible that more sophisticated versions of these methods will eventually come to our rescue. But the forms of corpus analysis available to most researchers today have little to offer, even when applied to large data sets.

One problem is that “indicator phrases” related to topics of PMP interest are typically highly ambiguous and context-sensitive; this is the “word-concept problem” for corpus methods discussed by (Chartrand 2022). Consider explanation, for instance. Terms like

‘explains’ and its cognates sometimes indicate genuine explanations (in the sense of reasons why a mathematical fact is true), but they more often indicate other notions: explaining what a term means, how to perform a calculation, why a notation was chosen, how to interpret a theorem, and so on. (I suspect that anyone with experience searching mathematical texts for PMP-relevant passages will be deeply familiar with this issue.) Conversely, there are many ways to raise explanatory concerns in alternative language. The same can be said for other philosophically interesting notions: beauty, simplicity, depth, understanding, naturalness. Some work based on such methods frames itself adversarially as PMP done right:

[T]he methodology employed here provides an empirical way to study mathematical practices on a broader scale than the traditional methods of philosophy of mathematics, such as the method of case studies. Relying on a few case studies might provide an inaccurate picture of what mathematical practices are really like, for the selected case studies may simply be outliers. ...The methods of text mining and corpus analysis used in this paper can provide a more accurate picture of what mathematical practices are like than the case study method. (Mizrahi 2020, 565)

In Mizrahi’s case, the methods in question are essentially a count of variants of *account*, *explain*, *explicate* and *elucidate* in five mathematics journals (compared against a count of perhaps even more problematic justification-indicator phrases, such as *demonstrate*, *justify*, *prove* and *show*). Contra Mizrahi, it’s difficult to see how such a blurry lens could project an accurate image of mathematical practice. Indeed, none of Mizrahi’s three randomly selected corpus examples clearly involve explanation in the relevant sense, though the author doesn’t comment on this (560).

Such shortcomings are to be expected given the study’s setup, as is recognized by those who have thought carefully about empirical methods in philosophy. “[T]he approach adopted by Mizrahi.... raises the most concerns [of several corpus-based studies of explanation], as it interprets the use of words like ‘explain’... as indicating that the text excerpt that contains this segment is an explanation. But obviously people do not label their discourse as they are engaging in it” (Chartrand 2022, 8). There are of course more sophisticated ways to attempt this sort of project. One can choose focal terms more carefully, tailor hypotheses more narrowly, or (given a small enough corpus and large enough workforce) individually check whether each data-mined item falls under the concept of interest. But the chances of

practicably obtaining strong evidence for an interesting philosophical conclusion seem small in general.

What's needed to propel such methods to greater success is a language-processing system capable of taking in vast amounts of text and ascertaining the meanings it contains to a high degree of accuracy. The current generation of frontier language models is already quite good at this; the main bottleneck, perhaps soon to be overcome, is the capability for typical academic researchers to feed such systems big-data-sized corpora and receive detailed analyses in a reasonable time. With such a setup in place, a researcher could query a model with an open-ended prompt (e.g. "Identify all claims about explanation in the primary philosophical sense, whether expressed in that language or not, and describe some noteworthy trends in the results") and expect reasonable output. That will be an interesting day for PMP when it comes. Given the current state of deep learning systems and the vast resources being directed toward their further development, I expect this research paradigm to look quite different in five years.

3.5 Experiment and Ideology

In light of the issues discussed above, the holy grail of PMP research would seem to be a well-designed empirical study in which a large sample of mathematicians are directly asked relevant questions and statistics can be done on the results. Seeking to improve on earlier such efforts (Inglis & Aberdein 2015, 2016), Mejía Ramos et al. (2021) recently attempted this.

Mejía Ramos and colleagues obtained 760 judgments from 38 mathematicians on the relative explanatory values of nine proofs (of a simple statement about the roots of a cubic polynomial), using a pairwise comparison setup. Happily for PMPers, the mathematicians studied strongly agreed about which proofs were more and less explanatory, with an inter-rater reliability coefficient of .947. This result should go some way toward allaying any lingering worries that mathematicians' judgments about explanation are too variable and contradictory to meaningfully study.

Curiously, however, Mejía Ramos et al. instructed their subjects to apply an *ontic* notion of explanation in their judgments, and "not to focus on how [the proofs] might be received by a particular audience" (585). Since there's live debate about whether the ontic conception of explanation applies widely in mathematics, this was a curious methodological choice. In practice, the subjects seem to have ignored the prompt, assigning very dissimilar ratings to

two proofs which gave variants of the same argument in different language. This arguably supports a non-ontic picture, but only by accident—better not to have courted confusion in the first place. (D’Alessandro & Lehet 2024 defends a non-ontic theory which, it’s argued, makes good sense of Mejía Ramos et al.’s findings.)

In general, experimentalists ought to avoid designing studies around controversial ideology in ways likely to confound results or muddle interpretations. Doing so perfectly may be challenging, given that unproblematic common ground is sometimes difficult to find in philosophy. But empirical work can at least aim to minimize the excess theoretical baggage it carries. PMPers should in turn be alert to these issues.

4 Conclusion

There’s much more to say about these and other methodological issues: enough, indeed, to fill volumes. Let me close with a plea to PMP practitioners. As fellow travelers on the road to understanding, let us think harder, and talk more openly, about what it is that we do (or aspire to do) when we engage in PMP work. In particular, let’s not lose sight of the goal of assembling compelling and systematic evidence about the practices we aim to capture—the kind that stands a chance of telling us how mathematics really works, and not just helping us publish papers in a faddish field.

The other sciences have been contending with mortal questions about evidence and methodology for some time now, and they’re the better for having faced the music and made some painful changes. It’s time for us to learn from their example. Fortunately, that’s the sort of thing we do best.

References

- Antos, Carolin and Mark Colyvan. 2024. "Explanation in descriptive set theory." In K. Robertson and A. Wilson (eds.), *Levels of Explanation*, Oxford: Oxford University Press, 289-309. <https://doi.org/10.1093/oso/9780192862945.003.0014>.
- Carter, Jessica. 2019. "Philosophy of mathematical practice—motivations, themes and prospects." *Philosophia Mathematica* 27, 1-32. <https://doi.org/10.1093/phimat/nkz002>.
- Chartrand, Louis. 2022. "Modeling and corpus methods in experimental philosophy." *Philosophy Compass* 17(6), e12837. <https://doi.org/10.1111/phc3.12837>.
- D'Alessandro, William and Ellen Lehet. 2024. "A noetic account of explanation in mathematics." *Philosophical Quarterly*. <https://doi.org/10.1093/pq/pqae137>.
- Hafner, Johannes and Paolo Mancosu. 2005. "The varieties of mathematical explanation." In Paolo Mancosu, Klaus Froyen Jørgensen and Stig Andur Pedersen (eds.), *Visualization, Explanation and Reasoning Styles in Mathematics*, Berlin: Springer, 215-250.
- Hanna, Gila and Brendan Larvor. 2020. "As Thurston says? On using quotes from famous mathematicians to make points about philosophy and education." *ZDM* 52, 1137-1147. <https://doi.org/10.1007/s11858-020-01154-w>.
- Inglis, Matthew and Andrew Aberdein. 2015. "Beauty is not simplicity: An analysis of mathematicians' proof appraisals." *Philosophia Mathematica* 23, 87-109. <https://doi.org/10.1093/phimat/nku014>.
- Inglis, Matthew and Andrew Aberdein. 2016. "Diversity in proof appraisal." In Brendan Larvor (ed.), *Mathematical Cultures*, Birkhäuser: Switzerland, 163–179. https://doi.org/10.1007/978-3-319-28582-5_10.
- Maddy, Penelope. 1988. "Believing the axioms. I." *Journal of Symbolic Logic* 53, 481-511. <https://doi.org/10.2307/2274520>.
- Maddy, Penelope. 1988. "Believing the axioms. II." *Journal of Symbolic Logic* 53, 736-764. <https://doi.org/10.2307/2274569>.

Maddy, Penelope. 1993. "Does V equal L ?" *Journal of Symbolic Logic* 58, 15-41.
<https://doi.org/10.2307/2275321>.

Mancosu, Paolo. 1999. "Bolzano and Cournot on mathematical explanation." *Revue d'Histoire des Sciences* 52, 429-455. <https://doi.org/10.3406/rhs.1999.1364>.

Mancosu, Paolo. 2000. "On mathematical explanation." In Emily Grosholz and Herbert Breger (eds.), *The Growth of Mathematical Knowledge*, Dordrecht: Kluwer, 103-119.
https://doi.org/10.1007/978-94-015-9558-2_8.

Mancosu, Paolo. 2001. "Mathematical explanation: Problems and prospects." *Topoi* 20, 97-117. <https://doi.org/10.1023/A:1010621314372>.

Mancosu, Paolo. 2008. "Mathematical explanation: Why it matters." In Paolo Mancosu (ed.), *The Philosophy of Mathematical Practice*, New York: Oxford University Press, 134-150. <https://doi.org/10.1093/acprof:oso/9780199296453.003.0006>.

Mancosu, Paolo (ed.). 2008. *The Philosophy of Mathematical Practice*. New York: Oxford University Press.

Mancosu, Paolo and Johannes Hafner. 2008. "Unification and explanation: A case study from real algebraic geometry." In Paolo Mancosu (ed.), *The Philosophy of Mathematical Practice*, New York: Oxford University Press, 151-178.

Mejía Ramos, Juan Pablo, Tanya Evans, Colin Rittberg and Matthew Inglis. 2021. "Mathematicians' assessments of the explanatory value of proofs." *Axiomathes* 31, 575-599.
<https://doi.org/10.1007/s10516-021-09545-8>.

Mizrahi, Moti. 2020. "Proof, explanation, and justification in mathematical practice." *Journal for General Philosophy of Science* 51, 551-568. <https://doi.org/10.1007/s10838-020-09521-7>.

Resnik, Michael D. and David Kushner. 1987. "Explanation, independence and realism in mathematics." *British Journal for the Philosophy of Science* 38, 141-158.
<https://doi.org/10.1093/bjps/38.2.141>.

Rittberg, Colin Jakob. 2019. "On the contemporary practice of philosophy of mathematics." *Acta Baltica Historiae et Philosophiae Scientiarum* 7, 5-26.
<https://doi.org/10.11590/abhps.2019.1.01>.

Shrout, Patrick E. and Joseph L. Rodgers. 2018. "Psychology, science, and knowledge construction: Broadening perspectives on the replication crisis." *Annual Review of Psychology* 69, 487-510. <https://doi.org/10.1146/annurev-psych-122216-011845>.

Steiner, Mark. 1978. "Mathematical explanation." *Philosophical Studies* 34, 135-151.
<https://doi.org/10.1007/bf00354494>.

Zelcer, Mark. 2013. "Against mathematical explanation." *Journal for General Philosophy of Science* 44, 173-192. <https://doi.org/10.1007/s10838-013-9216-6>.

Acknowledgments

Thanks to audiences at the 2024 Philosophy of Science Association meeting in New Orleans, the 2024 Association for the Philosophy of Mathematical Practice meeting in Pavia, and the 2024 ExUnMa talk series for lively and interesting discussions. I'm especially grateful to Silvia De Toffoli, Chris Pincock and Colin Rittberg for their feedback and inspiration. Finally, cheers to the APMP and the larger philosophy of mathematical practice community for providing a welcoming intellectual home—among our discipline's best kept secrets, though I hope it won't stay that way forever.

Funding

None to declare.

Competing interests

None to declare.