



UNIVERSITY OF LEEDS

This is a repository copy of *Platonic Relations and Mathematical Explanations*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/163904/>

Version: Accepted Version

Article:

Knowles, R (2021) Platonic Relations and Mathematical Explanations. *The Philosophical Quarterly*, 71 (3). pp. 623-644. ISSN 0031-8094

<https://doi.org/10.1093/pq/pqaa053>

© The Author(s) 2020. This is an author produced version of an article published in *Philosophical Quarterly*. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

PLATONIC RELATIONS AND MATHEMATICAL EXPLANATIONS

BY ROBERT KNOWLES

Some scientific explanations appear to turn on pure mathematical claims. The enhanced indispensability argument appeals to these ‘mathematical explanations’ in support of mathematical platonism. I argue that the success of this argument rests on the claim that mathematical explanations locate pure mathematical facts on which their physical explananda depend, and that any account of mathematical explanation that supports this claim fails to provide an adequate understanding of mathematical explanation.

Keywords: mathematical explanation, enhanced indispensability argument, mathematical platonism, ontic dependence, mathematical structuralism, underdetermination.

I. INTRODUCTION

Some scientific explanations appear to turn on pure mathematical claims. According to an influential line of argument, these *mathematical explanations* evidence *mathematical platonism*—the view that abstract mathematical objects exist. The argument can be simply stated as follows. We should infer to our best scientific explanations. Some of these turn on pure mathematical claims, ascribing properties to abstract mathematical objects. So, we should believe in abstract mathematical objects. This is the *enhanced indispensability argument* (EIA).¹

¹ See Baker 2005. EIA ‘enhances’ the traditional indispensability argument, attributed to Quine 1948 and Putnam 1971, by doing without the controversial principle that confirmation is radically holistic. See Sober 1993, Maddy 2005, and Morrison 2012 for relevant criticism of the traditional argument.

EIA supporters take mathematical explanations to work by locating explanatory mathematical facts.² Critics take them to work by exploiting mathematics as a representational aid to pick out otherwise elusive explanatory physical facts.³ Progress in this debate requires a mature enough understanding of mathematical explanation to adjudicate between these views.

In recent years, a number of philosophers have defended accounts of mathematical explanation that appear promising for bolstering EIA: Christopher Pincock argues that certain explanations locate mathematical facts on which their explananda depend, via a *sui generis*, non-causal dependence relation; Sam Baron, Mark Colyvan, David Ripley, and Mark Povich suggest that mathematical explanations locate the mathematical facts on which their explananda counterfactually depend; and Aidan Lyon argues that mathematical explanations identify causally relevant mathematical properties.⁴ These accounts appear promising for bolstering EIA because they are ontic: according to them, mathematical explanations limit the network of objective dependencies in which their explananda are embedded, and locate mathematical facts within that network.⁵ In contrast, relatively few have defended accounts that seem promising for undermining EIA.⁶ It may therefore appear that the balance of evidence currently favours EIA. I aim to dispel this appearance by providing support for the following argument.

PREMISE 1

Any account on which mathematical explanations locate pure mathematical facts on

² E.g. Colyvan 2002, 2013; Baker 2005, 2009, 2017; Lyon 2012.

³ E.g. Daly & Langford 2009; Leng 2010, 2012; Rizza 2011; Saatsi 2011; Yablo 2012.

⁴ Lyon 2012; Pincock 2015b; Baron *et al.* 2017; Baron 2019; Povich 2019.

⁵ I assume that there are non-causal relations of explanatory dependence, in accordance with some dominant thinking about grounding (see Schaffer 2009; Rosen 2010; Audi 2012; Clark & Liggins 2012; Raven 2015). I help myself to, and use interchangeably, the associated locutions ‘in virtue of’, ‘depends on’, and ‘is grounded by’. Two further clarifications: First, I talk as though grounding can relate many different kinds of things; if it turns out to be more selective, my arguments can be regimented accordingly. Second, I do not presume grounding is unitary; for all I say, it may be a non-committal means of discussing a range of more specific relations (see Wilson 2014).

⁶ See Leng 2012 and Knowles & Saatsi 2019.

which their physical explananda depend fails to provide an adequate understanding of mathematical explanation.

PREMISE 2

Any account of mathematical explanation apt to bolster EIA implies that mathematical explanations locate pure mathematical facts on which their physical explananda depend.

CONCLUSION

Any account of mathematical explanation apt to bolster EIA fails to provide an adequate understanding of mathematical explanation.

Part of the support I provide for this argument is piecemeal: I identify the three above-mentioned accounts of mathematical explanation as confirming instances of the generalizations PREMISE 1 and PREMISE 2 (§§II–IV).

In §II, I show that Pincock's (2015b) account faces two serious problems. First, it posits striking yet inexplicable regularities. This is the *brute regularities problem*. Second, it renders our supposed success in identifying the mathematical facts on which physical phenomena depend a cosmic coincidence. This is the *empirical access problem*. These problems reveal that Pincock's account does not furnish an adequate understanding of mathematical explanation. It is thus a confirming instance of PREMISE 1.

In §III, I show that, for the increasingly popular *countermathematical* account to provide any understanding of mathematical explanation, it must posit a *sui generis* physical-on-mathematical dependence relation. The result is that it also faces the brute regularities and empirical access problems and is a further confirming instance of PREMISE 1.

In §IV, I show that Lyon's (2012) account is best understood in terms of the dependence of mixed mathematical/physical facts on physical facts. In contrast to the other accounts, the direction of dependence is reversed here, which means the brute regularities and

empirical access problems are avoided. However, the cost is the ability of Lyon's account to bolster EIA. Lyon's account is a confirming instance of PREMISE 2.

The argumentative moves in each case are responsive to the same general reasons. Capitalising on this, I will formulate a general argument for CONCLUSION, in the form of a dilemma (§V). Finally, I will show that, despite initial appearances, adopting structuralism about mathematical objects will not help to avoid this dilemma (§VI). All of this provides substantial (albeit defeasible) support for CONCLUSION.

II. ABSTRACT DEPENDENCE

Pincock (2015b) claims that certain scientific explanations account for the properties of collections of physical systems by appealing to the properties of objects more abstract than those physical systems. He calls these *abstract explanations*, and posits the *sui generis*, non-causal relation of *abstract dependence* to make sense of them. His chosen case study is an explanation in which the more abstract objects are mathematical.

If Pincock is right, then his account bolsters EIA. In this section, I show that Pincock's account faces the brute regularities and empirical access problems, and so fails to provide an adequate understanding of mathematical explanation. First, I show that Pincock's account leaves it underdetermined which mathematical fact is selected for by the abstract dependence relation. Then I argue that we must nevertheless take instances of abstract dependence to select for particular mathematical facts, as Pincock does. Finally, I show that this leads to the brute regularities and empirical access problems.

Pincock's case study is the mathematical proof that soap formations must satisfy Plateau's laws. Plateau's laws capture three striking regularities in soap film and bubble formations:

- (1) Soap formations consist of finite flat or smoothly curved surfaces smoothly joined together.
- (2) Within a soap formation, there are three possible meetings of surfaces: (i) no surfaces meet; (ii) exactly three surfaces meet along a smooth curve; (iii) exactly six surfaces (together with four curves) meet at a vertex.
- (3) When three surfaces meet along a curve, they do so at angles of 120° ; when four curves meet at a point, they do so at angles of $\approx 109^\circ$.

The proof, from Jean E. Taylor, can be divided into three parts (following Almgren & Taylor 1976). The first is the initial modelling phase, where a mathematical analogue of soap formations is defined, capturing the basic physical principle that soap formations minimize their total surface area. For a soap formation on a wire frame, the area-minimization leaves the frame's size unchanged. For a bubble formation, the area-minimization leaves the volume of enclosed air unchanged. These properties are captured by approximating soap formations with configurations of two-dimensional surfaces in \mathbb{R}^3 that are *minimal*: their total area cannot be decreased by certain small deformations that leave their frame or enclosed volume fixed. These configurations are *almost minimal sets*.

The second part of the proof shows that almost minimal sets that satisfy (1) also satisfy (2) and (3). The final part shows that almost minimal sets exist and satisfy (1). Overall, the proof shows that almost minimal sets satisfy Plateau's three laws, which is relevant to soap formations because the defining property of almost minimal sets models the area-minimization principle that governs soap formations.

Pincock (2015b) claims that this proof explains the fact that soap formations satisfy Plateau's laws by showing that it abstractly depends on the fact that almost minimal sets are minimal. For this to provide an adequate understanding of mathematical explanation, we must have a decent understanding of the abstract dependence relation. Pincock is clear that it

is objective, non-causal, and *sui generis*, but going beyond this very general characterization raises serious problems.

Abstract dependence is supposed to obtain between a fact about certain physical systems, and a fact about some more abstract objects. To elucidate this idea, Pincock (2015b: 865–866) appeals to the relationship between types and tokens. A piece of music (type) and a particular performance of it (token) share many properties. For example, assuming the performance is faithful and successful, they share many of their aesthetic properties. Yet there are certain more specific properties the token alone has, such as having a particular spatial location. In this way, the piece of music is more abstract than the performance of it, and the latter is an instance of the former. Similarly, we can say that something is an instance of an almost minimal set just in case it has *minimality*, and other more specific properties besides.

So, a condition on physical-on-mathematical abstract dependence is that the physical systems are instances of the mathematical objects, in the above sense. This is not a sufficient condition. To see this, imagine an accurate plastic model of a soap formation created by a teacher by combining plastic surfaces so as to satisfy one of Plateau's laws. This is an instance of an almost minimal set, but the model satisfies Plateau's laws because the teacher made it so. According to Pincock (2015b: 866–867), for abstract dependence to obtain, the formation of the physical systems must be governed by a process relevant to the condition for being instances of the more abstract objects. The relevant fact about soap formations is eligible because soap formations are instances of almost minimal sets, and because they are formed by a process of area-minimization that is relevant to their being instances of almost minimal sets.

This gives us an idea of what abstract dependence requires on the physical side. But what about the mathematical side? We know that the mathematical objects must have the

physical systems as instances, but this fails to distinguish between many distinct candidates. There are at least three sources of underdetermination.⁷

Properties: There are various definitions of ‘minimality’ in terms of which the explanation can be run. For example, Guy David (2013) provides a variation on the definition offered by F. J. Almgren Jr. (1975), invoking different sets to fix the operative notion of a small deformation. David says ‘a few other variants exist, but they would not be significantly different for what we want to say here’ (2013: 77), where among what he wants to ‘say’ is Taylor’s proof. Even if these alternatives are extensionally equivalent, explanation is more fine-grained than extensional equivalence. Abstract dependence is supposed to select for a particular fact involving certain mathematical objects having certain properties. So, there are at least as many candidates as there are distinct properties in terms of which the explanation can be run.

Bearers: For a given property, we face the further choice of bearers. For example, we could restrict our attention to the almost minimal sets contained within a particular subregion of \mathbb{R}^3 , such as the interior of a particular sphere, with no overall effect on the explanation. There are at least as many further candidates for abstract grounds as there are distinct collections of bearers in terms of which we can run the explanation.

⁷ Pincock (2015b: 877) admits that it is possible to define distinct kinds of mathematical objects and relate them to soap formations in ways that mimic Taylor’s proof. Perhaps he has in mind one or more of the sources of underdetermination I identify, or perhaps he has in mind some further source. In any case, he recognizes this as a problem (see fn. 8), and suggests that the remedy is a theory about what abstract dependence is and how it is distributed. If my discussion in this section is right, no ameliorative theoretical moves are forthcoming. In another paper (2015a), where Pincock applies his account to explanations within mathematics, he identifies an analogue of the same problem. There he proposes that the abstract ground for a given explanandum is the least more abstract fact among the candidates (2015a: 12). This solution presupposes that the various candidates for abstract grounds are partially ordered by their abstractness. It is far from clear that they are, at least for the candidates I identify. However, even if they are, and there is a least more abstract fact in the offing, why choose this, as opposed to, say, the most more abstract fact? This solution seems unduly arbitrary.

Interpretations: Given a selection of properties and bearers, we face a further choice of interpretation. For example, there are many distinct set-theoretic models of \mathbb{R}^3 , each of which has its own collection of almost minimal sets. The mathematics of almost minimal sets (topology and geometric measure theory) is algebraic, meaning it has no intended model. In light of this, no particular model \mathbb{R}^3 has claim to being the more natural home for the explanation. So, there are at least as many further candidates for the abstract grounds as there are models of \mathbb{R}^3 .

These three sources of underdetermination are independent and cross-cutting. In light of them, our understanding of the abstract dependence relation seems inadequate. By ruling out the alternative options, I will now argue that the only sensible recourse is to claim that instances of abstract dependence select for a particular mathematical fact from among the eligible candidates.

As I see it, there are three alternative options. The first is to argue that each candidate abstract ground is partial, so that only by enumerating every candidate have we provided the full abstract explanation. On this view, Taylor's explanation is incomplete. This cannot be right. Once we have understood the explanation, we stand to learn nothing of explanatory value by considering other candidate abstract grounds. Moreover, a philosophical theory should not pass judgement on the success of a scientific explanation. If practitioners deem it successful, a naturalistic philosophical account should take its success as a datum. There is also the further difficulty that, on this view, the explanation looks to be impossible to complete.

The second option is to argue that each candidate abstract ground is complete. On this view, each instance of abstract dependence is a case of massive overdetermination of a particularly problematic kind. To illustrate, contrast the present case with the much-discussed

case of two people simultaneously throwing a rock through a pane glass window. Each rock-throw is sufficient to cause the glass to break all on its own, but both rocks hit the window at the same time, so both events cause the window to break. This is a case of causal overdetermination, but it is unproblematic for at least two reasons.

The first is that its occurrence does not involve systematic coincidence. Even if it is a coincidence that the rocks were thrown at exactly the same time, or that the rocks hit the window at the same time, one-off coincidences of this kind shouldn't concern us. The second is that, despite the overdetermination, the explanation of why the window breaks is 'causally satisfying: there is a precise account of the causal powers of both rocks, and of the individual contribution of each rock to the shattering of the window. Removing one rock-throw has an easily definable result: the window shatters with less force' (Bernstein 2016: 30).

In contrast, on the present proposal, every case of abstract explanation will involve massive and systematic overdetermination. Further, the explanation is not satisfying in above sense. We have no sense of what each abstract ground is contributing, nor how the explanandum would change were any one of them to be removed.

A further worry is that, if all it takes for something more abstract than the explanandum to be a complete abstract ground is that it satisfies one of a range of very general structural properties, then abstract grounds come too cheaply. This is at odds with the ontic aspirations of Pincock's account. Objective dependence relations should demand a lot of their relata. Causation demands physical or modal connections between its relata, and grounding demands more intimate metaphysical connections. On the present proposal, abstract dependence merely requires certain kinds of similarity. It is hard to understand how this would amount to objective dependence in any particular direction. Moreover, if abstract dependence comes too easily, we would expect it to crop up everywhere. As Pincock says, '[u]nless there is some principled way to constrain the proliferation of abstract dependence

relations, there will be too many of them and so the value of abstract explanations will be diluted' (2015b: 877).

The final option is to argue that it is indeterminate which mathematical fact is selected. To say that Taylor's proof explains why soap formations satisfy Plateau's laws by identifying something on which the explanandum does not determinately depend on, but also does not determinately not depend on, does not help us understand how this mathematical explanation works. This is surely a caricature of a range of worked-out views that may provide some understanding; but the onus is on the proponent of Pincock's account to provide such a worked-out view. Without one, tethering our understanding of mathematical explanation to our understanding of how ontic indeterminacy interacts with ontic dependence seems like a bad way to go.

In light of the inadequacy of the above three options, we are forced to conclude that an instance of abstract dependence selects for a particular mathematical fact from among the eligible candidates. But this incurs an explanatory debt. If abstract dependence selects for one among a range of eligible mathematical candidates, it seems there should be a reason why.⁸ We have seen that the conditions abstract dependence imposes on its relata do not account for this. I can think of two further ways of gaining understanding of a relation. First, we might identify symptoms of its obtaining. Pincock suggests that the novelty and informativeness of the characterization of soap formations as almost minimal sets may be symptomatic of abstract dependence (2015b: 877–878). However, explanations run in terms of any of the eligible candidates for the abstract grounds would be novel and informative in these ways, so these symptoms fail to address the present concern. Indeed, I can think of no symptoms that would.

⁸ Compare Benacerraf: '[i]f the numbers constitute one particular set of sets, and not another, then there must be arguments to indicate which' (1965: 58).

Second, we might provide an analysis of the relation in more familiar terms. Abstract dependence is *sui generis*, which rules out a reductive definition; but we may be able to provide an illuminating non-reductive characterization. Pincock's invoking of the instantiation relation, along with his description of abstract dependence as non-causal and objective is as close as he gets to such a characterization, and we have already seen that these descriptors are far from illuminating.

At this point, it may be tempting to say that there is something which accounts for why abstract dependence selects for the mathematical relata it does, though we may never be in a position to know it. This last draw amounts to a desperate assurance that the explanatory debt is settled, in spite of the lack of any reason to think so. An account on which we may never properly understand the operative dependence relation is a poor foundation on which to build an understanding of mathematical explanation. It seems our only recourse is to stipulate that, as a matter of brute fact, abstract dependence selects for one among the many candidate abstract grounds. This brings abstract dependence within reach of our understanding, to the extent that it places no important features of it beyond our ken. Ultimately, however, our understanding is no better off.

For a given instance of abstract dependence underlying a mathematical explanation, there will be a brute fact of the matter about which mathematical fact it selects for. Stipulating that there is nothing to explain here does not dispel the feeling that there is. It is a striking regularity that soap formations satisfy Plateau's laws by virtue of one mathematical fact, rather than any of the other eligible candidates. But, in taking it as brute, we relinquish any means by which we might illuminate it. The same goes for each putative instance of abstract explanation. Thus, adopting Pincock's position involves positing a range of striking yet inexplicable regularities. This is the brute regularities problem.

Let us assume that Taylor's proof succeeds in picking out the unique abstract ground of our explanandum. That is, it locates the right properties of the right mathematical objects. The proof is couched in terms of these properties of these objects in the first place because they provide good approximate models of soap formations, by bearing certain structural similarities to them. These structural similarities must therefore have been a reliable guide to which mathematical fact the explanandum abstractly depends on. But we have been forced to accept that abstract dependence selects for abstract grounds as a matter of brute fact. *Ipsa facto*, it does not select for abstract grounds in virtue of any structural similarity they bear to the physical systems whose properties they determine. But then Taylor's success in identifying the abstract ground of the fact that soap formations satisfy Plateau's laws, guided as it was by structural similarities, is a fluke. More generally, assuming there are supposed to be many more cases of abstract explanation, Pincock's account implies that the reliability with which practitioners identify abstract grounds is a massive cosmic coincidence. This is the empirical access problem.⁹

Because it faces these problems, Pincock's account fails to provide adequate understanding of mathematical explanation. The brute regularities problem shows that it offers no understanding of how abstract grounds are related to what they ground. The empirical access problem shows that it offers no understanding of how practitioners succeed in providing abstract explanations. There are mysteries precisely where an account of mathematical explanation should illuminate. Note that my arguments here are sensitive only to the fact that Pincock's view implies that physical phenomena bear an objective, *sui generis* dependence relation to pure mathematical facts. The peculiarities of Pincock's account are irrelevant. If other accounts of mathematical explanation imply the same, we should expect

⁹ This objection is reminiscent of Field's (1989: 25–30, 230–239) variation on the epistemological objection to mathematical platonism, which arguably improves on Benacerraf's (1973).

them to face the brute regularities and empirical access problems. Pincock's account is a confirming instance of PREMISE 1 for entirely general reasons.

III. COUNTERFACTUAL DEPENDENCE

The notion that scientific explanation involves tracing relations of counterfactual dependence is growing in popularity, due in no small part to James Woodward and Christopher Hitchcock's development of an influential counterfactual analysis of causal explanation in science.¹⁰ Many authors have since argued that we can generalize Woodward and Hitchcock's account in pursuit of a monist counterfactual theory that covers causal and non-causal explanations alike.¹¹ In this spirit, Baron, Colyvan, Ripley, and Povich offer counterfactual accounts of mathematical explanation.¹²

I am sympathetic to this movement. However, there is more than one way of extending the counterfactual theory to mathematical explanation. One might take the mathematics in mathematical explanations to identify mathematical facts on which their explananda counterfactually depend; or one might take it to help identify non-mathematical facts on which their explananda counterfactually depend. The aforementioned authors all develop the former option.¹³ This countermathematical account seems promising for bolstering EIA. However, I will argue that its proponents must posit and evidence the existence of a *sui generis* relation of physical-on-mathematical non-causal dependence. Because of this, for all the reasons detailed in §II, the countermathematical account faces the brute regularities and empirical access problems, and thus fails to provide an adequate understanding of mathematical explanation. In other words, it is a further confirming instance of PREMISE 1.

¹⁰ Hitchcock & Woodward 2003; Woodward 2003; Woodward & Hitchcock 2003.

¹¹ Strevens 2008; Saatsi & Pexton 2013; Rice 2015; Reutlinger 2016; Povich 2018.

¹² Baron *et al.* 2017; Baron 2019; Povich 2019.

¹³ Knowles & Saatsi 2019 develop the latter.

Baron *et al.* (2017) and Baron (2019) use the following case study.¹⁴ North American periodical cicadas lie dormant in larval form for either 13 or 17 years, then emerge to eat, mate, and die. Why 13 and 17, specifically? Background ecological constraints restrict cicadas' life-cycles to between 12 and 18 years. Within that range, periods that maximize the time between co-emergence with nearby periodical predators will be advantageous. The number of years between the co-emergence of two periodical organisms is equal to the least common multiple of their life-cycles in years, and this is maximized when these numbers are coprime. There are good reasons for thinking that any nearby predators will have life-cycle periods of less than 12 years, and prime numbers are coprime with all numbers smaller than themselves. We therefore expect there to have been evolutionary pressure for cicadas to evolve prime-numbered life-cycle periods, and 13- and 17-year periods, specifically.

On the countermathematical account, the above explanation works by identifying the mathematical facts on which the explanandum counterfactually depends. It does this by implicating specific countermathematical claims, such as the following:

CM : If 13 and 17 were not prime, cicadas would not have 13- or 17-year life-cycles.

To make sense of this, we at least need a procedure for evaluating countermathematicals like *CM*. Baron *et al.* (2017) provide a two-step procedure relied upon by each of the aforementioned proponents of the countermathematical account. We first imagine a scenario in which the antecedent obtains, keeping the rest of the scenario as similar to actuality as possible without giving rise to a contradiction in the neighbourhood. In other words, we change just enough to produce a fragment of mathematics in which the antecedent obtains consistently. This step is the *twiddle*. We then rely on relevant physical laws to ascertain how

¹⁴ Introduced to the literature by Baker 2005.

the mathematical changes ramify in the physical system. This step is the *ramification*. For example, for *CM*, we imagine a scenario in which the multiplication function is different, such that 2 and 6 are factors of 13, making any further changes required to produce a consistent fragment of arithmetic (twiddle). Finally, we appeal to the laws of evolutionary theory to see how this change ramifies (ramification).

On the countermathematical account, countermathematicals such as *CM* must come out as true via the above procedure. This requires the assumption that the twiddle not only changes the mathematical facts, but also changes certain physical facts along with it. Only then will the physical laws yield the truth of the consequent. This is a substantive assumption, and requires the positing of a relation between the mathematical and physical facts.

Recognising this, Baron *et al.* (2017: 9) suggest that structure-preserving mappings (homomorphisms) may suffice. They won't. There is a homomorphism between my fingers and my toes, but that doesn't give us any reason to think that, if I had nine fingers, I would have nine toes! The point generalizes. Homomorphisms are relations of similarity, grounded in the properties of their relata. Changing the properties of one relatum only serves to break the similarity; it does not force the other relatum to change along with it.

One might reply as follows. When we evaluate a counterfactual of any kind, we have to make decisions about what to hold fixed and what to vary. All the above objection shows is that we haven't held enough fixed. If we hold the homomorphism between the physical and the mathematical fixed, then the twiddle will ramify in the desired way.¹⁵ This move is methodologically suspect. To see why, it will be helpful to consider the evaluation of more humdrum counterfactuals. Suppose I almost drop a very fragile cup on my tiled kitchen floor, but only just manage to catch it. The following counterfactual is plausible: If I had not caught the cup, it would have smashed. Following the standard procedure, we evaluate this

¹⁵ This is exactly what Baron 2019: 6 and Baron *et al.* 2017: 9–10 suggest.

counterfactual by imagining a situation as similar as possible to actuality, but where the antecedent is true. In this situation, it seems, I drop the cup, and the cup smashes.

One might think that, in the course of my evaluation, I have decided to hold fixed a range of facts: the fact that the floor is tiled, the fact that the cup is fragile, the fact that a very soft rug didn't suddenly appear on the floor, the fact that an angel wasn't waiting to catch the cup in my place, and so on *ad nauseam*. But this is to misunderstand the procedure. It is not a case of deciding what to hold fixed to get the truth-value we want; it is a matter of trying to work out what would remain unchanged, if the antecedent were true. I left the abovementioned facts unaltered because, based on what we know about the world, there is good reason to believe that my failing to catch the cup would not change these things.

In contrast, we have seen that, if the mathematical and physical domains are related by mere homomorphism, there is good reason to think that the homomorphism would not survive certain changes to the mathematical domain. To respond to this by suggesting we hold the homomorphism fixed seems completely *ad hoc*. Worse, it is hard to see how countermathematicals whose truth is guaranteed in this way could bear any explanatory weight. We might just as well explain the redness of my socks by claiming that it counterfactually depends on the redness of my shirt. After all (holding their sameness in colour fixed), if my shirt were not red, my socks would not be, either.

The above shows that we need a stronger relation to underpin the truth of the relevant countermathematicals. In particular, we must posit a relation of physical-on-mathematical dependence, and evidence its existence. If we have reason to think such a relation obtains, we have reason to think that changes in the mathematical facts will ramify as desired. The dependence relation cannot be one that obtains too easily. For instance, if the primality of 13 and 17 determines the properties of any homomorphic physical system, then changes in these properties will have widespread ramifications, and there will be no interesting counterfactual

connection between the primality of 13 and 17 and the cicadas' life-cycle periods. The dependence must be more demanding, such that changes in the relevant properties of 13 and 17 only ramify in changes to the cicadas' life-cycles.¹⁶

There is no off-the-shelf dependence relation one can appeal to here: the remit is far too specific. So, the proponent of the countermathematical account must posit a *sui generis*, non-causal dependence relation, the existence of which is presumably evidenced by the success of the explanations amenable to analysis in terms of it. We saw in §II that this path leads to the brute regularities and empirical access problems. Just as in the soap formation case, there are many candidate mathematical facts eligible for the role of determining the cicadas' life-cycle durations. Again, there appear to be at least three sources of underdetermination.

Properties: The familiar definition of primality can be stated as follows. For $a \in \mathbb{Z}^+$ where $a > 1$, a is prime iff, for any $b, c \in \mathbb{Z}^+$, $bc = a$ only if $b = 1$ or $c = 1$. However, there is an alternative definition that runs as follows, where $x|y$ means x is a factor of y . For $a \in \mathbb{Z}^+$ where $a > 1$, a is prime iff, for any $b, c \in \mathbb{Z}^+$, $a|bc$ only if $a|b$ or $a|c$. In fact, these definitions pick out distinct properties. The first defines a special case of *irreducibility*, while the second defines a special case of *primality*. For the positive integers, the properties coincide; but in more abstract algebraic structures they come apart. So, that 13 and 17 are irreducible and that 13 and 17 are prime are distinct facts, both of which imply the relevant facts about LCM-maximization. The explanation runs equally well by appeal to either.

Bearers: For a chosen property, we face a further choice of bearers. For example, by measuring life-cycles in months, we uniformly multiply by 12. Such a uniform translation

¹⁶ Baron *et al.* (2017: 9–10; fn. 10) recognize this, which is perhaps why they opt for homomorphism.

will preserve the LCM-maximizing structure, so we can just as well take the explanation to work by implicating a countermathematical such as the following: If, in addition to 12(1) and 12(13), 12(13) had factors 12(2) and 12(6), then cicadas would not have 12(13)-month life-cycle periods.

Interpretations: There are infinitely many set-theoretic models of the positive integers, and it makes no difference whatsoever to the explanation if we interpret the mathematics as about one or other of these models.

For the same reasons outlined in §II, this radical underdetermination forces us to stipulate that, as a matter of brute fact, the dependence relation in question selects for a particular mathematical fact from among the many eligible candidates. This posits a striking yet inexplicable regularity, and there will be further striking yet inexplicable regularities associated with each mathematical explanation amenable to the countermathematical account. Thus, the countermathematical account posits a range of striking yet inexplicable regularities.

On the countermathematical account, the cicadas explanation works by identifying the mathematical facts on which the explanandum counterfactually depends. The reason 13 and 17 are appealed to in the first place is that they help to form an adequate model of the physical system. The considerations at play in the initial modelling must therefore be a reliable guide to what the explanandum counterfactually depends on. These considerations include our desire to capture the hypothesised structural relationship between the cicadas' life-cycles and the life-cycles of nearby predators, namely the minimization of co-emergence. Perhaps we also consider the year to be a biologically significant unit of time, which may influence our choice in units (Baker & Colyvan 2011: 329).

However, the counterfactual dependence in question is supported by a *sui generis* dependence relation that selects for its mathematical relata as a matter of brute fact, meaning the natural numbers 13 and 17 are not selected in virtue of their ability to capture structural features of the physical system, nor in virtue of the privileged status of the year. Biologists' success in identifying the mathematical facts on which the cicadas' life-cycle periods counterfactually depend is therefore a coincidence. The countermathematical account faces the brute regularities and empirical access problems, and so fails to provide an adequate understanding of mathematical explanation. It is a further confirming instance of PREMISE 1.

IV. CAUSAL RELEVANCE

Lyon (2012) analyses mathematical explanations in terms of Frank Jackson and Philip Pettit's (1990) theory of program explanation, with the explicit intention of supporting EIA. To illustrate program explanation, imagine that water in a sealed glass container is heated to boiling point. Why does the glass shatter? There are two kinds of explanation. We can gesture towards the underlying causal process, culminating in particular water molecules striking the glass with momenta collectively sufficient to break it. This is the *process explanation*. Or we can point to the water's being at 100°C. This is the *program explanation*. Jackson and Pettit's theory aims to show how these two explanations are related.

The relation is a modal one. The temperature property had to be realized by some arrangement of molecules and some distribution of momenta among them, such that some molecules or other would have struck the glass with momenta collectively sufficient to break it. Because we know the temperature property was instantiated by the water, we can be sure that some causal process or other produced the explanandum. We say that the temperature property *programs for* the causes of the explanandum, and is thereby *causally relevant* to it.

Why do we need program explanations, if we are so sure there are process explanations in the offing? The reason is twofold. First, underlying causal processes are often recondit, so it is beneficial to have a means of exploiting their existence for explanatory ends without having to describe them explicitly. Second, program explanations yield explanatory information that process explanations do not. Even if we were able to trace the trajectory of the molecules that shattered the glass, doing so would miss the fact that, even if these particular molecules were not responsible, some other molecules would have been. The program explanation improves on the process explanation by implying that, whatever the underlying causal goes on, so long as the temperature property was instantiated, some causal process or other will have culminated in the glass shattering. In this way, program explanations reveal the *modal robustness* of their explananda.

Lyon claims mathematical explanations work by locating a causally relevant property of mathematical objects. For example, in the cicadas case, the instantiation of *being prime* by 13 and 17 is supposed to program for the causes of the explanandum. Unfortunately, Lyon fails to explain how this might work. Nor does any promising account seem forthcoming (Saatsi 2012). Such an account requires appeal to a dependence relation to supply the requisite modal force. In the temperature example, we appeal to *realization*: we say that the temperature property *had to be realized by* something sufficient for the relevant causes to be instantiated. But the fact that 13 and 17 are prime does not obviously demand anything of cicadas, in the way that the temperature property demands something of the water.

The temperature case suggests that programming occurs internal to the physical system, from a higher-level property instantiated by it, to lower-level causes instantiated in it. If we are willing to give up on the contention that pure mathematical properties program for causes, there is a natural way of getting the primality of 13 and 17 involved with programming. There is a homomorphism between the physical system and the mathematical

domain that preserves the minimization of the frequency of co-emergence as the LCM-maximization of 13 and 17 with integers smaller than 12. This mapping is determined by our decision to measure lifecycle periods in years, along with the other background constraints imposed by the explanation. Call this mapping Φ . In virtue of it, the cicadas' life-cycle periods instantiate the higher-level, multiply-realizable properties *being mapped $_{\Phi}$ to 13* and *being mapped $_{\Phi}$ to 17*.

These properties tick all the boxes. Any period that instantiates one will be part of a system featuring a minimization of the frequency of overlap between it and shorter periods. Some causes or other will have led to this, so the instantiation of one of these properties programs for the causes of the explanandum. In particular, its stable instantiation within a biological system programs for evolutionary pressure towards that stability. Moreover, the existence of the positive integers is required to enter into the mapping on which the instantiation of these properties depends. So we have a metaphysic that explains how mathematical properties can be causally relevant to physical phenomena, and appears promising for supporting EIA.

This metaphysic does not entail that physical facts depend on mathematical facts. The instantiation of the impure mathematical programming properties depends on the overlap-minimizing structure of the physical system and the LCM-maximizing structure of the mathematical domain, respectively. Because of this, it is not a problem that there are many distinct mathematical candidates for capturing the relevant properties of the physical system. We can happily say that the physical system instantiates a distinct programming property for each candidate, since this will not result in any troubling systematic overdetermination. (Compare: the fact that the water in the temperature example has a distinct programming property corresponding to each temperature measurement scale does not over-determine the explanandum.) We are not forced to posit any brute regularities, so our success in identifying

the programming properties is not mysterious. The present proposal therefore appears to offer some understanding regarding how mathematical explanations work.

Unfortunately, despite initial appearances, this account is unfit to support EIA. In locating our causally relevant impure mathematical properties, we have inadvertently located a purely physical property that can do all of the desired explanatory work. We avoid the brute regularities and empirical access problems by accepting that the instantiation of the programming properties is partially grounded by the overlap-minimizing structure of the physical system. All the work done by the impure mathematical property can be done by this overlap-minimizing property. It is a multiply-realizable property of the physical system that programs for the causes of the explanandum, but its instantiation does not require the existence of any mathematical objects. Moreover, it is mutually-recognized among parties to the debate. Indeed, the proponent of Lyon's account needs it to ground the relevant mapping.

The critic of EIA can achieve the same level of understanding of mathematical explanation by appeal to the same theory of explanation, while claiming that the mathematics in mathematical explanations merely serves to represent a physical higher-level structural property of the physical system. Importantly, the failure to support EIA is a consequence of relinquishing the claim that the physical explananda of mathematical explanations depend on pure mathematical facts. Lyon's account confirms PREMISE 2.

V. A DILEMMA

Recall my master argument:

PREMISE 1

Any account on which mathematical explanations locate pure mathematical facts on which their physical explananda depend fails to provide an adequate understanding of mathematical explanation.

PREMISE 2

Any account of mathematical explanation apt to bolster EIA implies that mathematical explanations locate pure mathematical facts on which their physical explananda depend.

CONCLUSION

Any account of mathematical explanation apt to bolster EIA fails to provide an adequate understanding of mathematical explanation.

In §§II–IV, I identified two sources of confirmation for PREMISE 1, and one for PREMISE 2. I showed that positing a *sui generis* relation of physical-on-mathematical dependence fails to provide an adequate understanding of mathematical explanation (§II). I showed that the increasingly popular countermathematical account must posit a *sui generis* relation of physical-on-mathematical dependence, and so fails to provide an adequate understanding of mathematical explanation (§III). And I showed that, while the causal relevance account seems able to provide some understanding, it fails to support EIA precisely because it does not posit physical-on-mathematical dependence (§IV). These findings are significant in their own right. But one might think they only provide piecemeal support for CONCLUSION. Not so. My arguments are responsive to entirely general reasons. Capitalising on this, we can give a general argument in favour of CONCLUSION in the form of a dilemma.

Any account of mathematical explanation apt for supporting EIA must be ontic. It must characterize mathematical explanations as limning the network of objective dependencies in which the explanandum is embedded, and locate explanatory mathematical facts within this network.¹⁷ Such an account must choose between two options. First, take the mathematical facts invoked by mathematical explanations to depend on their physical

¹⁷ I take this to be obvious and will not argue for it here. If I am wrong, then CONCLUSION must be weakened as follows: Any ontic account of mathematical explanation apt to bolster EIA fails to provide an adequate understanding mathematical explanation. This is still a significant result

explananda. Second, take their physical explananda to depend on the mathematical facts they invoke.¹⁸

Platonic mathematical facts do not depend on contingent physical facts, so mounting the first horn with respect to a given explanation involves locating an impure mathematical fact that depends on the physical explanandum via an independently-established relationship. If anything about a mathematical domain is explanatory with respect to a physical phenomenon, it is its structure, and only if the physical phenomenon is related to the mathematical domain via some kind of mapping. So, our independently-established relationship will be some kind of mapping between the physical and the mathematical domain. However, mappings obtain in virtue of the structures exhibited by their relata. As such, any explanatory relationship the mathematical structure bears to the physical explanandum (via the mapping) will be mediated by the structure of the physical system, and the latter will be just as eligible for bearing the proposed explanatory relationship to the explanandum. So, there will be a mutually-recognized physical property that does the required explanatory work. Mounting this horn results in an account of mathematical explanation that cannot support EIA, whatever its other virtues.

Mounting the second horn with respect to a given explanation involves locating some mathematical fact on which the physical explanandum depends. To provide an adequate understanding of this dependence, the account will have to say something about how it selects for its mathematical relatum. Since mathematical objects are abstract, this will be in terms of the mathematical objects' fulfilment of a certain theoretical role, such as capturing structural features of the physical system. However, for any specified theoretical role, there will be many different collections of mathematical objects able to fill it. For the reasons given in §II, this forces us to stipulate that, as a matter of brute fact, the proposed dependence

¹⁸ I am omitting the complication that we can also choose whether the relevant dependence is partial or whole, which ultimately makes no difference to my argument.

selects for a particular mathematical relatum from among the eligible candidates. This generates the brute regularities and empirical access problems, and so mounting this horn destroys our understanding of mathematical explanation.

The first horn vindicates PREMISE 1, the second horn vindicates PREMISE 2, and our choice of horns seems to be a forced decision between exhaustive and mutually exclusive alternatives. Along with the three confirming instances identified in §§II–IV, this provides considerable (defeasible) support for CONCLUSION.

VI. THE SIREN CALL OF STRUCTURALISM

One might worry that my dilemma has limited scope. Perhaps the second horn only threatens a traditional form of platonism, on which the intrinsic natures of mathematical objects determine the mathematical relations they bear to one another. On another well-established form of platonism, the entities answering to our mathematical theories are abstract structures that are independent of, and explanatorily prior to, any possible instantiation of them by any particular system of objects. On this *structuralism*, a mathematical singular term refers to a position in a mathematical structure whose nature is exhausted by the relations it bears to other positions in the structure.¹⁹

At a distance, it appears that adopting structuralism will blunt the second horn of my dilemma. For a given explanation, if the supposedly different candidates for mathematical grounds are merely different instantiations of a unique mathematical structure, then we can say that the posited dependence relation selects for the structure itself, blocking the underdetermination that leads to the brute regularities and empirical access problems. However, on closer inspection, things are not so simple. Even if we assume that there is a

¹⁹ The structuralism considered here is developed and defended in detail in Shapiro 1997. There are many other ways of spelling out the intuitive idea that mathematics is a science of structure. While interesting and worthy of attention, they are not relevant to our discussion. See Reck & Price (2000) for an excellent survey.

unique mathematical structure picked out by each of our non-algebraic mathematical theories (the *uniqueness assumption*), we can only block two of the three sources of underdetermination I have identified. Worse still, there is reason to doubt the uniqueness assumption. If it is false, then we are no better off in adopting structuralism.

For concreteness, consider the cicada explanation outlined in §III, which appeals to properties of natural numbers. On the uniqueness assumption, there is a unique structure answering to our theory of natural numbers: *being a sequence closed under a successor function with a unique first element*. Natural numbers are positions in this structure. When a system of objects instantiates the natural-number structure, objects in that system ‘occupy’ the natural-number positions. Now, a collection of positions in a structure can be treated as a system of objects which can itself instantiate a distinct structure. So, we can consider various different sequences of positions in the structure described by set theory as instantiating the natural-number structure. On structuralism, this is precisely what we do when we provide set-theoretic models of the natural numbers. On this view, the existence of many different set-theoretic models of the natural numbers does not underdetermine the *interpretation* of the mathematics in the cicadas explanation. The one and only correct interpretation of this mathematics is in terms of the natural-number structure itself.

A similar move blocks underdetermination in our choice of objects within a given interpretation. I suggested that, by applying a uniform translation, we end up with a distinct collection of mathematical objects with which we can run the explanation equally well. However, we can only do this because the result of a uniform translation of the natural numbers is a sequence of natural numbers that instantiates the natural-number structure. In the example I gave, 12 occupies the 1-position, 24 occupies the 2-position, and so on, where the successor function s is defined as $s(x) = x + 12$. Measuring life-cycles in months does not yield a distinct candidate for the mathematical ground of the explanandum. The explanatory

properties will be those had by the 13-position and 17-position in the natural-number structure, under any uniform translation.

Unfortunately, none of this helps with the third source of underdetermination. We have the choice of appealing to the fact that 13 and 17 are irreducible or the fact that 13 and 17 are prime. Both can be used to run the explanation equally well, yet these are distinct properties. The uniqueness assumption doesn't help. Nor should we think this is a special case. We saw in §II that there are different properties we might appeal to in running the explanation of why soap formations satisfy Plateau's laws. More generally, for a given physical phenomenon, it seems naïve to assume there is a unique mathematical property in terms of which it is best modelled and explained. So, there remains a source of underdetermination on which to hang my arguments from §II, keeping the second horn of my dilemma sharp. One might nevertheless see this as progress. Structuralism has blocked two out of three sources of underdetermination. Perhaps with some further ingenuity, we can block the third. But such optimism is premature.

We have seen that two sources of underdetermination can be blocked on the uniqueness assumption. But the uniqueness assumption is controversial. In fact, the very ingredients of structuralism that allow us to block underdetermination, on the uniqueness assumption, provide compelling reasons to doubt it. The ingredients are the flexible nature of mathematical objects, as both positions in structures and objects able to occupy those positions. It is natural to individuate structures on the basis of isomorphism, so that isomorphic structures are identical, and non-isomorphic structures are distinct. However, if we can make sense of a position of one structure occupying a position of another structure, then we can make sense of a permutation of the positions of a given structure. For example, take the natural-number structure and permute the first and second position. We end up with a new structure that is isomorphic to the original. In this way, it seems, there are infinitely-

many distinct yet isomorphic structures equally eligible for being the natural-number structure. In light of arguments like this, Stewart Shapiro, a key defender of the present form of structuralism, concedes that ‘an ontological realist cannot simply stipulate that there is at most one structure for each isomorphism type’ (2006: 143).²⁰

Structuralism, combined with an account of mathematical explanation that posits physical-on-mathematical dependence, faces underdetermination in which mathematical property the posited dependence selects for, and which of the infinitely many structures answering to the relevant mathematical theory is selected for. These sources of underdetermination are independent and cross-cutting. Following my arguments in §II, this combination of views faces the brute regularities and empirical access problems, so adopting structuralism fails to blunt the second horn of my dilemma.

VII. CONCLUSIONS

I do not want to overreach. I have not shown that EIA fails. Nor have I shown that it is impossible to account for mathematical explanation in a way that supports platonism. I have provided considerable (albeit defeasible) support for the following claim.

CONCLUSION

Any account of mathematical explanation apt to bolster EIA fails to provide an adequate understanding of mathematical explanation.

In doing so, I have dispelled the appearance that the balance of evidence currently favours EIA. As things now stand, the balance of evidence counts against EIA. I mentioned in §I that there are accounts of mathematical explanation that are promising for undermining EIA. Robert Knowles and Juha Saatsi (2019) provide an account of mathematical explanation according to which they explain by locating physical facts on which their explananda

²⁰ See Assadian 2018 for more on this and related problems for structuralism.

counterfactually depend. Mary Leng (2012) provides an account according to which mathematical explanations identify structural features of the physical system in virtue of which their explananda had to occur. And, in light of our discussion in §IV, we can add Lyon's (2012) causal relevance account to this list.

Perhaps these are genuine rivals. Perhaps they offer different but compatible theoretical perspectives on the same phenomenon. Either way, they are compatible with the view that the mathematics in mathematical explanations merely serves to represent otherwise elusive explanatory physical facts. As yet, there is no reason to think they run into difficulties that undermine their capacity to provide genuine understanding of how mathematical explanations work. So, on balance, we have reason to side with the critics of EIA.²¹

REFERENCES

- Almgren, F. J. (1975) 'Existence and Regularity Almost Everywhere of Solutions to Elliptic Variational Problems with Constraints', *Bulletin of the American Mathematical Society*, 81/1: 151–4.
- Almgren, F. J. Jr and Taylor, J. E. (1976) 'The Geometry of Soap Films and Soap Bubbles', *Scientific American*, 235/1: 82–93.
- Assadian, B. (2018) 'The Semantic Plights of the Ante-Rem Structuralist', *Philosophical Studies*, 175/12: 3195–215.
- Audi, P. (2012) 'Grounding: Toward a Theory of the In-Virtue-Of Relation', *Journal of Philosophy*, 109/12: 685–711.

²¹ I gratefully acknowledge funding from the British Academy, without which this paper would not have been possible. Thanks to Scott Shalkowski, Jess Isserow, Will Gamester, Juha Saatsi, Seamus Bradley, David Liggins, Chris Daly, and two anonymous referees from this journal. Their comments on various iterations of this paper were invaluable. I am grateful to Christopher Pincock for enlightening discussion. Thanks also to audiences at the University of Cambridge, the University of Manchester, the University of Chieti-Pescara, and the University of Santiago de Compostela for their insightful questions. Special thanks to Emily Shercliff for patient discussion and unwavering support.

- Baker, A. (2005) 'Are there Genuine Mathematical Explanations of Physical Phenomena?', *Mind*, 114/454: 223–38.
- Baker, A. (2009) 'Mathematical Explanation in Science', *British Journal for the Philosophy of Science*, 60/3: 611–33.
- Baker, A. (2017) 'Mathematics and Explanatory Generality', *Philosophia Mathematica*, 25/2: 194–209.
- Baker, A. and Colyvan, M. (2011) 'Indexing and Mathematical Explanation', *Philosophia Mathematica*, 19/3: 323–34.
- Baron, S. (2020) 'Counterfactual Scheming', *Mind*, 129/514: 535–62.
- Baron, S., Colyvan, M., and Ripley, D. (2017) 'How Mathematics Can Make a Difference', *Philosophers' Imprint*, 17/3: 1–19.
- Benacerraf, P. (1965) 'What Numbers Could not Be', *The Philosophical Review*, 74/1: 47–73.
- Benacerraf, P. (1973) 'Mathematical Truth', *Journal of Philosophy*, 70/19: 661–79.
- Bernstein, S. (2016) 'Overdetermination Underdetermined', *Erkenntnis*, 81/1: 17–40.
- Clark, M. J. and Liggins, D. (2012) 'Recent Work on Grounding', *Analysis*, 72/4: 812–23.
- Colyvan, M. (2002) 'Mathematics and Aesthetic Considerations in Science', *Mind*, 111/441: 69–74.
- Colyvan, M. (2013) 'Road Work Ahead: Heavy Machinery on the Easy Road', *Mind*, 121/484: 1031–46.
- Daly, C. and Langford, S. (2009) 'Mathematical Explanation and Indispensability Arguments', *Philosophical Quarterly*, 59/237: 641–58.
- David, G. (2013) 'Regularity of Minimal and Almost Minimal Sets and Cones: J. Taylor's Theorem for Beginners', in G. Dafni, R. J. McCann, and A. Stancu, A. (eds.) *Analysis*

- and Geometry of Metric Measure Spaces*, 67–118. Providence, RI: American Mathematical Society.
- Field, H. (1989) *Realism, Mathematics & Modality*. Oxford: Blackwell Publishing.
- Hitchcock, C. and Woodward, J. (2003) ‘Explanatory Generalisations, Part II: Plumbing Explanatory Depth’, *Noûs*, 37/2: 181–99.
- Jackson, F. and Pettit, P. (1990) ‘Program Explanation: A General Perspective’, *Analysis*, 50/2: 107–17.
- Knowles, R. and Saatsi, J. (2019) ‘Mathematics and Explanatory Generality: Nothing but Cognitive Salience’, *Erkenntnis*, DOI: <https://doi.org/10.1007/s10670-019-00146-x>.
- Leng, M. (2010) *Mathematics and Reality*. Oxford: Oxford University Press.
- Leng, M. (2012) ‘Taking it Easy: A Response to Colyvan’, *Mind*, 121/484: 983–95.
- Lyon, A. (2012) ‘Mathematical Explanations of Empirical Facts, and Mathematical Realism’, *Australasian Journal of Philosophy*, 90/3: 559–78.
- Maddy, P. (2005) ‘Three Forms of Naturalism’, in S. Shapiro, (ed.) *The Oxford Handbook of Philosophy of Mathematics and Logic*, 437–59. New York: Oxford University Press.
- Morrison, J. (2012) ‘Evidential Holism and Indispensability Arguments’, *Erkenntnis*, 76/2: 263–78.
- Pincock, C. (2015a) ‘The Unsolvability of the Quintic: A Case Study in Abstract Mathematical Explanation’, *Philosophers’ Imprint*, 15/3: 1–19.
- Pincock, C. (2015b) ‘Abstract Explanations in Science’, *British Journal for the Philosophy of Science*, 66/4: 857–82.
- Povich, M. (2018) ‘Minimal Models and the Generalized Ontic Conception of Scientific Explanation’, *British Journal of the Philosophy of Science*, 69/1: 117–37.

- Povich, M. (2019) 'The Narrow Ontic Counterfactual Account of Distinctively Mathematical Explanation', *British Journal for the Philosophy of Science*, DOI: <https://doi.org/10.1093/bjps/axz008>.
- Putnam, H. (1971) *Philosophy of Logic*. New York: Harper.
- Quine, W. V. (1948) 'On What There Is', *The Review of Metaphysics*, 2/5: 21–38.
- Raven, M. J. (2015) 'Ground', *Philosophy Compass*, 10/5: 322–33.
- Reck, E. H. and Price, M. P. (2000) 'Structures and Structuralism in Contemporary Philosophy of Mathematics', *Synthese*, 125/3: 342–83.
- Reutlinger, A. (2016) 'Is There a Monist Theory of Causal and Noncausal Explanations? The Counterfactual Theory of Scientific Explanation', *Philosophy of Science*, 83/5: 733–45.
- Rice, C. (2015) 'Moving Beyond Causes: Optimality Models and Scientific Explanation', *Noûs*, 49/3: 589–615.
- Rizza, D. (2011) 'Magicicada, Mathematical Explanation, and Mathematical Realism', *Erkenntnis*, 74/1: 101–14.
- Rosen, G. (2010) 'Metaphysical Dependence: Grounding and Reduction', in R. Hale and A. Hoffman (eds.) *Modality: Metaphysics, Logic, and Epistemology*, 109–36. Oxford: Oxford University Press.
- Saatsi, J. (2011) 'The Enhanced Indispensability Argument: Representational versus Explanatory Role of Mathematics in Science', *The British Journal for the Philosophy of Science*, 62/1: 143–54.
- Saatsi, J. (2012) 'Mathematics and Program Explanations', *Australasian Journal of Philosophy*, 90/3: 579–84.

- Saatsi, J. and Pexton, M. (2013) 'Reassessing Woodward's Account of Explanation: Regularities, Counterfactuals, and Noncausal Explanations', *Philosophy of Science*, 80/5: 613–24.
- Schaffer, J. (2009) 'On What Grounds What', in D. J. Chalmers, D. Manley, and R. Wasserman (eds.) *Metametaphysics: New Essays on the Foundations of Ontology*, 347–83. Oxford: Clarendon Press.
- Shapiro, S. (1997) *Philosophy of Mathematics: Structure and Ontology*. Oxford: Oxford University Press.
- Shapiro, S. (2006) 'Structure and Identity', in F. MacBride (ed.) *Identity and Modality*, 34–69. Oxford: Oxford University Press.
- Sober, E. (1993) 'Mathematics and Indispensability', *Philosophical Review*, 102/1: 35–57.
- Strevens, M. (2008) *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- Wilson, J. (2014) 'No Work for a Theory of Grounding', *Inquiry*, 57/5–6: 1–45.
- Woodward, J. (2003) *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
- Woodward, J. and Hitchcock, C. (2003) 'Explanatory Generalizations, Part I: A Counterfactual Account', *Noûs*, 37/1: 1–24.
- Yablo, S. (2012) 'Explanation, Extrapolation, and Existence', *Mind*, 121/484: 1007–29.

University of Leeds, UK