

HOW
MATHEMATICS
CAN MAKE A
DIFFERENCE

*Sam Baron, Mark Colyvan, David
Ripley*

University of Western Australia

University of Sydney

University of Connecticut

© 2016, Sam Baron, Mark Colyvan, and David Ripley

www.philosophersimprint.org/017/003

Abstract

Standard approaches to counterfactuals in the philosophy of explanation are geared toward causal explanation. We show how to extend the counterfactual theory of explanation to non-causal cases, involving extra-mathematical explanation: the explanation of physical facts (in part) by mathematical facts. Using a structural-equation framework, we model impossible perturbations to mathematics and the resulting differences made to physical explananda in two important cases of extra-mathematical explanation. We then address some objections to our approach.

1. Introduction

Ideally, a philosophical theory of explanation ought to be fully general, providing an account of explanation wherever it is found (26). Explanations, however, occur in a wide variety of places (at the very least, in the sciences, in mathematics, and in the humanities), and it's not obvious how to come up with a theory of explanation that can cover such a wide range. Here, we'll attempt a small step in the direction of unification, aiming at the area between science and mathematics. What we seek is a theory of explanation that covers cases of extra-mathematical explanation — explanations of non-mathematical or "empirical" facts by mathematical ones. Extra-mathematical explanation is an important phenomenon in its own right, and has risen to prominence recently due to the role that it plays in arguments for the existence of mathematical objects.¹ To date, though, there has been no attempt to roll explanation of this kind into a general theory of scientific explanation.

In this paper we assume that there are genuine cases of extra-mathematical explanation in science, and seek to develop a theoretical understanding of them. In particular, we aim to lay the foundations for generalising a counterfactual theory of explanation in order to cover extra-mathematical cases of scientific explanation. Accordingly, the aims of this paper are largely orthogonal to the ongoing debate

1. See, for example, (3; 5; 13; 15; 37).

over the existence of mathematical objects, which leads some to call into question the legitimacy of extra-mathematical explanations. We begin with a very general characterisation of counterfactual reasoning, before addressing some basic worries with the extension of this reasoning to include mathematics. Following that, we sketch a procedure for evaluating the counterfactuals implicated in cases of extra-mathematical explanation, one that yields their correct truth-values. Finally, we show how to model cases of extra-mathematical explanation using the structural equation framework.

2. Counterfactuals and Counterpossibles

From the outset it is important to note that this is not a paper about the semantics of counterfactuals. We will not be providing an account of how to evaluate the counterfactuals operative within cases of extra-mathematical explanation by situating them within an interventionist framework, or within a Lewis-Stalnaker semantics. Instead, we abstract away from the particular semantic theories that have been offered to date, and focus on three notions that surface, in some way or other, in all such theories. These are the notions of “holding fixed”, “twiddling” and “ramifying”. Very roughly, in order to evaluate a counterfactual conditional, we take three steps. First, we choose some class of facts to be invariant under counterfactual variation; these facts are “held fixed”. Second, we allow the facts corresponding to the antecedent of the counterfactual of interest to vary: we “twiddle” those facts. Third, we consider the downstream implications for the facts that we are not holding fixed of letting the antecedent vary: we see how the twiddle “ramifies” through these facts.

So, for instance, consider the following familiar counterfactual: if Suzy had not thrown the rock, the bottle would not have broken. First, we make a choice about what to hold fixed. We might, for instance, choose to hold fixed the past up to a time just before Suzy throws (as David Lewis [20] suggests). Second, we allow Suzy’s throw to vary by supposing that Suzy did not throw. So we choose not to hold Suzy’s throw fixed, choosing instead to twiddle it. Third, we carry the choice

about what we have held fixed and what we have allowed to vary through the remaining facts we have not held fixed, looking out, in particular, for what happens to the breaking of the bottle. So, for instance, if we hold fixed the past but not the future, we can then let Suzy’s throw vary and see what ramifications this variance has on the future of the bottle. Of course, this is not the only choice about what to hold fixed in this case, and we have not offered any specific recommendations on what to hold fixed or why. The point is just that some choice of what to hold fixed and what to twiddle must be made and the ramifications of these choices must be considered.

Note that the choice of what to hold fixed is undoubtedly sensitive both to context and to whatever broad reasons one might have for evaluating a counterfactual in the first place (roughly: to the pragmatics of counterfactuals). Note also that we cannot hold everything fixed. For instance, suppose we hold fixed all of the facts about the past, the future, and Suzy’s throw. Then there will be no consistent way to make sense of varying the facts about Suzy’s throw *given that we’re holding those facts fixed*. Or, to take another example, suppose we hold fixed facts about the past and facts about the future, but allow Suzy’s throw to vary. Clearly we have held fixed too much, for there will be no ramifications at all of Suzy’s throw through the “free facts”. Or, at least, there will be no sense in which variations in Suzy’s throw consistently ramify through facts about the future, which are being treated as free in this example, given, again, that we’re holding those facts fixed. The point is that we must always be careful to specify what we’re holding fixed when attempting to evaluate a counterfactual of any stripe.

These points about what, and what not, to hold fixed have direct bearing on the present project. For the task of deciding what to hold fixed is part of a more general project of developing a reasonable procedure for evaluating counterfactuals. Our goal in this section is to sketch such a procedure for extra-mathematical cases.² The procedure we out-

2. Baker (2) provides a useful discussion of counterfactuals and mathematics. However, Baker does not press the relevant counterfactuals into service in

line provides an account of what to hold fixed in extra-mathematical cases and shows how to determine the ramifications of any mathematical twiddling. First, however, we must address two immediate concerns one might have about the very concept of holding mathematical facts under counterfactual supposition.

2.1 Mathematical Twiddling

You may well be worried about twiddling mathematics. So we pause here to survey the situation and show that what we're proposing is a simple application of familiar ideas about counterfactuals. Just as in non-mathematical cases, we must hold some things fixed while allowing other things to vary. If we hold too much fixed, we will run into needless contradiction; too little, and we will miss the important connections we're after. Before we look at how this plays out in the cases we're interested in, we'll look at a few ways in which things might go wrong. Counterfactual reasoning about mathematics is perhaps less familiar than other kinds of counterfactual reasoning, so some examples will help us get a feel for them. But in short, nothing different is going on here; we are merely applying a standard way of thinking about counterfactuals to mathematical counterfactuals, just as we might apply it to biological counterfactuals or theological counterfactuals or whatever.

2.1.1 Don't hold too much fixed

Mathematical facts tend to be rather tightly integrated with each other. Because of this, it is important, when we twiddle one fact, not to hold too many others fixed. To see what might go wrong if we do, consider the following kind of reasoning:

What would happen if $5 + 7$ were 13? Well, since $5 + 7$ is 12, 12 would be 13. Subtracting 12 from both sides, 0 would be 1. But then all numbers would be identical: for any m, n , $m \times 0 =$

grounding a theory of extra-mathematical explanation, as we do here.

$n \times 0$, and since 0 would be 1 under the supposition in question, we have it that $m = n$. And if all numbers were identical, any number of horrible things would follow.

(A version of this objection can be found in [41, p. 172].) We of course agree that if all numbers were identical, any number of horrible things would follow. The problem with the above reasoning is that it imports too much mathematical *reality* into what is, after all, an *unreal* setting. The whole point of counterfactual thinking is to consider cases in which things are *different* from how they actually are. Compare the following reasoning:

What would have happened if John had worn a blue jacket, instead of the green one he actually wore for the photo shoot? Well, since he wore a green jacket, the blue jacket would be a green jacket, and this would mean that green is blue.

This reasoning should convince nobody. The problem is clear: too much of what is *actually* the case (the actual colour of a particular jacket) has been imported into the *non-actual* situation under consideration. What really would have happened if John had worn a blue jacket is that he *wouldn't* have been wearing the green one; rather than rearranging the colour wheel, we should rearrange the locations of jackets.

Of course, we are not *barred* from holding fixed that John wore the green jacket; if we are genuinely interested in what would happen if his green jacket were at the same time a blue jacket, then the above reasoning is fine. The point is that, in almost all cases, this is *not* what we're concerned with. Absent some very special conversational context, we should not hold so much fixed.

The same goes for the mathematical case. The mathematical reasoning above holds too much fixed; just as the full state of John's sartorial display wouldn't still hold if he wore a blue jacket, not all of the mathematics appealed to in the given reasoning would still hold if $5 + 7$ were 13. Again, it is not as if we are *barred* from holding the remainder of the mathematics fixed; it is just that, absent some special context, it

is almost certainly not what we're really after, if we're wondering what things would be like were $5 + 7$ to be 13.

There may seem to be a remaining difference. It is clear just what would have given way in the jacket example: if John wore the blue jacket, the green jacket would be elsewhere. But it is not clear just what would have given way in the mathematical case: would $5 + 7$ not have been 12, or would cancellation via subtraction have stopped working, or would it no longer be the case that $m \times 0 = n \times 0$ for all m, n , or what? To the extent that this is a real difference between the cases (rather than, say, a difference in the degree of detail with which the cases are presented), we do not think it undermines our point. If anything, it reinforces it: evaluating counterfactuals requires *decisions* about what to hold fixed. If we hold fixed that $5 + 7$ is 12 while entertaining its being 13 as well, then something else must give way. Even if we turn out to have *more* decisions to make, or if they are more difficult, in mathematical cases (and it is not clear to us that this is so), they are nonetheless no different in kind.

2.1.2 *Necessity*

We suppose throughout this paper that mathematical truths are *necessarily* true. Our proposed mathematical twiddles, then, are impossible. In conversation, we have sometimes encountered this as an objection to our proposal, so we pause here to sketch our take on the situation. As far as we can see, there are two possible concerns, only one of which is worth taking seriously.

If the concern is one about our proposing to *do* something impossible, then we can set the worry quickly aside; it is based on a misunderstanding. 17 is prime, and so 17 is necessarily prime. If we were to consider what would happen if 17 were composite, our proposal is to 'twiddle' this fact. But of course we are not proposing to *make* 17 composite to see what happens. As nice as it might be to be able to run such an experiment, we can't. Nor can anyone else. And of course the difficulty isn't mere lack of funding; it simply can't be done. But this is

no different from non-mathematical cases. It is not practically possible to change the past, and yet we can explore counterfactuals with false antecedents about the past. Whatever twiddling is, it does not require actually changing the relevant facts, nor does it require being *able* to change them. This is all as it usually is.

There is another worry in the area, though, that is more serious. This is that, owing to the necessity of mathematics, we cannot even *entertain* what it would be like for mathematics to be any way other than the way it is. Although this is a serious worry, here we can do no better than provide a flatfooted response. We deny that only possible things are entertainable. We can entertain water being H_3O ; we can entertain being something other than human; and we can entertain 17's being composite. We can entertain Goldbach's conjecture being true, and we can entertain its being false, although one of these is impossible. And so on.

To say more here would be to wade into a debate that, while worthwhile, would take us much too far afield. So, recognizing that we have stuck our necks out, we will leave the matter there.³

2.1.3 *The metalinguistic worry*

Finally, we consider what is possibly the most common worry about twiddling mathematical facts. This is the worry that we are not really twiddling mathematical facts at all, but only twiddling how we describe the mathematics, which itself remains as it was. When we purport to consider a situation in which 17 is composite, this worry has it, we are really only considering a situation in which '17' refers to a composite number, rather than referring to 17 (or one in which 'composite' picks out a property 17 actually has, rather than the property

3. It may be that some reluctance to believe in our ability to entertain the impossible is, at least in part, motivated by a lack of theoretical resources for understanding such an ability. This paper, we hope, helps provide some such resources. But there are other motivations as well, and those we will not address here. For more on the relationship between conceivability and modal notions, see (12; 44).

of being composite). So while we are trying to twiddle mathematics, we are instead only twiddling linguistic practice.⁴

This worry, we think, loses much of what plausibility it has when we turn to a wider variety of cases. Just as we can consider what would be the case if 17 were composite, we can consider what would be the case if the Riemann hypothesis were true, and what would be the case if it were false. One of these is impossible. The metalinguistic worry accounts for whichever one is impossible by claiming that, in that one, we are really considering what would be the case if ‘the Riemann hypothesis’ referred to a different claim, rather than the Riemann hypothesis. But that’s clearly not what we are doing; we are simply considering the truth or falsity of the Riemann hypothesis itself. The metalinguistic hypothesis cannot accommodate this. Nonetheless, there is a worthwhile question here: *how* do we manage to continue to refer to things like 17 and the Riemann hypothesis, even while supposing they are other than they are?

As it happens, there is a parallel question about the usual counterfactual cases. Let’s turn to these for a moment, so we can see whether the usual responses to this worry will apply in mathematical cases as well. In non-mathematical counterfactual cases, the worry would look like this: when we “twiddle”, say, whether Malcolm Turnbull became Prime Minister (PM) of Australia, we are not really considering what would have happened if Malcolm Turnbull himself had not become PM. Instead, we are merely considering what would have happened if we had used ‘Malcolm Turnbull’ to refer to someone who is not the PM, rather than to Turnbull himself (e.g. to Joe Hockey).

We do not nowadays take this kind of worry particularly seriously. But it would be a mistake to think that this is because it is obvious how to respond; it is not. We don’t worry like this about counterfactuals nowadays, because of two broad families of responses we have

4. This worry and the worry discussed in §2.1.1 are in serious tension with each other, but we have encountered both in conversation and think they are both worth addressing.

inherited. These responses, we think, are good. But they were not obvious. We pause to rehearse them ever so briefly.

How can it be that we are still talking about *Turnbull*, even in considering him as other than PM? On one kind of response, this is because the person in the counterfactual situation we are considering bears enough similarities to the actual Turnbull: perhaps, for example, they physically resemble him and have similar political credentials. On another kind of response, we simply *stipulate* that we are talking about Turnbull; it is part of the specification of the counterfactual situation we are considering that it is one in which *Turnbull himself* is not PM. A situation in which Turnbull still is PM, but in which we use ‘Malcolm Turnbull’ differently, is simply a different situation than the one that’s been specified — our stipulation is about political facts, not linguistic ones.

For our purposes, either of these responses, or any combination of them, will do. Both adapt completely without difficulty to the mathematical case. First, the similarity approach: even if 17 were composite, it would still have many of the properties that it in fact has. For example, it would still occur between 16 and 18; not all of 17’s properties are so closely tied to its primeness. As another example, it would still be the number of the Star in a standard tarot deck; not all of 17’s properties are mathematical properties.

Second, the stipulation approach: in exploring what would happen if 17 were composite, there is no possibility that we’re talking about something other than 17 itself, because it is simply part of the specification of the case to be considered that we are talking about 17. As Kripke (19) says (emphasis in original):

“Possible worlds” are *stipulated*, not *discovered* by powerful telescopes. There is no reason why we cannot *stipulate* that, in talking about what would have happened to Nixon in a certain counterfactual situation, we are talking about what would have happened to *him* [p. 44].

Nor do any new reasons become available because we are talking

about 17 rather than Nixon, or entertaining an impossible situation rather than a possible one.⁵

Whichever way the metalinguistic misinterpretation is to be avoided in ordinary counterfactual cases, it can be avoided in the same way in mathematical counterfactuals. Both standard responses carry over unmodified: first, many of the properties of mathematical objects will remain unchanged in the counterfactual circumstances we are entertaining; and second, our stipulations are mathematical in nature, not linguistic.⁶

2.2 *Holding Fixed Within Mathematics*

We turn now to providing a procedure for evaluating the counterfactuals implicated in extra-mathematical explanations. It is important to be clear about what we are *not* trying to do. We're not trying to argue that the counterfactuals implicated in cases of extra-mathematical explanation are true. Rather, we're just going to assume that these counterfactuals are true and then give a way of evaluating these counterfactuals that yields their correct truth-values. This will then be parlayed into a general procedure for evaluating the counterfactuals in these cases.

In what follows, it will be useful to have a basic case of extra-mathematical explanation on hand, along with some candidate counterfactuals for that case. As our example we will use what is perhaps

5. Of course, whether we *can* entertain an impossible situation remains contentious. But we have already taken a stand on that front (see §2.1.2).

6. A related objection holds that the properties of mathematical objects are essential properties. According to this line of thought, 17 is essentially prime. If this is right, then a fully Kripkean stipulation strategy won't help, since we can't stipulate our way around essences. This requires more discussion than we can give it here, but here are two lines of thought in response. First, our original objection to the metalinguistic strategy stands even in the face of this objection: it still cannot account for considerations about the truth or falsehood of the Riemann hypothesis. So this is not a way of rescuing the metalinguistic hypothesis, only a way of narrowing the options for explaining how reference can work in these cases. Second, it is far from clear that all properties had by mathematical objects, even all mathematical properties, are essential to those objects (4).

the most familiar case of extra-mathematical explanation: Alan Baker's case of the North American cicadas (3).

There are two sub-species of North American periodical cicadas that have prime-numbered life cycles of 13 and 17 years respectively. Why those life-cycle lengths and not others? The explanation we explore appeals to an optimality model that has four components: (i) a range of ecological constraints that restrict the life-cycle length of the cicadas to within the range 12–18 years; (ii) the assumed presence of predators with periodical life cycles; (iii) facts about 13 and 17, in particular that they're both prime numbers; and (iv) a mathematical fact regarding primes and common multiples. The model tells us that the optimal way for an organism with a periodic life cycle to avoid predators with periodic life cycles is for that organism to possess a prime-numbered life cycle. That's because prime-numbered life cycles minimise the chance of overlapping with periodical predators possessing numerically nearby life-cycle lengths. For instance, a cicada with a 14 year life cycle will overlap with predators possessing 1-, 2-, 7-, and 14-year life cycles. A 13-year cicada, by contrast, will overlap with predators possessing life-cycle lengths of 1 and 13 only (and, in the 13 case, only when the life cycles are synchronised). That prime numbers have this property is contained within the mathematical fact in (iv).

The most straightforward way to capture the explanatory role played by mathematics in this case in counterfactual terms is to focus on twiddling the properties of the numbers involved: 13 and 17. If, say, 13 had been different — in particular, if it were not prime — then these cicadas would not have had 13-year life-cycle lengths. This is because (iii) would no longer be true, and so could not combine with (i), (ii), and (iv) to deliver the result that having 13-year life cycles is an optimal way to avoid predation. Because of this:

(CF₁) If, in addition to 13 and 1, 13 had the factors 2 and 6, North American periodical cicadas would not have 13-year life cycles.

Similarly:

(CF₂) If, in addition to 17 and 1, 17 had the factors 2 and 6, North American periodical cicadas would not have 17-year life cycles.

Recovering the truth of (CF₁) and (CF₂) is only half of the battle. We must also recover the falsity of counterfactuals like:

(CF₃) If, in addition to 4 and 3, 12 had the factors 5 and 7, North American periodical cicadas would not have 13-year life cycles.

(CF₄) If, in addition to 13 and 1, 13 had the factors 19 and 23, North American periodical cicadas would not have 13-year life cycles.

For if (CF₃) and (CF₄) are true, then that would seem to cast doubt on the importance of (CF₁) and (CF₂) for this case. In order to explain what to hold fixed, however, so as to render (CF₁) and (CF₂) true, whilst rendering (CF₃) and (CF₄) false, we need to do two things. First, we need to say a bit about how to make sense of the antecedent of these conditionals: the mathematical differences at issue. Then we need to track this into the extra-mathematical case by saying a bit about the consequent.

Focus on (CF₁). The “twiddle” to 13 that we are imagining involves changing its factors. Suppose we do this by giving 13 the factors 2 and 6. It would seem that a contradiction will quickly arise. For we are supposing that $2 \times 6 = 13$. But $2 \times 6 = 12$, and $12 \neq 13$. So it is both true that $2 \times 6 = 13$ and that $2 \times 6 \neq 13$. Similar contradictions quickly infect the rest of the number-theoretic structure under consideration. If $2 \times 6 = 13$, then $1 \times 6 = 13/2$. But $1 \times 6 = 6$, and $6 \neq 13/2$. So it is both true that $1 \times 6 = 6$ and true that $1 \times 6 \neq 6$. Similarly, if $2 \times 6 = 13$, then any multiple of 13 should have 2 and 6 as factors. So 26 should have 2 and 6 as factors. But 26 doesn’t have 6 as a factor. So 26 both does and does not have 6 as a factor. And so on it goes.

That contradictions arise in the situation imagined is not necessarily a devastating problem. There are robust and precise inconsis-



tent mathematical structures, explored and discussed, for example, in (25; 40; 30; 31; 33; 35). These are based on paraconsistent logics, in which true contradictions are manageable. However, we do not see the need, in many cases, to deploy a paraconsistent logic simply because we are twiddling mathematics; we might twiddle mathematics while holding classical logic fixed.⁷ To do this, we must *avoid*, rather than *manage*, the contradictions arrived at in the previous paragraph; the way to do this is to hold less of the mathematics fixed, so that contradictions do not arise.

We should not go too far, however; we still want to hold fixed as much as we can with respect to the natural numbers. What we’re ultimately interested in, recall, are the ramifications of twiddling 13. We are not interested in the ramifications of twiddling any other number. In other words, we want to be able to carry out a “surgical strike” on 13 that enables us to gauge the consequences of altering this number for physical reality in as much isolation as possible from alterations to anything else within mathematics. Here’s our suggestion: work backwards from the desired twiddle. First, twiddle 13 and hold some portion of the number theory structure fixed. Does a contradiction result? If yes, then relax the amount you’ve held fixed and re-twiddle. Does a contradiction result? If yes, then relax the amount you’ve held fixed and re-twiddle. Does a contradiction result? If yes... And so on. Stop when you get to the maximal amount you can hold fixed within mathematics without inducing a contradiction. If there is more than one maximal amount, then pick the maximal amount that interests you, and let the interests be set by your context of evaluation.

Without going through the details, here is the result of applying this suggestion to the cicada case. One can hold all of number theory fixed except for the twiddles to 13 if one is prepared to change the way that multiplication works. For there is another function, multiplication*, that takes all of the same inputs and yields all of the same outputs as multiplication except in one special case. Whereas

7. In other settings we might well consider twiddling logical principles (7).

multiplication never takes in 2 and 6 and yields 13, multiplication* does exactly that. Moreover, whereas multiplication takes in 2 and 6 and yields 12, multiplication* does not. Below is a graphical representation of a fragment of the multiplication function and a fragment of the multiplication* function.



Multiplication* will preserve the same theorems as multiplication, and imbue the natural numbers with the same structure, except for whatever disruption is involved in changing the factors of 13; obviously we have to be willing to allow that much disruption in order to make the counterfactual supposition in the first place. Moreover, the structure will be consistent just if multiplication* does not take one set of numbers as input and map those same numbers onto two different outputs. Because functions are so easy to come by, we can be assured that there is some function that behaves exactly this way, and so no contradictions will arise by twiddling multiplication so that it matches multiplication*. This is the surgical strike that we seek: by letting 13 vary whilst demanding parity between multiplication and multiplication* in all other matters, we can successfully isolate the change to 13's

factors and consider the implications this change on its own might have (if any) beyond mathematics.

You might be concerned that in ironing out one mathematical contradiction, we simply invite more, so the ironing-out process goes on forever. Or at least it goes on until we arrive at some very non-standard arithmetic, where multiplication, addition and pretty-much everything is changed.⁸ But we do not need to resolve all the looming contradictions. We just need to resolve those relevant to assessing the counterfactual at hand. After all, that's all we do in the Suzy and the rock case. We do not go all the way back to the big bang or even to Suzie's birth in order to achieve consistency. We iron out the immediate inconsistencies and leave it there. But somewhere in the background there will be further inconsistencies looming. Suzy moved her arm in a throwing motion, yet the rock did not move? She willed her arm to move, but it didn't? We simply set these problems aside because they are not relevant to the assessment of the counterfactual of interest. We entertain her failing to throw the rock, and we don't much care what went on immediately before this. There is nothing different in the mathematical case. We deal with any immediate contradictions and leave it there.⁹ The contradictions might go on forever, but we do not need to deal with them all: we simply push the lump in the carpet away from the area we care about, and that's the end of it.

2.3 *Beyond Mathematics*

So that is our basic procedure for handling the antecedent of a counterfactual implicated in a case of extra-mathematical explanation: hold

as much fixed as you can within mathematics compatible with the twiddle, without inducing a contradiction. The next step is to connect this up to the consequent, and to thus track counterfactual changes to mathematics into the physical realm. Again, we're after an account of what to hold fixed. The first thing to notice is that we're not dealing with just one structure. We're dealing with two structures: a mathematical structure (such as the natural numbers) and a physical structure (the physical world). Accordingly, in order for changes within mathematics to ramify into the physical world, there must be some link between the mathematical and physical structures, such that a change to the mathematical structure implies a corresponding change in the physical structure. This link must be held fixed as part of the counterfactual supposition in order for the mathematical twiddle to properly ramify.

The link that we propose to hold fixed is a morphism. The broad idea is that mathematical systems and physical systems share particular structural features. In the extreme case, a mathematical structure may be isomorphic to some physical structure. But, as Bueno and Colyvan (11) point out, the mapping relation is rarely that tight: often some looser mapping will be implicated instead, such as a homomorphism or a monomorphism. The details, for now at least, are not important. What matters is this: there is always a structural parity between mathematical and physical structures to some degree. This matters because holding fixed the fact that the mathematical structure maps the physical structure is enough to allow twiddles to the mathematics to ramify across to the physical system at issue.¹⁰ When we hold fixed the fact

8. Thanks to Rachael Briggs and a referee of this journal for pressing us on this issue.

9. In a way, our example of the contradiction generated by multiplication (once we suppose that 13 is composite) is a little misleading. After all, it is not clear that any of this is relevant to the counterpossible of interest, so there's no need to even consider the contradictions arising there. We presented the details of how we might iron out such a contradiction by way of example. But, in fact, there is no reason to think that such further contradictions are relevant.

10. An alternative to the mapping account would be to treat the relation between mathematical and physical facts as a logical relation. So, for instance, it may be that mathematical facts can be used to deduce various empirical facts. More generally, then, entailment of some order may be used to carry traffic out of mathematics. Alan Baker (3; 5) toys with a view along these lines). Again, exactly which consequence relation is needed is a matter for debate.

A third option, recently suggested by Pincock (29), is to link mathematical and empirical facts using constitution relations, the idea being that phys-

that there is a morphism between a mathematical and a physical structure, what we're effectively demanding is that changes to those parts of a mathematical structure mapped onto a physical structure must be reflected in the physical structure at issue. The physical structure must, as it were, keep up with any twisting or bending in the mathematical structure. Of course, the degree to which the physical structure keeps step with the mathematical structure depends a great deal on the strength of the morphic relationship. The stronger the morphism — the more structure it preserves — the more the physical structure must bend in response to changes within the mathematical structure.

One final thing: in order to carry the changes from mathematics through to the physical world, we must make a decision about what to hold fixed with respect to physical reality. As is standard, we give the laws priority, and recommend holding fixed as much as possible

ical systems are 'constituted by' or otherwise depend ontologically upon mathematical facts. Such a view lends itself toward a "Heavy Duty Platonism" which Field (16, pp. 186–193) contrasts with "moderate platonism":

According to moderate platonism [...] *relations between physical things and numbers are conventional relations that are derivative from more basic relations that hold among physical things alone.* The heavy duty platonist rejects this, taking the relation between physical things and numbers to be a brute fact, not explainable in other terms. (If one likes to flaunt one's heavy duty metaphysics, one can say that there is a mysterious relation of platonic participation between physical things and numbers. But the position is the same whether or not one flaunts it.) [emphasis in the original]

These three accounts of the relation between mathematical and empirical facts are by no means exhaustive or, indeed, exclusive: it may be that the link between the mathematical and the physical is constituted by some combination of the above relations between mathematical and empirical facts, or another relation entirely. (Further options include identity [mathematical and empirical facts are identical with one another] or lawful dependence [mathematical and empirical facts are related by physical laws]. One might also be tempted to take the link itself to be a relation of counterfactual dependence. That option, however, threatens to be circular in the current context.) But there must be *some* link if we are to make sense of counterfactuals in the context of extra-mathematical explanation at all. For if there is no such link, then twiddle mathematics all you like — this will never have implications for anything empirical.

compatible with those laws, whilst allowing changes to the physical world to ramify appropriately. This, we take it, typically means holding fixed the laws and the past up to a particular point in time, and then allowing counterfactual changes to ramify into the future.

2.4 *An Evaluation Procedure*

We are now in a position to outline our general strategy for evaluating the counterfactuals implicated in cases of extra-mathematical explanation. First, hold fixed the morphism between the mathematical structure S that appears in the counterfactual and the physical structure P . Second, make a change to mathematics while holding fixed as much as one can without inducing a contradiction. Finally, consider the ramifications of the change by looking at the way(s) in which the physical structure P twists in response to the twiddling in S in order to preserve the morphism.

To see the procedure in action, consider again the following four counterfactuals:

(CF₁) If, in addition to 13 and 1, 13 had the factors 2 and 6, North American periodical cicadas would not have 13-year life cycles.

(CF₂) If, in addition to 17 and 1, 17 had the factors 2 and 6, North American periodical cicadas would not have 17-year life cycles.

(CF₃) If, in addition to 4 and 3, 12 had the factors 5 and 7, North American periodical cicadas would not have 13-year life cycles.

(CF₄) If, in addition to 13 and 1, 13 had the factors 19 and 23, North American periodical cicadas would not have 13-year life cycles.

As noted, we want (CF₁) and (CF₂) to be true and (CF₃) and (CF₄) to be false. To evaluate these counterfactuals, we must first identify the mathematical and physical structures at issue. The mathematical struc-

ture is the structure of the natural numbers. The physical structure is time measured in years. The explanation is underwritten by the representation of years via the natural numbers. That, then, is the morphism we must hold fixed. Now, consider (CF_1) . To evaluate this counterfactual, we start in the mathematics. We hold fixed as much as we can by changing multiplication to behave like multiplication*. This leaves 13's factors as desired. This gives us a structure, S^* , that is just like the natural numbers, except that 13 is not prime, and factorises via 2 and 6. Because we are holding fixed the relationship between the mathematical and physical structures, the physical structure that is now being mapped onto S^* must twist to keep up with the counterfactual change. The result is that an interval of 13 years is now divisible into six two-year segments, or into two six-year segments. It follows from this that a cicada with a 13-year life cycle will overlap with predators that have two-year and six-year life cycles and thus that 13 is not an optimal way to avoid predation. So cicadas won't evolve 13-year life cycles. So (CF_1) is true. Exactly the same story can be told for (CF_2) , *mutatis mutandis*.

This brings us to (CF_3) . Again, we start by holding the morphism fixed. We then change multiplication's behaviour so that 12 takes on the additional factors 5 and 7. This change is then carried down into the physical structure of years, forcing the division of years to keep step. The result is that 12-year intervals become divisible into 5 and 7 year segments. Though the morphism forces the world to keep up with the mathematical twiddle, nothing changes for the cicadas. For it remains the case that 13 and 17 are the optimal strategies for avoiding predation. All that has happened is that evolving a 12-year life cycle is now an even worse strategy for avoiding predation than it was before, since a 12-year cicada will overlap with even more different kinds of predators. So (CF_3) is false.

Finally, consider (CF_4) . Hold fixed the morphism. Now make the counterfactual change to the mathematics. The world keeps up its end of the bargain, and so a 13-year lifespan is now divisible into 19- and 23-year intervals. The cicadas don't budge: 13 remains the optimal strategy for avoiding predation by organisms that have life cycles up to

18 years. Of course, if there are 19- or 23-year predators, then 13 is no longer optimal. However, there are ecological constraints on the cicada case that rule out these predators. So (CF_4) is false.

We anticipate two responses to our proposed evaluation procedure. The first response relates to the idea that we should hold fixed the morphic relationship between the natural numbers and the length of years. The trouble comes this way: the natural numbers stand in morphic relationships to a large number of physical structures. If all of these morphisms are held fixed, then changes to mathematical facts, when they ramify, might well result in massive changes to the universe. Indeed, the changes could be so great that (CF_1) and (CF_2) are falsified. Here's one example of what we mean: Suppose that we make a change to the natural numbers. Suppose, however, that this results in a change to one of the universal constants, which is also represented by some natural number. For instance, suppose that it results in the second power having the properties of the third power, turning the inverse square law into something that behaves as an inverse cube law. Suppose further that this change is needed to evaluate a counterfactual like (CF_1) . The results will be rather bad: if the inverse square law behaves like an inverse cube law, then galaxies won't form. The upshot is that things like cicadas won't exist, and so no interesting facts about their life cycles will follow.

Once again, however, this is to hold too much fixed. By holding all of the morphic relationships between the natural numbers and physical structures fixed, we end up demanding rather global changes to physical reality in order to force it to keep step with changes to mathematics. So we cannot hold so much fixed. In order to get the right truth-values for the counterfactuals above, we must hold fixed only the morphic relationship between the natural numbers and the life-cycle lengths of cicadas in years. We must allow all other morphic relationships to break upon twiddling the mathematics, so that the changes implied by the twiddle don't massively rewrite physical reality.

The second worry follows immediately from this last point: what justifies holding fixed the morphic relationship between mathemati-

cal and physical structures at all? Providing such a justification might seem especially difficult given that we have just granted that one cannot hold all morphisms fixed; one must hold fixed only the morphisms directly related to the counterfactual supposition at issue. It is not our goal, however, to provide such a justification. Recall, our goal is not to argue that these counterfactuals are in fact true. That is a job for the proponent of a counterfactual theory of extra-mathematical explanation (and perhaps for our future time-slices). Rather, we are operating under the assumption that the counterfactuals *are* true, and then outlining what must be held fixed in order to yield that result so as to delineate a general evaluation procedure. To ask whether it is reasonable to hold these facts fixed when evaluating counterfactuals is to call into doubt the truth of the counterfactuals at issue. And so we pass the buck.

Nonetheless, what we have said here is by no means idle. The proponent of a counterfactual theory of extra-mathematical explanation is now in a much better position to know what they must defend to get their theory off the ground. They must argue that the strategy for evaluating counterfactuals sketched here, or something very much like it, is the correct way to evaluate the counterfactuals implicated in cases of extra-mathematical explanation, since that would put them in a strong position to hold that the counterfactuals are true.

3. Structural Equation Modelling

We turn now to the integration of extra-mathematical counterfactuals into the broader framework of structural equation modelling. The model we will use to represent the dependencies in this case is in Figure 3. To keep things simple, each node in the diagram is a full proposition; each can be true or false. The model of the dependencies in question we are using is also extremely simplified, but this is okay; the point here is not to give a particularly complete or accurate picture of the details of the cicada case itself, but instead to give an example of how the kinds of structural equation tools deployed in ordinary counterfactual cases can be used in mathematical counterfactual cases as

well.

The letters in the diagram should be understood as follows (we use $a \perp b$ to mean that a and b are *coprime*: that their greatest common denominator is 1):

- | | |
|---|---|
| A | If $a \perp b$, the distance between their common multiples is $a \times b$. |
| B | If $a \not\perp b$, the distance between their common multiples is $< a \times b$. |
| C | Coprime numbers have larger distances between their common multiples than nearby non-coprimes. |
| D | When animals reproduce with cycles x and y , they reproduce simultaneously at common multiples of x and y . |
| E | Animals with coprime life cycles reproduce simultaneously less frequently than those with nearby non-coprime life cycles. |
| F | Cicadas will evolve life cycles coprime with those of their predators. |
| G | Cicadas will evolve life cycles that intersect with those of their predators the least frequently. |
| H | Environmental constraints fix cicada life cycles to be between 12 and 18 years. |

I	There is a cicada predator with life cycle 2.
J	There is a cicada predator with life cycle 3.
K	$12 \not\perp 2$ and $12 \not\perp 3$
L	$14 \not\perp 2$
M	$15 \not\perp 3$
N	$16 \not\perp 2$
O	None of 12, 14, 15, 16 are coprime with both 2 and 3.
P	13 is prime.
Q	17 is prime.
R	If a number is prime, it is coprime with every other number.
S	13 and 17 are each coprime with both 2 and 3.
T	Cicadas will evolve life cycles of 13 and 17.

Each node in the diagram takes value 1 or 0, according to whether the proposition it represents is true or false, respectively. We suppose that all the exogenous nodes take value 1. The structural equations we use are simple: each endogenous node takes the minimum of the values of the nodes that feed into it. So, for example, $C = \min(A, B)$ and $O = \min(K, L, M, N)$. In some cases, these equations are more realistic than in others; again, our goal is to sketch the structure of an approach like this, rather than to defend a particular account of the cicadas themselves.

Now, let's consider some twiddles. Suppose, for example, that we twiddle the handle at *I*, removing the cicada predator with life cycle 2. Then this change would ramify to *T*. If the cicadas didn't have a predator with life cycle 2, then their own life cycles wouldn't be 13 and 17. For example, there might be a species with life cycle 14, at least if there is no predator with life cycle 7 either.

Or suppose we twiddle the handle at *S*. Since this is an *internal* handle, twiddling it breaks the upstream links; the equation $S = \min(P, Q, R)$ that is part of the specification of the model no longer needs to be respected. But the downstream effects of this twiddle will still be felt; the change again ramifies to *T*. If 13 and 17 weren't each




Figure 1: The cicadas

coprime with both 2 and 3, then the cicadas' life cycles would not be 13 and 17.

The important point, for our purposes, is how *alike* those two twiddles are. The main difference between these cases, from the point of view of the model, is that I is exogenous while S is endogenous. The fact that I is non-mathematical and S mathematical simply does not affect how they fit into this model, or how we handle them via the model. The structural equation modelling toolkit is completely insensitive to the difference between mathematical and non-mathematical cases; as a result, it works just as well on mathematical cases as it does on non-mathematical cases. We have developed the model with more nodes than just these; each node, as usual, provides an opportunity for twiddling. Some nodes are purely mathematical (e.g. Q), some purely non-mathematical (e.g. J), and some mixed (e.g. F).

This model emphasizes some of the details of the mathematical part of the cicada case, while seriously simplifying the non-mathematical bits. We don't mean to suggest that the non-mathematical aspects of this case (environmental constraints, predator behaviour, evolution, etc.) are unimportant. We simply want to make clear that the mathematics involved in this case is not monolithic. There are any number of different twiddles that can be made even within the mathematics of this case, and they will ramify differently throughout the system we've sketched.

3.0.1 *The Honeybees*

Once we see how to model the cicada case so as to make twiddles in the mathematical and non-mathematical components, it is a straightforward matter to model other cases of extra-mathematical explanation in just the same way. We turn to another familiar case: the honeybees, as outlined in (22).

Honeybees produce hexagonal honeycomb cells. Why? The optimality explanation for this phenomenon contains (i) ecological constraints on honeybees, namely: pressure to be as efficient as possible in

producing and storing nectar; (ii) information concerning the trade-off between energy consumption and wax production; and (iii) mathematical results concerning the most efficient way to tessellate a 2D surface into regions of greatest area with least total perimeter. The model tells us that the optimal honeycomb shape is hexagonal. That's because the use of hexagons is the most efficient way to tessellate a surface area into cells capable of storing the greatest amount of nectar. Accordingly, hive-bees have evolved to exploit this optimal trade-off between energy and efficiency with regard to nectar storage and production.

To test the mathematics in this optimality model, we can test the geometric results of the honeycomb theorem which proves the optimality of hexagons for tessellating a 2D surface area (17). To do that, we consider the following counterpossible:

(CF₃) If the optimal way to tessellate a surface into regions of greatest area with least total perimeter had not been via hexagons, hive-bees would not have built hexagonal honeycomb cells.

CF₃ is a coarse-grained version of something more fine-grained, such as manipulating the perimeter-to-surface-area function, or directly manipulating the properties of hexagons by altering the interior angles of hexagonal polygons.

It is straightforward to develop a model of the honeybee case that is as detailed as the cicada model in Figure 1. Instead of doing that, however, we will demonstrate a much more compact model. Rather than dividing up the various mathematical and empirical facts that may be twiddled and spreading them across nodes, encoding each using equations, we will simplify matters by treating the mathematical facts and the empirical facts as two binary nodes that feed into the explanandum of interest. The model is displayed in Figure 2. To produce the honeybee case, use the following translation schema:




Figure 2: Honeybees

A	Hexagons are the most efficient method of tessellating a surface into regions of equal area with least total perimeter.
B	Hive-bees are under evolutionary pressure to produce the largest honeycomb cells using the least wax.
C	Hive-bees will evolve to produce hexagonal honeycomb.

And use the following structural equations: $A = 1$; $B = 1$; $C = \min(A, B)$.

By simplifying the honeybee case in this manner, we effectively treat any mathematical twiddle on a par with any other mathematical twiddle with respect to C . So too for the non-mathematical twiddles. This can be useful if we want to look at the explanations in a very coarse-grained way, so as to perhaps compare the mathematical and non-mathematical twiddles with one another. Mathematical twiddling happens as in the cicada case. Suppose, in the honeybee case, we set A to 0 by making it the case that circles, and not hexagons, are the most efficient way to divide a surface into regions of greatest area with least total perimeter. Then, by the structural equation for C , it no longer follows that the honeybees will evolve to produce hexagonal honeycomb.

3.0.2 Generalized Optimality Modelling

One of the advantages of using a counterfactual theory of explanation to handle standard non-mathematical cases of explanation is that it can facilitate the process of honing in on those things that are really making a difference to whatever it is we are trying to explain. The same is true in cases of extra-mathematical explanation.

Both the honeybee case and the cicada cases are instances of a classic mathematical problem with a wide variety of applications. This is the problem of optimising some quantity subject to a specified constraint. For example, in conservation biology we might be interested in minimising the chance of extinction of a particular threatened species, subject to a fixed budget constraint. In physics we might be interested in maximising the electric potential, subject to a spatial constraint. In economics we might be interested in maximising a share portfolio return, subject to an investment constraint. By looking at the mathematical features of optimality modelling as a general phenomenon, it is possible to identify mathematical facts that, were they to change, would have substantial ramifications through all cases of extra-mathematical explanation that deploy optimality modelling, of which there are many (6; 34).

Here we give a very simple example of such a problem in order to get a feel for some of the relevant mathematics involved. Consider the problem of maximising the area of a rectangular paddock subject to the constraint of a fixed and specified length of fencing material. Let the sides of the optimum paddock be x and y respectively, and let c denote the length of fencing material available. The problem is thus one of finding x and y such that xy is a maximum subject to the constraint that $2x + 2y = c$. One very powerful general method for solving problems such as this is the method of Lagrange multipliers.¹¹ We solve the original optimisation problem by introducing a third variable, λ , and an auxiliary function of three variables (our original variables x and y

¹¹. See, for example, (1).

along with the new variable λ). This auxiliary function is known as *the Lagrange function*:

$$\Lambda(x, y, \lambda) = xy + \lambda(2x + 2y - c). \quad (1)$$

We then find the stationary points of Λ . The first partial derivatives here are $\partial\Lambda/\partial x = y + 2\lambda$, $\partial\Lambda/\partial y = x + 2\lambda$, $\partial\Lambda/\partial\lambda = 2x + 2y - c$. Setting these to zero yields: $\lambda = -c/8$, $x = y = c/4$. This gives us an area for the paddock in question of $c^2/16$.¹²

A few points are worth noting in relation to this method of solving optimisation problems under constraint. First, in finding a solution to our problem, we have also proven that the rectangle in question is a square (i.e. $x = y$). Indeed, this is so, irrespective of the amount of fencing material available. This is a very general result that drops out of the mathematics in question. Next, note that we have here a general solution to a whole class of problems. The mathematics doesn't care whether we're interested in areas of paddocks, electrical potentials, or profit margins. All problems that can be formulated as above are solved in one fell swoop along with our paddock problem.

Finally, we note that the method itself is very general in a couple of senses. First, it applies to any function $f(x, y)$ of two variables to be maximised subject to any constraint function $g(x, y) = c$ of two variables, so long as f and g have continuous first partial derivatives. In every case, we form the Lagrangian $\Lambda(x, y, \lambda) = f(x, y) + \lambda(g(x, y) -$

$c)$ and find the latter's stationary points.¹³ Moreover the method is general in the sense that it can straightforwardly be adapted to higher dimensions: to maximising and minimising functions of n variables subject to a constraint function of n variables. These points about the generality of the method are important for what follows; not only can we simultaneously solve optimisation problems from different areas of science, but we can solve *non-isomorphic* problems (i.e. ones with different constraint functions and different functions to be optimised) by the same means. Moreover, the reason the method works is itself mathematical.¹⁴ The solution above, for example, does not depend on the causal properties of fences, electrical fields, markets, or the like.

The mathematics underlying the optimality results that drive the cicada case and the honeybee case are much the same as the mathematics just described. In each case we are maximising some quantity, while holding fixed (or, in slightly more complicated cases, minimising) some other quantity. We might be maximising paddock area, honey storage capacity, or years between potential overlap with predators, and trying to achieve this while holding fixed the length of fencing, the wax available to build hive cells, or ecological resources, respectively. The mathematics is blind to such matters. We can thus represent these cases in more detail, if needed: we can add a node through which the mathematical and physical facts get filtered. This is easiest to represent in the honeybee case. The case is modelled in

12. It's not too hard to see why this method works. The basic idea is that any maximum of f occurs when the gradient of f lines up with the gradient of g and thus one gradient is a multiple of the other. The Lagrange multiplier λ is just the multiplicative constant in question.

Of course, in this particular example we could solve the problem using more elementary methods. We could, for instance, use the constraint function to produce an expression for y in terms of x and substitute this into the function to be maximised. We could then find the value of x at the maximum and use this in the constraint function to find the corresponding value for y . But we can't always do this; the method of Lagrange multipliers is more general.

13. In fact the method of Lagrange multipliers provides only a necessary condition for optimisation — not all stationary points of the Lagrange function are optima. Necessary and sufficient conditions for the solution to the optimisation problem exist, in terms of determinants of the relevant bordered Hessian matrix of second partial derivatives of the Lagrangian, but we can set this complication aside. This simplified partial solution is enough to serve our present purposes.

14. C.f. the virtues of the general method of solving homogeneous linear differential equations discussed in (14).




Figure 3: Generalized Optimality

Figure 3. The translation schema is this:

- | | |
|---|---|
| A | Hexagons are the most efficient method of tessellating a surface into regions of equal area with least total perimeter. |
| B | Hive-bees are under evolutionary pressure to produce the largest honeycomb cells using the least wax. |
| C | Hexagons are the optimal method for producing honeycomb cells. |
| D | Hive-bees will evolve to produce hexagonal honeycomb cells. |

And the equations are as follows: $A=1$, $B=1$, $C = \min(A, B)$, $D = C$. A pared-down version of the cicada case that includes an optimality node can be modelled using Figure 3 by switching to the following alternative translation schema:

- | | |
|---|--|
| A | Coprime numbers have the biggest distances between their common multiples. |
| B | Cicadas will evolve life cycles that intersect with those of their predators the least frequently. |
| C | Prime life cycles are the optimal strategy for avoiding predation. |
| D | Cicadas will evolve prime-numbered life cycles. |

By adding the C node, we now have the ability to twiddle broad facts about optimality. One way to do this is to change the way that optimality calculations work by twiddling the Lagrangian mathematical basis. So, for instance, we might twiddle the Lagrange function by twiddling the constraint function. If we do this, then all the relevant optimality models will yield different results. What we will see, and this is the important point, is that changing the mathematics in this way is to twiddle the C node in both the honeybee graph and the cicada graph at the same time and in the same way and with the same upshot: namely, $D = 0$. The mathematics of optimisation then unifies both of these explanations (and many more besides), whilst also preserving the counterfactual dependence of empirical results on mathematical facts (namely, facts about Lagrange multipliers). What we can see, then, is that the mathematical results about primes or about hexagons are, in some sense, secondary to the mathematics underlying the optimality process used to recover the explanandum in each case. The more general interest of this is that we are able to find this “master handle” in much the same way that such handles are found in ordinary cases of causal modelling, further demonstrating the common refrain of this paper: whatever you want to say about counterfactuals in ordinary cases, you can say the very same sorts of things in an extra-mathematical context. There really is no reason for concern about the generalisation of counterfactual theories of explanation over to extra-mathematical explanations.

4. Concluding Remarks

In this paper, we've shown that the familiar tools of structural equation models and counterfactual dependence are flexible enough to apply to extra-mathematical explanations in much the same way as they apply to causal explanations.

There is a lot, of course, that we haven't done. We haven't defended these familiar tools as an account of explanation at all. For that, see (28; 42; 43). We also haven't given a semantics that can handle the counterfactuals we've appealed to — and since these counterfactuals have impossible antecedents, the accounts of (20; 36) won't work. Accounts that *will* work, however, are not hard to come by; see (8; 9; 10; 18; 21; 23; 24; 27; 32; 38; 39).

Finally, we haven't explored the possibilities for making good on our opening motivation: to find a single account of explanation that can help us understand all kinds of explanation, from an explanatory proof of a mathematical theorem all the way to explanations of billiard ball velocities. But this is as it should be, for now; unified theories get built up piece by piece, not all at once. We trust that this exploration of extra-mathematical explanation turns out to have been a step in the right direction.

Acknowledgments

Thanks to Franz Berto, Rachael Briggs, Rohan French, Alan Hájek, Ed Mares, Daniel Nolan, and Graham Priest for helpful discussions about counterpossibles and Joy Britten for pointing out an error in an earlier draft. Versions of this paper were presented at the workshop "Determinism, Probability, and Conditionals" at the Australian National University in July 2015 and at the workshop "Explananza" at the University of Western Australia in December 2015. We are indebted to the audiences at these presentations for very fruitful discussions and for pressing us on several issues. Mark Colyvan's work on this paper was funded by an Australian Research Council Future Fellowship (grant number FT110100909).

References

- [1] Adams, Robert A. *Calculus of Several Variables*. Don Mills, Ontario, Canada: Addison-Wesley, 1987
- [2] Baker, Alan "Does the Existence of Mathematical Objects Make a Difference?" *Australasian Journal of Philosophy* 81/2 (2003): 246–264
- [3] ——— "Are There Genuine Mathematical Explanations of Physical Phenomena?" *Mind* 114/454 (2005): 223–238
- [4] ——— "Mathematical Accidents and the End of Explanation" In Otávio Bueno and Øystein Linnebo (eds.), *New Waves in Philosophy of Mathematics*. Basingstoke, UK: Palgrave Macmillan, 2009: 137–159
- [5] ——— "Mathematical Explanation in Science" *British Journal for the Philosophy of Science* 60/3 (2009): 611–633
- [6] Baron, Sam "Optimisation and Mathematical Explanation: Doing the Lévy Walk" *Synthese* 191/3 (2014): 459–479
- [7] Baron, Sam and Mark Colyvan "Time Enough for Explanation" *Journal of Philosophy* 113/2 (2016): 61–88
- [8] Beall, Jc et al. "On the Ternary Relation and Conditionality" *Journal of Philosophical Logic* 41/3 (2012): 595–612
- [9] Bernstein, Sara "Omission Impossible" *Philosophical Studies* 173/10 (2016): 2575–2589
- [10] Brogaard, Berit and Joe Salerno "Remarks on Counterpossibles" *Synthese* 190/4 (2013): 639–660
- [11] Bueno, Otávio and Mark Colyvan "An Inferential Conception of the Application of Mathematics" *Noûs* 45/2 (2011): 345–374
- [12] Chalmers, David J. "Does Conceivability Entail Possibility?" In Tamar S. Gendler and John Hawthorn (eds.), *Conceivability and Possibility*. Oxford: Oxford University Press, 2002: 145–200
- [13] Colyvan, Mark *The Indispensability of Mathematics*. Oxford: Oxford University Press, 2001
- [14] ——— "Mathematics and Aesthetic Considerations in Science" *Mind* 111/441 (2002): 69–74
- [15] ——— "There Is No Easy Road to Nominalism" *Mind* 119/474 (2010): 285–306

- [16] Field, Hartry *Realism, Mathematics and Modality*. Oxford: Basil-Blackwell, 1989
- [17] Hales, Thomas C. "The Honeycomb Conjecture" *Discrete and Computational Geometry* 25/1 (2001): 1–22
- [18] Jago, Mark "Impossible Worlds" *Noûs* 47/3 (2013): 713–728
- [19] Kripke, Saul A. *Naming and Necessity*. Cambridge, Massachusetts: Harvard University Press, 1980
- [20] Lewis, David *Counterfactuals*. Oxford: Blackwell, 1973
- [21] Lycan, William G. *Real Conditionals*. Oxford: Oxford University Press, 2001
- [22] Lyon, Aidan and Mark Colyvan "The Explanatory Power of Phase Spaces" *Philosophia Mathematica* 16/2 (2008): 227–243
- [23] Mares, Edwin D. "Who's Afraid of Impossible Worlds?" *Notre Dame Journal of Formal Logic* 38/4 (1997): 516–526
- [24] Mares, Edwin D. and André Fuhrmann "A Relevant Theory of Conditionals" *Journal of Philosophical Logic* 24/6 (1995): 645–665
- [25] Mortensen, Chris *Inconsistent Mathematics*. Dordrecht: Kluwer Academic Publishers, 1995
- [26] Nickel, Bernhard "How General Do Theories of Explanation Need to Be?" *Noûs* 44/2 (2010): 305–328
- [27] Nolan, Daniel "Impossible Worlds: A Modest Approach" *Notre Dame Journal of Formal Logic* 38/4 (1997): 535–572
- [28] Pearl, Judea *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press, 2000
- [29] Pincock, Christopher "How to Avoid Inconsistent Idealizations" *Synthese* 191/13 (2014): 2957–2972
- [30] Priest, Graham "Inconsistent Models of Arithmetic Part I: Finite Models" *Journal of Philosophical Logic* 26/2 (1997): 223–235
- [31] ——— "Inconsistent Models of Arithmetic Part II: The General Case" *The Journal of Symbolic Logic* 65/4 (2000): 1519–1529
- [32] ——— *Beyond the Limits of Thought*. Oxford: Oxford University Press, 2002
- [33] ——— "Mathematical Pluralism" *Logic Journal of the IGPL* 21/1 (2013): 4–13
- [34] Rice, Collin "Moving Beyond Causes: Optimality Models and Scientific Explanation" *Noûs* 49/3 (2015): 589–615
- [35] Shapiro, Stewart *Varieties of Logic*. Oxford: Oxford University Press, 2014
- [36] Stalnaker, Robert C. "A Theory of Conditionals" In Nicholas Rescher (ed.), *Studies in Logical Theory*. Oxford: Blackwell, 1968: 98–112
- [37] Steiner, Mark "Mathematics, Explanation, and Scientific Knowledge" *Noûs* 12/1 (1978): 17–28
- [38] Vander Laan, David "The Ontology of Impossible Worlds" *Notre Dame Journal of Formal Logic* 38/4 (1997): 597–620
- [39] ——— "Counterpossibles and Similarity" In Frank Jackson and Graham Priest (eds.), *Lewisian Themes: The Philosophy of David K. Lewis*. Oxford: Oxford University Press, 2004: 258–276
- [40] Weber, Zach "Transfinite Cardinals in Paraconsistent Set Theory" *The Review of Symbolic Logic* 5/2 (2012): 269–293
- [41] Williamson, Timothy *The Philosophy of Philosophy*. Oxford: Oxford University Press, 2007
- [42] Woodward, James *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press, 2003
- [43] Woodward, James and Christopher Hitchcock "Explanatory Generalizations, Part 1: A Counterfactual Account" *Noûs* 37/1 (2003): 1–24
- [44] Yablo, Stephen "Is Conceivability a Guide to Possibility?" *Philosophy and Phenomenological Research* 53/1 (1993): 1–42