

A Formal Apology for Metaphysics

Sam Baron*

Abstract

There is an old meta-philosophical worry: very roughly, metaphysical theories have no observational consequences and so the study of metaphysics has no value. The worry has been around in some form since the rise of logical positivism in the early twentieth century but has seen a bit of a renaissance recently. In this paper, I provide an apology for metaphysics in the face of this kind of concern. The core of the argument is this: pure mathematics detaches from science in much the same manner as metaphysics and yet it is valuable nonetheless. The source of value enjoyed by pure mathematics extends to metaphysics as well. Accordingly, if one denies that metaphysics has value, then one is forced to deny that pure mathematics has value. The argument places an added burden on the sceptic of metaphysics. If one truly believes that metaphysics is worthless (as some philosophers do), then one must give up on pure mathematics as well.

1. INTRODUCTION

There is an old meta-philosophical worry: very roughly, metaphysical theories have no observational consequences and so the study of metaphysics has no value. The worry has been around in some form since the rise of logical positivism in the early twentieth century but has seen a bit of a renaissance recently, with new versions of it appearing in the work of Maclaurin and Dyke (2012) and Ladyman and Ross (2007).¹ My goal is to offer an apology for metaphysics in the face of this type of concern. Metaphysics, I will argue, is valuable for the same reasons that pure mathematics is valuable. To be clear: my claim is not that metaphysics is just like pure mathematics, and so because pure mathematics is a valuable activity, so too for metaphysics. My claim, rather, is that the reasons for valuing pure mathematics generalise to metaphysics as well, despite the activities bearing important dissimilarities to one another. The paper is structured as follows. In Section 2 I sketch the shape of the apology before presenting the apology in detail in Section Three. After that I consider some objections in Section Four and wrap up in Section 5.

2. THE APOLOGY

My apology focuses on metaphysical theories that lack observational consequences. A theory lacks observational consequences when there is, in principle, no way to con-

*Department of Philosophy, University of Western Australia. *Email*: samuel.baron@uwa.edu.au

¹Ney (2012) and Price (2009) also offer versions of the worry, but with a Carnapian twist.

firm or disconfirm that theory via observation. Exactly what a metaphysical theory is, however, is not a question that I intend to answer in the first instance (though I will say more on this in §2). Instead, I will proceed initially by pointing to the theories that philosophers tend to study under the auspices of metaphysics, and that they therefore consider to be metaphysical theories. Examples of theories that the apology is supposed to cover thus include theories of composition, truthmaking, persistence, personal identity, grounding and causation (to name a few), all of which are considered to be metaphysical theories and all of which lack observational consequences.

Philosophers have long worried about the study of theories that, in principle, lack observational consequences. The first attempt to make this worry precise, which used the twin cannons of logical positivism and verificationism, met with a sticky end. Recently, however, a new version of the worry has arisen, one that does not require the same positivistic baggage. The revamped objection starts from naturalism, rather than positivism. Roughly stated, the objection is this. According to naturalism, the only source of knowledge of the actual world is science. But science produces knowledge by gathering observations that confirm or disconfirm theories. So if metaphysical theories lack observational consequences, then there is no way to produce knowledge of the actual world by gathering observations that confirm or disconfirm those theories. It follows that metaphysical theories play no role in the production of knowledge regarding the actual world. But metaphysics, which is the study of such theories, is valuable only if it produces knowledge of the actual world. So metaphysics is not valuable.²

Maclaurin and Dyke put forward a version of this argument. They maintain that because metaphysics “makes no difference to scientific investigations, it cannot claim, as it does, to be part of the pursuit of knowledge about the objective world” (Maclaurin and Dyke, 2012, 299). By ‘makes no difference’ they mean that it “cannot be harnessed for practical effect” (Maclaurin and Dyke, 2012, 299), which is understood in terms of observational consequences: only theories that, in principle, have observational consequences can be harnessed for practical effect.

Ladyman and Ross (2007) also offer a version of the argument. The argument is based on their *Principle of Naturalistic Closure* (PNC). According to PNC, “no hypothesis that the approximately consensual current scientific picture declares to be beyond our capacity to investigate should be taken seriously” (Ladyman and Ross, 2007, 29). Where a metaphysical theory is within our capacity to investigate only if it has observational consequences. This becomes clear when Ladyman and Ross understand the notion of ‘capacity’ in terms of extracting information from regions of spacetime using instruments:

... let us be clear that ‘capacity’ is to be read in a strong modal sense. In saying that something is beyond our capacity to investigate we do not just mean that it’s beyond our practical capacity—because we would have to last too long as a species, or travel too far or too fast or use a probe no one now has any idea how to build. We refer instead to parts of reality from which science itself tells us information cannot, in principle, be extracted

²Note that by ‘knowledge of the actual world’ I mean non-modal knowledge of the merely actual world. This will become important later on, when I argue that there is modal knowledge of the way the actual world could be that can be gained by doing metaphysics.

for receipt in our region of spacetime or in regions of spacetime to which we or our instruments could in principle go. (2007, 29)

In the end, Ladyman and Ross (2007, 29–33) maintain that the only metaphysical theories that we should take seriously are ones that contribute to progress in science in some way, and thus to the accumulation of scientific knowledge. Their ban on metaphysical theories thus follows from an underlying commitment to naturalism (which is embedded into their PNC).³

There are, broadly, four responses to the naturalistic objection against metaphysics available. First, one might deny the underlying assumption of naturalism, and maintain that there are routes to gaining knowledge of our universe that don't proceed via science. Second, one might argue that the metaphysical theories at issue do, in fact, have observational consequences. Third, one might deny that metaphysics is valuable only if it provides knowledge of the actual world. Fourth, one might concede the objection and argue for reformation: metaphysics of the relevant kind should be discontinued and replaced with a more upstanding enterprise.

Like most metaphysicians these days, I accept the underlying naturalism. I also believe that metaphysics lacks observational consequences and that this is a defeasible indication of the detachment of metaphysics from science. Finally, I don't think metaphysicians should capitulate to the objection by consigning metaphysics to the flames. This rules out the first, second and fourth responses. My goal, then, is to develop the third response by arguing that metaphysics is valuable even if it fails to produce knowledge of the actual world.

There are many potential sources of value for metaphysics. But not just any source of value will be dialectically effective when defending metaphysics. If, for instance, one claims that metaphysics is valuable because philosophers enjoy doing it, then that is unlikely to convince the critic of metaphysics that it should be continued. Instead, what needs to be shown is that metaphysics enjoys a kind of value that even the critic of metaphysics has reason to take seriously. It is for this purpose that I will look toward pure mathematics. Pure mathematics detaches from science in the same way that metaphysics does. Pure mathematical theories lack observational consequences. And yet even the most staunch, naturalistically-minded critic of metaphysics should agree that pure mathematics is a valuable activity. Taking pure mathematics as a model for justifying metaphysics, then, I will argue that the reasons for believing that pure mathematics is valuable generalise to metaphysics. Thus, if one accepts that pure mathematics is valuable, then there is pressure to accept that metaphysics is valuable as well.

The critic of metaphysics cannot then escape the apology by taking issue with the type of value claimed for metaphysics. She cannot maintain that while metaphysics is valuable in some sense, the type of value it carries is not of the right kind to justify metaphysical research. For if she denies that the value I am claiming for metaphysics is enough to save it, she thereby imperils pure mathematics. The argument thus places an added burden on the sceptic of metaphysics. If one truly believes that metaphysics

³Ladyman and Ross and Maclaurin and Dyke seem to have in mind 'natural science' when they speak of science. That is, science of the mind-independent aspects of the world. But, of course, a great deal of science is about the mind-dependent aspects of the world and, indeed, of the mind itself, as is the case in much of psychological and cognitive science. As we shall see later on, focusing on the natural sciences makes it easier to see metaphysics as broadly useless.

is worthless (as some philosophers do), then one must either show why the argument I present fails, or one must give up on pure mathematics as well. Either way, by throwing its lot in with mathematics, I will have strengthened metaphysics against outside attack.

3. PRESENTING THE APOLOGY

In order to set the scene for the apology, I will begin by saying a bit about why it is that naturalists should agree that pure mathematics is valuable. After that, I will argue that the basis for valuing pure mathematics generalises to metaphysics.

3.1. NATURALISM AND MATHEMATICS

Naturalists who are sceptical of metaphysics are not usually sceptical of mathematics as well. Indeed, Ladyman and Ross maintain that “mathematics and science have undoubtedly borne fruits of great value” (2007, 16). Maclaurin and Dyke are also willing to give mathematics a free pass, writing that they “have no concerns about ordinary mathematical problem-solving” (2012, 293).⁴

Clearly, naturalists shouldn’t be sceptical of *applied* mathematics. Applied mathematics plays a central role in science and so, given the naturalist’s standing commitment to taking scientific methodology seriously, they should fully endorse mathematics of this kind. A naturalist may, however, be sceptical of pure mathematics that has no applications.

But such scepticism is not warranted. Naturalists should value pure, unapplied mathematics because it is an important source of applied mathematics. Indeed, a great many of the most important pieces of applied mathematics began their lives as pieces of pure, unapplied mathematical research. As Gowers puts the point “. . . the history of mathematics is littered with examples of areas of research that were initially pursued for their own sake and later turned out to have a completely unexpected importance” (2000, 7).

Gowers offers three examples. The first of these is the development of non-Euclidean geometry. The investigation of non-Euclidean geometry was carried out within pure mathematics before any application was conceived. However, the development of non-Euclidean geometries subsequently enabled Riemann to develop four-dimensional curved geometries which, as Gowers notes, “seemed to be an example of pure mathematics *par excellence* until it turned out to be exactly what Einstein needed for his general theory of relativity” (2000, 7).

The second example is number theory. In his apology for number theory, Hardy (1969) begins by noting that number theory has no applications. Indeed, this is a fact of which he was particularly proud. Useful mathematics can be used in war, he said, and so “a real mathematician”—a mathematician who works in pure, unapplied mathematics—“has his conscience clear” (Hardy, 1969, 140). It was surprising, then, when Rivest, Shamir, and Adelman (1978) used number theory to develop an encryp-

⁴Maddy’s naturalism also extends to both science and pure mathematics (see Maddy, 1992, 1997, 2000). This brand of naturalism has come to play a fairly central role in the philosophy of mathematics (see Colyvan, 2001). In so far as naturalists of this stripe are critical of metaphysics, the apology presented here places pressure on that critical attitude.

tion method that subsequently became the backbone for internet security, in point-to-point and wifi encryption.

The third example is knot theory. A mathematical knot is a closed loop that has been twisted into some configuration (imagine the two ends of a piece of string that are fused together, and then imagine that the string is twisted in various ways to form a knot-like structure). The unknot is a simple loop. A knot is isotopic to the unknot when, through a sequence of moves in which the loop is untwisted—known as Reidemeister moves—a knot can be transformed into the unknot. Research in knot theory focuses on (i) working out when two knots are isotopic and thus can be transformed into one another through a series of Reidemeister moves and (ii) determining which knots are isotopic to the unknot and which knots are not. Study in knot theory was, until recently, what Colyvan (2007) might call ‘mathematical recreation’: an interesting puzzle pursued purely out of mathematical interest with no foreseeable applications. That is, until Jones (1985) and Witten (1988; 2012) showed how to apply knot theory to fundamental physics, in Quantum Field Theory and String Theory.

Now, a naturalist who remains sceptical of pure mathematics may respond to these examples by drawing a distinction between two kinds of pure mathematics. On the one hand, there is pure mathematics that, while unapplied, has obvious potential for applications. On the other hand, there is pure mathematics that is unapplied, and has no obvious potential for applications. She might concede that the value she accords applied mathematics carries over to pure mathematics of the first kind, while denying that there is a similar pressure to value pure mathematics of the second kind.

However, the pure mathematics of non-euclidean geometry, number theory, and knot theory were all developed without any obvious applications. And these are not isolated cases. Bernstein (1979) offers ten more examples, including the investigation of Weierstrass functions which, only recently, have been used to model certain features of black holes (see, e.g., Gibbons and Vyska, 2012).⁵

Again the naturalist might respond by conceding the value of pure mathematics that has no *obvious* potential for applications, while maintaining that we should be sceptical of mathematics that is unlikely to become applied. However, as both Gowers (2000, 7) and Bernstein (1979, 252) argue, we are spectacularly bad at predicting which parts of pure mathematics have the potential to become applied, and which parts do not. And so we shouldn’t put much stock in our intuitive judgements about the likelihood of a piece of mathematics becoming applied.

But perhaps there are some areas of pure mathematics about which we can make reasonable judgements regarding the likelihood of application. The higher reaches of set theory may be like this. However, even for areas of pure mathematics that have a vanishingly small likelihood of being applied, the naturalist should value those areas. That’s because pure mathematics of this kind is applicable to other areas of mathematics. These areas of mathematics are, in turn, applicable to further areas and so on until we make it all the way back to scientific application. Indeed, there are chains of applicability leading from the most applied aspects of mathematics, to the most pure. Pure mathematics is needed to illuminate applied mathematics. The value that we place on applied mathematics therefore flows back into the pure domain via these chains of applicability. Colyvan makes this point in terms of the indispensability of mathematics.

⁵See Hamming (1980), Wigner (1960) for further examples.

He writes that:

If transfinite set theory is indispensable for analysis and analysis is indispensable for physics, then I say transfinite set theory is indispensable for physics ... this is the justification for the higher reaches of set theory that I endorse. (Colyvan, 2007, 113)

Colyvan's defense of transfinite set theory is bolstered by the fact that mathematics is a deeply interconnected field of research. Gowers (2000, 10–14) provides one example of this, showing how results from number theory, differential algebra and geometry all inform one another in unexpected ways. Further examples of the interconnectedness of mathematics are not difficult to find, and include the proof of Fermat's last theorem using ring theory; the use of set theory for a range of mathematical notions such as functions, measures, graphs, rings, and groups; and the wide applicability of group theory throughout mathematics. Because of these interconnections within mathematics "any attempt to purge mathematics of its less useful parts would almost certainly be very damaging to the *more* useful parts as well" (Gowers, 2000, 9).

In sum then, even if the naturalist does not *as a matter of fact* value pure mathematics, they should. Work in pure mathematics transitions into applied mathematics in a manner that is quite unpredictable, and so a blanket policy of endorsing pure mathematical research is the most reasonable way to maximise the potential for new applications. Furthermore, research in pure mathematics underwrites work in applied mathematics, due to the deep interconnectedness of mathematics as a discipline. So even if some branch of pure mathematics has little chance of ever being directly applied, it is justified via a chain of applications to more applied areas of mathematics.

3.2. KNOWLEDGE

So far I have argued that a naturalist should believe that pure mathematics is valuable. But like metaphysics, pure mathematics has no observational consequences. The naturalist, then, needs to argue that pure mathematics is valuable in a manner that metaphysics is not. Otherwise, the naturalistic endorsement of pure mathematics will spill over onto the metaphysical case as well.

To this end, the naturalist might appeal to knowledge. Pure mathematics, she might argue, produces knowledge. Of what? Of mathematical facts, facts regarding a peculiarly mathematical subject matter. Such facts include facts about numbers, sets, functions, groups, classes, geometries, and so on. Metaphysics, by contrast, does not produce knowledge. And so mathematics is valuable in a way that metaphysics is not.

But metaphysics *does* produce knowledge. Of what? At a minimum: knowledge of models. A metaphysical theory is a class of metaphysical models. A metaphysical model is analogous to a scientific model. Very roughly, on the semantic conception of models, a scientific model is an abstract object that represents an actual or hypothetical physical system. A metaphysical model is also an abstract object. However, a metaphysical model represents a peculiarly metaphysical subject matter. What is a peculiarly metaphysical subject matter? According to Paul (2012, 6), metaphysics focuses on the nature of things in the most general terms. Metaphysics thus seeks to tell us about the nature of laws, of parts, of persistence, of causation and more. This is in

contrast to science, which focuses on discovering the instances of those natures. So, for example, a scientific model might represent what causes what, whereas a metaphysical model represents what causation is. Similarly, a scientific model might represent what properties there are, whereas a metaphysical model represents what properties themselves are.

Of course, as Paul recognises, the distinction between metaphysical and scientific models may not be sharp. Science does sometimes seek to provide models of the nature of things (as is the case with time), and metaphysics does sometimes seek to tell us what things there are (as is the case with some debates in composition, where the disagreement is partly over how many objects exist in the world). Still, the basic distinction that Paul draws is a useful one, and it will serve us in what follows. I will thus assume that metaphysical models are models of the natures of things, in Paul's sense.

I have said that metaphysical models are abstract objects, but abstract objects consisting of *what*? Again, we can turn to Paul for help:

... a metaphysical theory can be understood as a class of models, where the models are composed of logical, modal and other relations relating variables that represent n-adic properties, objects and other entities. For example, a theory of composition can be thought of as a class of models of the composition relation such that some x s compose a y if and only if the activity of the x s constitutes life. The models we can take to be the theory are structures of abstract objects that represent activity-constituting objects standing in necessitation relations to abstract objects that represent composites or wholes of the activity constituting objects. The theory is a class of (suitably abstract) models, where these models are isomorphic to various instances of the activity-constituting relations between parts and wholes. (2012, 12)

A model has two components on this account. First, it features variables that represent the particular features of the world that are implicated in whatever it is we are trying to model, features such as events, objects, persons, minds and so on. Second, it features relations between those variables. As Paul notes, these can be logical relations—like entailment or identity—or modal relations—such as necessitation or counterfactual dependence. Presumably, however, they can also be mathematical relations, such as functions; or hypermodal relations, such as grounding and essential dependence; or even temporal relations, such as simultaneity and the earlier-than relation. Taken together, the variables and the relations give a model its structure.⁶

Philosophy journals are full of papers that can be broadly construed as providing us with knowledge of metaphysical models. For instance, we know what a model of causation is like if counterfactual dependence is sufficient for causation. Similarly, we have models of composition where the composition relation is unrestricted, restricted

⁶As another example of a model of this kind, Paul (2012, 14) points to the counterfactual theory of causation. One very simple class of metaphysical models of causation can be specified as follows. In each such model, there are variables that represent events. These events are related by temporal relations of precedence, along with relations of counterfactual dependence that connect the events one to the other. Structural equation models and their associated neuron diagrams are one type of model along these lines. But one can also produce less explicitly mathematical models, via a suitable modal logic and a semantics for counterfactual conditionals.

or absent entirely. We have models of properties according to which properties are universals or tropes, and we have models of grounding according to which grounding is an existing relation in the world. We also have models of modality, in the form of Lewis's modal realism and its competitors.

Paul believes that metaphysics goes further and also produces knowledge of how things really are, by establishing the truth or accuracy of a class of metaphysical models. The method by which the truth of a class of models is established is via cost/benefit analysis on the theoretical virtues of the models. The use of cost/benefit analysis in this context is something that Maclaurin and Dyke (2012) sharply criticise. I won't enter into this debate here except to point out that, as Maddy (1988a; 1988b) shows, the debate over the correct axioms for set theory is largely carried out by focusing on the elegance, utility and explanatory power of different axioms. Scepticism about the use of cost/benefit analysis thus threatens to carry over to pure mathematics as well.

Suppose, however, that metaphysics produces no knowledge of how things really are actually. Even if this is correct, we can sometimes gain knowledge of epistemic possibilities.⁷ For instance, because we have models of composition where the composition relation is unrestricted, and because our current knowledge of the world doesn't rule out that such models are correct, we know what it would be like for the actual world to be one in which mereological universalism is true.⁸ This makes metaphysical modelling even more analogous to scientific modelling. Models of quantum mechanics (such as the many-worlds interpretation) also provides us with knowledge of epistemic possibilities. Even if we cannot settle the case as to whether a given epistemic possibility is the way things are actually, we have still gained knowledge of the possibility itself via the model.

It may also be the case that metaphysics provides knowledge of non-epistemic possibilities, as Lowe (1998) suggests, though this is more controversial. The trouble is

⁷The view that metaphysics produces knowledge of epistemic possibilities is reminiscent of Russell's account of the value of philosophy more generally in *The Problems of Philosophy*. He writes:

Philosophy, though unable to tell us with certainty what is the true answer to the doubts which it raises, is able to suggest many possibilities which enlarge our thoughts and free them from the tyranny of custom. Thus, while diminishing our feeling of certainty as to what things are, it greatly increases our knowledge as to what they may be; it removes the somewhat arrogant dogmatism of those who have never travelled into the region of liberating doubt, and it keeps alive our sense of wonder by showing familiar things in an unfamiliar aspect. (Russell, 1912, 137)

⁸What is knowledge of an epistemic possibility? This depends on what epistemic possibilities are. There is a fair bit of controversy surrounding what it takes for something to be an epistemic possibility (see Huemer, 2007 for discussion). But for any account one cares to give, there will be a way to understand what knowledge of such possibilities amounts to. For instance, suppose that P is an epistemic possibility for S if P is consistent with everything that S knows. Then what it means for mereological universalism to be an epistemic possibility for a given metaphysician, M, is for mereological universalism to be compatible with what M knows. When M gains knowledge of mereological universalism they are thereby gaining knowledge of a theory that is compatible with their current state of knowledge. Note that epistemic possibilities are usually defined as possibilities for a particular person S. In the case of metaphysics, it makes more sense to focus on epistemic possibilities that are defined over the state of knowledge more generally and over all rational agents engaged in a particular field of enquiry. So we could say that P is an epistemic possibility for all scientists and metaphysicians if P is compatible with what all scientists and metaphysicians know. Call this a *robust* epistemic possibility. These are the kinds of epistemic possibilities that metaphysics can tell us something about.

that non-epistemic possibilities are usually thought to be metaphysical possibilities. If, however, one is sceptical of metaphysics, then perhaps one ought to be sceptical of metaphysical possibility as well. Note, however, that the naturalist under consideration does not reject metaphysics *tout court*. Metaphysical theories that possess observational consequences are deemed to be acceptable. The naturalist therefore believes that some metaphysical possibilities are in good standing—namely the actual ones. So it is not clearly open to her to be sceptical of metaphysical possibility in general.⁹

For present purposes, however, knowledge of the models themselves and of epistemic possibilities will suffice, and so I will set the non-epistemic possibilities aside. Assuming, then, that metaphysics and pure mathematics both produce knowledge, the naturalist cannot object to metaphysics on the grounds that it produces no knowledge. Instead, in order to break the symmetry between pure mathematics and metaphysics, the naturalist must argue that the knowledge that pure mathematics provides is valuable in a way that the knowledge that metaphysics provides is not.

There are four potential sources of value for mathematics. First, knowledge in pure mathematics is intrinsically valuable. It is intrinsically valuable just because any knowledge is intrinsically valuable. Second, knowledge in pure mathematics may be secure in a way that other knowledge isn't, potentially making it the most robust kind of knowledge there is. Third, knowledge in pure mathematics is significant. Results in mathematics illuminate mathematics for us in profound ways. Finally, knowledge in pure mathematics feeds into applied mathematics. It is the 'seed bank' for the development of scientific models of reality.

The first justification for pure mathematics invites little by way of discussion. Obviously, the reasoning regarding intrinsic value carries over to the metaphysical case straightforwardly. Instead, I will focus on the remaining three sources of value for knowledge in pure mathematics, drawing connections in each case to metaphysics.

3.3. SECURITY

One way to defend the idea that mathematical knowledge is peculiarly secure is to focus on proof. One might argue that mathematical results can be proved and that proof delivers security. Compare this with scientific results. Belief in a result produced by the scientific method does not have the same sense of security about it. When one comes to believe that a certain claim about physical reality is true based on the best science of the day, that belief may be overturned down the track (and probably will be!).

Metaphysics, it might be thought, is much more like science than it is like math-

⁹One could appeal to a notion of absolute possibility instead of metaphysical possibility and leave it open as to whether absolute possibility is a distinctively metaphysical modality. However, Clarke-Doane (2017) makes trouble for the idea there is a viable notion of absolute possibility, maintaining that any such notion is open to indefinite extensibility arguments. Better, I think, to specify the possibilities negatively, as any possibilities that are not nomic possibilities. On such a view, metaphysical models are similar to Lange's (2013) distinctively mathematical explanations. A distinctively mathematical explanation works "by showing how the explanandum arises from the framework that any possible causal structure must inhabit" (Lange, 2013, 505). The two kinds of model thus provide information about the same kinds of possibilities. The advantage of going this way is that, as Lange argues, distinctively mathematical explanations are a central aspect of science. So the naturalist has good reason to take these nomic impossibilities seriously.

ematics in this respect. Accordingly, whatever value accrues to mathematics in virtue of the security of mathematical knowledge, that value is not possessed by whatever knowledge might be gained from doing metaphysics.

But, in fact, mathematics and metaphysics are quite similar when it comes to the relative security of knowledge in the two domains. To see this, we need to differentiate between two notions of ‘proof’. The first sense of proof is just a matter of deducing some theorem from a set of axioms. The second sense of proof is closer to truth. Proving a theorem in this sense amounts to showing that the theorem is true because it follows from axioms that are deemed to be somehow ‘self-evident’ (more on this in a moment).

Now, if ‘proof’ is understood in the first sense, then metaphysical and mathematical results are equally open to being ‘proved’. Axiomatic systems for grounding, fundamentality, and mereology have been provided (see Fine 2012; Turner 2016; and Casati and Varzi 1999 respectively). There are also tense-logical axioms that can be used to axiomatise metaphysical accounts of time (see Prior, 1957). Results about grounding, parthood and time can be proved within these models based on their associated axiomatic systems. Indeed, pretty much any metaphysical model can be formalised and ‘theorems’ about the model can then be proved. It is just that, for the most part, metaphysicians tend not to proceed in this fashion. I suspect this because metaphysicians are more interested in working out what is *really the case* and so are more concerned with something like the second sense of proof outlined above.

Of course, in both the mathematical and the metaphysical cases proofs in the first sense are subject to a background logic. A theorem may be provable in, say, classical logic, but fail in intuitionistic logic or paraconsistent logic. Because there is a substantial debate over what the ‘correct’ logic is, it follows that such theorems may not be entirely future-proof. Developments in philosophical logic, or indeed in meta-mathematics, may alter the landscape of what follows from what in a substantial way. So it is unclear just how secure mathematical knowledge in this first sense of proof really is. At any rate, this doesn’t alter the fact that whatever degree of security there is to be had with respect to mathematical knowledge, that carries over to the metaphysical case as well.

This brings us to the second sense of proof. Exactly what to make of this second notion of proof depends on what it means to say that an axiom is ‘self evident’. One option is to construe ‘self-evident’ in something like the Fregean sense, which Jeshion (2000, 953) characterises as follows:

A proposition p is *self-evident* if and only if clearly grasping p is [a] sufficient and compelling basis for recognition of p ’s truth.

If we understand the security of mathematical knowledge in terms of the self-evidence of the underlying axioms, then it is doubtful that knowledge in metaphysics is secure in this sense. The most that we can say for the axiom-like claims that a metaphysician might use as the basis for building a metaphysical model is that the claims are ‘plausible’ or ‘intuitive’. But there is massive disagreement about even the most basic metaphysical claim within pretty much every area of metaphysics.

But there is widespread disagreement in mathematics as well. For instance, as Maddy (1988a; 1988b) argues, there is an ongoing and complex debate surrounding what the axioms of set theory are supposed to be. This is not restricted to set

theory. As Clarke-Doane (2017, 241) puts the point “for core claims in every area of mathematics—from set theory to analysis to arithmetic—*there are some* non-error-theorists who deny those claims.” Such disagreement within mathematics is difficult to square with the supposed self-evidence of the axioms.

The idea that the axioms are self-evident is out of step with mathematical practice in another way. The trouble is that the justification for the axioms is typically not a matter of their self-evidence, but rather, is based on the fact that the axioms can be used to derive other propositions that are, arguably, self-evident (Clarke-Doane, 2017, 242). This method of justifying mathematical axioms is therefore based on a process of reflective equilibrium. We start with a set of putative mathematical beliefs and then look at ways to systematise those beliefs using general, axiomatic principles. The axioms we should believe are the ones that provide us with the best systematisation.¹⁰

In light of these two worries about the Fregean self-evidence of the axioms, it makes sense to focus on mathematical propositions more generally (theorems and axioms) and to appeal to some other notion of self-evidence that can take full account of mathematical practice. The trouble, however, is that there doesn't seem to be a viable notion of self-evidence available. Clarke-Doane (2017, 242–244) considers a range of options. Self-evidence might be a matter of unanimous agreement on a proposition, he suggests, or a matter of widespread agreement on a proposition, or a matter of a proposition being generally found to be plausible. In each case, Clarke-Doane argues that the disagreement within mathematics is too widespread and fundamental to support any plausible notion of self-evidence. So there is no way to argue from a notion of self-evidence to the conclusion that mathematical knowledge is peculiarly secure. Mathematics is just as open to disagreement as metaphysics is.

In sum, then, if the safety of mathematical knowledge corresponds to the way that mathematical theorems can be proved in the first sense described above, then metaphysical knowledge possesses this security as well. If the safety of mathematical knowledge is to do with proving a claim from self-evident propositions (or showing that a claim can be used to prove self-evident propositions) then metaphysical knowledge isn't secure in this sense, but neither is mathematical knowledge. Either way then, there is parity.

3.4. SIGNIFICANCE

We come now to the significance of mathematical results. The reason why knowledge in pure mathematics is special as compared to metaphysical knowledge is that mathematical facts are significant. For present purposes, I will use the account of significance offered by Hardy:

We may say, roughly, that a mathematical idea is ‘significant’ if it can be connected, in a natural and illuminating way, with a larger complex of other mathematical ideas. (1969, 89)

Of course, not every mathematical fact is significant in this sense. Hardy gives, as an example of insignificant mathematics, the mathematics of chess moves. While Hardy accepts that the mathematics of chess is genuine mathematics, the mathematical

¹⁰A similar picture of mathematics is offered by Shapiro (2009, 198–204).

facts that we might prove within chess are not significant. This is evidenced by the lack of connections between those theorems and other areas of mathematics. Despite the presence of insignificant mathematics, pure mathematics *as a whole* is valuable, because *some* results in pure mathematics are significant in Hardy's sense.

But let us pause for a moment and consider why it is that significant results are possible in mathematics. Such results are possible because mathematics is a deeply interconnected field. So a result proved in one domain has the ability to ramify through a number of different areas of mathematics in profound ways.

But notice that metaphysics is interconnected in much the same way. The stock-in-trade of metaphysicians—causation, time, ground, property, modality, mereology, constitution, and so on—all bear important connections to one another. So when one builds a metaphysical model, that model rests atop a web of conceptually interrelated notions. If one is then able to demonstrate something about one of those notions using such a model, or if one is able to build an entirely unique model that makes use of one of those notions, then this can—and often will—have ramifications throughout the broader space of known metaphysical models.

Consider, for instance, Lewis's (1986) modal realism. Whatever else one thinks of Lewis's modal realism, it is an impressive and detailed theory, one that represents a space of concrete worlds each of which is spatiotemporally disconnected from our own. By developing modal realism, Lewis managed to give possible worlds semantics new life in philosophy. That semantics, and indeed, the modal realist theory itself, had widespread ramifications throughout metaphysics, and beyond, into linguistics, logic and more. Lewis's a priori demonstration of how modal realism works and his argumentative mapping of modal realist models resulted in an understanding of possible worlds that has profoundly influenced metaphysics, giving rise to new metaphysical models of causation, the mind, properties, and more.

Here's another example: the work on grounding by Bennett (2011), Fine (2012), Rosen (2010), Schaffer (2016; 2009) and others. These philosophers, it is fair to say, have developed a range of different models of grounding. This has led to new applications of grounding to build new metaphysical models of other phenomena, including time, causation, mind/body relations, and composition. Grounding has illuminated connections between things that we could not see before. Even if one thinks that grounding is not a real part of reality, the investigation of grounding models within metaphysics has deepened our knowledge of what grounding is like, and that knowledge has ramified through metaphysics in interesting ways.

So metaphysics is significant in Hardy's sense. The knowledge of models produced within metaphysics can be illuminating and fruitful in the same way that mathematical results can be. Of course, not all knowledge in metaphysics is like this. Lots of it isn't. But, as indicated, the same is true in pure mathematics. There's non-significant work in both fields.

3.5. APPLICATION

This brings us to (arguably) the most important reason for believing that the knowledge produced in pure mathematics is special. As we saw in Section 3.2, pure mathematics feeds into applied mathematics. Mathematics gets applied within science in two main ways. First, it gets applied as a *tool* for developing a scientific theory. For in-

stance, mathematics might be used as a mere calculational framework for explaining a physical phenomenon; the mathematics itself is not assumed to play a substantial role in representing or explaining the physical system being modelled. Second, mathematics gets applied in science as a way of explaining the structure of some physical system (for more on the two roles for mathematics, see Baker 2005; Colyvan 2001; Saatsi 2011). This sometimes takes the form of what Lange (2013; 2016) calls distinctively mathematical explanation: cases in which the mathematics itself is carrying the explanatory load.¹¹

As we also saw in Section 3.2, the traffic between pure and applied mathematics ultimately justifies pure mathematics. In its most general form, the justificatory structure of the situation is this:

1. Some knowledge in pure mathematics can be used to make gains outside of mathematics.
2. These gains are gains worth having.
3. So some knowledge of pure mathematics is knowledge worth having.

Notice that the basic justificatory strategy outlined above doesn't hang on the fact that mathematics is applied within science *per se*. Rather, the strategy requires only that mathematics is applied outside of mathematics in a way that is valuable to us. For example, mathematics is also applied at the level of meta-logic. Assuming that we value meta-logic, then this source of value flows back to pure mathematical knowledge as well. Similarly, if mathematics were extremely useful for producing art, and had no function in science, and if we valued art, then that value would provide a further reason to develop pure mathematics. Because I am focusing on naturalism, what ultimately matters is application within science. Nonetheless, the structure of the justification itself permits the discovery of value for pure mathematics based on its application elsewhere.

However, the justificatory story needs to be expanded to show that knowledge in mathematics quite generally is valuable, even when it doesn't feed into applications. For that, a bridging argument is required. We considered two such arguments in Section 3.2. First, we are very bad at predicting which parts of pure research in mathematics might become applied. So it is reasonable to pursue a wide range of pure research projects in mathematics. In short, the path from 'some mathematical knowledge is worth having' to 'all mathematical knowledge is worth having' is paved by ignorance. Second, there are chains of applicability leading from the most applied aspects of mathematics to the most pure. These chains of applicability provide justification for even the highest reaches of set theory.

Now, what we need to consider in the metaphysical case is whether research into metaphysics yields gains outside of metaphysics. In order to give some shape to our discussion of this issue, it is useful to draw a distinction between 'pure' and 'applied' metaphysics. Pure metaphysics is metaphysical research that stays entirely within metaphysics, and never strays beyond. Metaphysics of the kind that I have been discussing—research into metaphysical theories that lack observational consequences—is a type of pure metaphysical research. It fails to reach beyond metaphysics and, in particular, fails

¹¹Applied mathematics is extraordinarily successful. Indeed, for some, mathematics is unreasonably successful in science, given its origins in pure mathematics (see Hamming, 1980; Wigner, 1960).

to connect up with science. Applied metaphysics is metaphysical research that extends beyond metaphysics in some manner. One type of applied metaphysics is the application of metaphysics within science. But there can be other types of applied metaphysics as well. As in the mathematical case, if a piece of metaphysics is useful for something that we value, then that value will infect metaphysics regardless of whether the source of value is scientific or not. So long as knowledge from metaphysics is applied in a useful way outside of metaphysics, that is sufficient to bring mathematics and metaphysics into broad parity with regard to the justificatory story outlined above. To convince the naturalist, however, it must be shown that metaphysics can be applied within science as well.

As in the mathematical case, we don't need every piece of pure metaphysics to become applied. All we need is for some of it to become applied. We can then defend research in pure metaphysics more generally in the same two ways that we defend research in pure mathematics. We don't know which parts of metaphysics will come to be useful down the track, and so we should allow a wide range of projects to flourish. Moreover, as already discussed, metaphysics—like mathematics—is a deeply interconnected field. Because of this, we can defend the importance of pure metaphysics based on its application within applied metaphysics. So, for instance, with respect to the application of the metaphysics of causation discussed below, the particular model of causation being applied draws on other notions for which we require further metaphysical models. Causal models usually have temporal precedence built-in and so models of time are called upon. Similarly, causal models are usually models of relations between objects, events or properties, and so models of each of these will be needed to guide the application of such models. These models, in turn, draw on other concepts and other models, such as models of grounding or laws and on it goes. At some point in the chain, for some metaphysical model or other, the models that we produce have no applications outside of metaphysics. Nonetheless, these models are being used to illuminate the more applied aspects of metaphysics in an important way. We can thus discern similar chains of applicability in the metaphysical case to the ones that Colyvan uses to justify pure mathematics.

There are, broadly, two dimensions of applied metaphysics. First, there is metaphysics that is used outside of metaphysics, but within philosophy. Call this 'internally applied metaphysics'. Second, there is metaphysics that is used outside of philosophy. Call this 'externally applied metaphysics'. Within each dimension, metaphysics can—in principle, at least—be applied in the same two ways as mathematics. First, it can be applied as a tool or framework for understanding a phenomenon, without accurately representing or explaining that phenomenon. Second, it can be used to accurately represent or outright explain some phenomenon. Both applications promise to provide a source of value for pure metaphysics.

It is straightforward to make the case that metaphysics is internally applied. Here's one obvious example: personal identity and ethics. Theories of personal identity do not have any observational consequences. Persons may be animals, bodies, minds, conventions, or some mixture. There is no observation we could make that would confirm one theory over another. Such theories, however, have applications in ethics. One of the focuses of ethics is determining what 'I' ought to do or what 'my' ethical responsibilities are. Theories of personal identity can be applied within ethics to provide an account of what this 'I' or 'My' is. For example, the psychological continuity theory of

personal identity implies that I am the same person as someone in the future, or past, who bears important mental relationships to me (such as continuity of consciousness). This theory can be applied within a theory of normative ethics to circumscribe the target of ethical obligations we owe to ourselves and others.

Applying the metaphysics of personal identity in this fashion can lead to significant gains within normative ethics. For example, a better understanding of what a person *is* can help to dismantle objections against a given normative theory. One salient example of this is Parfit's *Reasons and Persons*. The focus of that book is the application of the metaphysics of personal identity to normative ethics. Parfit argues for a reductionist approach to persons:

We are not separately existing entities, apart from Our brains and bodies, and various interrelated physical and mental events. Our existence just involves the existence of our brains and bodies, and the doing of our deeds, and the thinking of our thoughts, and the occurrence of certain other physical and mental events. Our identity over time just involves (a) Relation R-psychological connectedness and/or psychological continuity-with the right kind of cause, provided (b) that this relation does not take a 'branching' form, holding between one person and two different future people. (1984, 216)

By adopting this approach to persons, Parfit is able to demonstrate that certain approaches to normative ethics (such as egoism) are self-defeating. He also manages to provide a strong defense of utilitarianism on the basis of his reductionist theory of personal identity.

Whether or not Parfit's arguments are successful, the application of personal identity to the debate surrounding normative ethics has been extremely influential. Looking at normative ethics in this fashion has enabled us to develop new versions of utilitarianism (such as the one that Parfit defends). It also gave new scope to the debate about what we owe to future generations. Precisely what we owe to future individuals turns, in some sense, on how we construe personal identity. Parfit argues that we have the same obligations to the temporally distant needy as we do to the spatially distant needy based on how we understand both what it is to be a person, as well as the relationships between persons. Again, we might not agree with the theory of normative ethics that Parfit develops, but it is hard to deny that the theory itself, which involves the application of metaphysics, changed the face of research in this area.

Another example of internally applied metaphysics is the ongoing work within experimental philosophy conducted by Knobe (2009), Knobe and Fraser (2008), Hitchcock and Knobe (2009) and others. Particular models of causation (such as counterfactual and process-based theories), and particular aspects of causal modelling (such as causal omission) have been used to form the basis of empirical experiments that aim to examine the intuitions and concepts that people possess. This work has then fed back into philosophy more generally, as a basis for rethinking philosophical methodology.

What about externally applied metaphysics? In particular, are there any cases in which metaphysics comes to be applied in science? Recently, Hawley (2016) has outlined some examples of what she calls 'applied metaphysics', some of which involve the application of metaphysics within science. She cites three particular examples: the use

of ontology within computer science; the investigation of social kinds within social science; and the investigation of natural kinds within biology, medicine and psychology. All three cases are examples in which scientists appear to be doing metaphysics. Computer scientists build models of objects and categories in order to provide the resources for artificial intelligence to successfully navigate the world, or make judgements about it. Social scientists actively study the nature of various social kinds as a basis for better understanding societies, cultures and political systems. Psychologists are constantly attempting to determine which mental health disorders are natural kinds and which are not, for the purposes of better equipping the DSM.

While all three cases are interesting, they do not involve a particular metaphysical theory being picked up and then used in science. Instead, certain topics that metaphysicians have traditionally investigated are being explored in other fields. That said, the examples that Hawley identifies do help to build a case for the possibility that a particular metaphysical theory might make its way into science. For example, as Hawley notes, different theoretical approaches to natural kinds may have different implications for how we classify mental health disorders. Thus, the development of new theories of natural kinds has the potential to shape this research, especially if one of those theories provides insights that push science along. Similarly, theories of objects and of properties may conceivably feed into ongoing work on ontologies within computer science, and provide future directions for the production of artificial intelligence and for software engineering more generally.

Still, it would be nice to have some more concrete cases of externally applied metaphysics that focus on science. The difficulty with finding such cases is that, historically speaking, the division between metaphysics and science is quite new. So, when looking at the history of science, it is unclear what is a metaphysical theory and what is a scientific theory. Trying to find traffic between two disparate fields, then, is challenging, because up to a point there is no traffic to speak of; there is simply a single discipline doing different things.

That being said, it is possible to identify two cases in which pure metaphysical research has been applied within science. The first case comes from the metaphysics of mind. Consider the extended mind thesis. According to the extended mind thesis, the mind literally extends beyond the skull. Thus, objects in the environment can be construed as partly constitutive of the mind. Smart phones become a case in point. These days it is common to outsource a range of cognitive tasks to one's smart phone. For instance, instead of trying to remember something, one might make note of it in one's phone. According to Clark and Chalmers (1998), when an object comes to play the functional role of memory, it can be properly thought to constitute a part of the mind. In this way, one's mind extends outward and includes part of the phone, just as much as it might include some part of the brain. It is arbitrary to restrict the mind to only certain parts of our bodies.

The extended mind thesis, on its own, does not have any observational consequences. It is compatible with all of the empirical data we have regarding neurophysiology, psychology, and cognition. Indeed, if we compare two worlds, one in which a smart phone is literally a part of the mind, and another in which the smart phone is used in the same way, but is not a part of the mind, it is quite difficult to see what the empirical difference between two such worlds would be.

In the last decade, however, there has been an explosion of research outside of

metaphysics and, indeed, outside of philosophy that applies the extended mind thesis. The extended mind thesis has been applied in the context of linguistics and cognitive science to aid in the explanation of language acquisition (see Atkinson, 2010; Geert, 2008). It has been used to help explain and understand group cognition in social psychology and cognitive science (see, e.g., Theiner et al., 2010). It has also been used to help understand and explain memory (cf. Barnier et al., 2008). To give a sense of the scope of application of this thesis in science, the original Clark and Chalmers paper has now been cited almost 4000 times since 1998, making it one of the most cited articles in philosophy. Imagine if research within pure metaphysics had been shut down in 1997 and the extended mind thesis had never been published. We would have lost something valuable: a piece of pure metaphysical research that subsequently came to be applied within science in important ways.

The extended mind case is one in which a piece of metaphysics is brought into science to represent or explain a particular phenomenon in the world, and so it is analogous to the use of a pure mathematical theory in science toward a similar end.¹² The next case is an example in which a metaphysical theory is used as a tool or calculating framework for doing science. The case comes from the metaphysics of causation. David Lewis's (1973) theory of causation forges a strong connection between causation and counterfactual dependence. Lewis's theory, taken on its own terms, has no observational consequences. The theory is also widely regarded to fail as a theory of causation, because of the numerous counterexamples that it faces. And yet, Lewis's approach to causation shaped the development of causal modelling practices within science.

The primary approach to causal modelling is the structural equation modelling framework. One of the central figures in the development of causal modelling modelling is Judea Pearl. Pearl, a statistician and computer scientist, developed much of the formalism for causal modelling, and, along with his collaborators, is largely responsible for the widespread use of causal modelling in a huge range of scientific applications.

The early development of Pearl's causal modelling framework was heavily influenced by both Lewis's counterfactual theory of causation and Lewis's analysis of counterfactuals more generally (see, e.g., Pearl, 1997, 1995). In his later work, the connection between causation and counterfactuals forged by Lewis constitutes a central plank in Pearl's structural equation framework.

Indeed, it is difficult to overstate the contribution that the metaphysics of causation has on Pearl's work through Lewis. Pearl's semantic analysis of counterfactuals is based on Lewis's (Pearl, 2000, 238–243) and many of the core examples that are used to develop metaphysical models of causation—such as omission, preemption and overdetermination—are used by Pearl as a basis for testing his own theory (Pearl, 2000, 309–313). The importance of causal preemption and omission is also evident in later literature on causal modelling, in which preemption cases are used to test variations on Pearl's approach, as well as the application of causal models in particular contexts (see, e.g., Love et al., 2016; Stephan and Waldmann, 2018). Researchers outside of philosophy are not just drawing on the features of the structural equation approach to develop

¹²Such as the use of knot theory to explain why Sara cannot untie her shoe laces. Sara cannot untie her shoelaces because the laces are knotted into a trefoil knot, and the trefoil knot is not isotopic to the unknot. So there is no sequence of moves that Sara could make to get the knot out. See Lange (2013) for discussion.

their theories either, they, like Pearl, are reaching back to Lewis’s work to inform their views (see Love et al. 2016, who explicitly draw on Lewis’s account of causation).

According to Lewis’s counterfactual theory of causation, causation is analysed as the ancestral of a chain of counterfactual dependencies. Lewis’s exact account of causation is not imported into the causal modelling framework. What is imported, however, is the deep connection between counterfactual dependence and causation. In its barest form, a causal model is a set of equations with the following form (Pearl, 2000, 27):

$$x_i = f_i(pa_i, u_i), i = 1, \dots, n$$

To explain the formalism: pa_i corresponds to the set of variables that ‘determine’ the value of X_i (more on this notion of ‘determination’ in a moment) and U_i is an error term which represents any aspect of the physical situation that is not being duly represented by the model. Basically, what the formalism above encodes is a system of dependencies between variables. One of these variables will represent the particular physical phenomenon that is being modelled (the X_i); the others will represent aspects of the physical system that are thought to be responsible for that phenomenon (the pa_i). The equations themselves encode the way in which the pa_i are causally responsible for the x_i . According to Pearl, the structural equations correspond to counterfactual relationships. Pearl makes this clear as follows (note that by the counterfactual relationship of “is determined by”, Pearl seems to mean ‘would’ counterfactuals, especially given what he goes on to say in the rest of the quotation).

The equality signs in structural equations convey the asymmetric counterfactual relationship of “is determined by,” and each equation represents a stable autonomous mechanism. For example, the equation for Y states that, regardless of what we currently observe about Y and regardless of any changes that might occur in other equations, if variables $(Z, Z_2, Z_3, \epsilon_Y)$ were to assume the values $(x, z_2, z_3, \epsilon_Y)$, respectively, then Y would take on the value dictated by the function f_Y . (Pearl, 2000, 69)

The connection between causation and counterfactual dependence preserved in Pearl’s model is crucial to both understanding how the models work, and to using them to model physical systems. Without it, the structural equations lack an interpretation, and so it is unclear even what the family of equations corresponding to the model would be. Pearl’s theory thus owes a heavy debt to Lewis’s. Not that Pearl ever says otherwise. Pearl is quite up front about the relationship between the causal modelling framework and philosophical work on causation in metaphysics that precedes it.

The causal modelling framework serves as “the most popular approach to causal analysis in the social sciences” (Bollen and Pearl, 2013, 301). The causal modelling framework is also used extensively in epidemiology,¹³ psychology,¹⁴ engineering,¹⁵ and in the analysis of traffic conditions.¹⁶ It is the counterfactual core of the theory

¹³Greenland et al. (1999), Greenland and Brumback (2002), Kluve (2004), Oakes (2004), Vander-Weele (2016).

¹⁴Stephan and Waldmann (2018), Lombrozo (2010).

¹⁵Love et al. (2016).

¹⁶See Cheng et al. (2018), Golob (2003), Mitra (2016).

that is partly responsible for the now widespread use of the causal modelling framework. The connection between counterfactuals and causation captured within the framework enables the use of causal models as a basis for understanding interventions. A causal model can be made to represent an actual physical system by setting the values of the variables to reflect the actual makeup of the system, and by setting the structural equations so as to reflect known dependencies. Once a causal model has been used to capture an actual system, we can consider interventions into the system. Each intervention involves changing the value of one of the variables in the model in order to gauge the output of the model itself. The results of interventions are determined by the counterfactuals underpinning the structural equations within the model. Interventions on causal models are a powerful way to gather empirical predictions from a causal set up, and constitute an important basis for both designing and modelling experiments.

Of course, Lewis's theory is *not* a theory of causal modelling. But then, Riemann's theory of four dimensional geometry was not a theory of space-time. Moreover, just as only certain aspects of Lewis's theory are retained within a causal modelling framework, only certain aspects of Riemann's theory are carried over into general relativity. In neither case is a theory outside of science—be it mathematical or metaphysical—carried holus-bolus over into an empirical setting. Rarely are things ever so neat as all that. Nonetheless, in both cases a core strand of pure research is subsequently used in application within science to great effect.

4. OBJECTIONS

This concludes my initial presentation of the apology. To recap, the naturalist should accept that pure mathematics is valuable. But, pure mathematics, like metaphysics, has no observational consequences. To stop the endorsement of pure mathematics from generalising to metaphysics, then, the naturalist must find some difference between the mathematical and metaphysical cases. But I have argued that there is no difference: mathematics and metaphysics both provide us with knowledge, and enjoy the same benefits (such as they are) when it comes to security, significance and applicability. So the naturalist should believe that metaphysics is valuable. I anticipate five objections.

4.1. METAPHYSICS IS MATHEMATICS

First objection: I have said that metaphysics produces knowledge of models which can feature mathematics. But, one might argue, pure mathematics also produces knowledge of models which feature mathematics. So, one might reply, I have claimed that metaphysics provides the same knowledge as pure mathematics. It is no wonder, then, that metaphysics enjoys the same justifications as pure mathematics!

However, while some metaphysical models may be mathematical models, it doesn't follow that knowledge within metaphysics just is pure mathematical knowledge. It is useful to differentiate between two senses of the phrase 'mathematical model'. First, a 'mathematical model' might be any model that both uses mathematics and is a model of some peculiarly mathematical feature, such as numbers or sets. Second, a 'mathematical model' might be any model that uses mathematics but is not a model of a mathematical feature.

In so far as pure mathematics produces knowledge of models, it usually produces models of the first kind: mathematical models of some peculiarly mathematical feature. Some of the models produced in metaphysics may also be of this type. For instance, Lewis (1991) collapses the distinction between mereology and set theory. If Lewis is right, then mereology is mathematics (or at least one part of mathematics is mereology, depending on how we read the Lewis project). So mereological models may be mathematical models in the first sense. But most models produced within metaphysics are not like the mereological case. Metaphysical models may feature mathematics but they are still typically models of peculiarly metaphysical things, like causation.

Thus, when I say that metaphysics produces knowledge of models, this claim is compatible with metaphysics producing knowledge of mathematical models in the second but not the first sense, which is the distinct purview of pure mathematics. This means that there is a distinction between the knowledge that pure mathematics provides and the knowledge that metaphysics provides. In this way, metaphysics is similar to science. Both metaphysics and science produce models of non-mathematical things that *use* mathematics. However, in both cases such models are not models *of* a mathematical feature.

4.2. OVER-GENERALISATION

Second objection: my argument proves too much; it overgeneralises. If the apology I have offered goes through, then the production of absurd metaphysical models is also valuable. For instance, I might offer a theory of causation according to which x causes y if and only if x is a potato and y is a potato. One could develop this theory so that it could not be refuted or established by any empirical observation. One could also develop it to an arbitrary degree of rigour. But, one might contend, if my argument works, then research into this theory is valuable. A journal should, in principle, be willing to publish a paper on my potato theory.

I am willing to bite the bullet and admit that research into such theories has some value. But I don't think that this is a problem. It is open to the pure mathematician to investigate any piece of mathematics that interests them. So, for instance, one could investigate a mathematical model in which the addition function is non-commutative (i.e., $1+2 \neq 2+1$) and the only numbers are 1, 0 and 2. Doing so is likely to seem a bit crazy to most mathematicians. But there is nothing wrong about a mathematician spending their time on such research. It may just not be as valuable as other research, and so there is little motivation to do so. The same is true in the metaphysical case. One could explore a pretty wild metaphysical model, but doing so may not yield much that is interesting. Not every piece of metaphysical research is of equal value. Some models are more important to explore than others.

4.3. THE SUCCESS OF MATHEMATICS

The third objection focuses on the success of applied mathematics. Applied mathematics is much more successful than applied metaphysics. The degree of success matters for legitimacy one could argue. Pure mathematics is valuable only because of the degree of success enjoyed by applied mathematics. If applied mathematics were much less successful, there would be a question mark hanging over the value of pure mathematics

as well.

But is that right? Is applied metaphysics a less successful enterprise than applied mathematics? The massive impact of Lewis's counterfactual theory of causation on science should give one pause to reconsider. The same is true of the extended mind case, which, as discussed, is enjoying stunning success in the applied domain.

Once we bring the causation and extended mind cases into view, it is not too difficult to think of other cases as well. Consider, for instance, the debate over relationism and substantivalism about space and time. Very roughly, the substantivalist maintains that space and time exist independently of the objects and events that are located in space and time. Relationists, by contrast, maintain that space and time do not exist independently of the objects and events located in space and time. Both substantivalism and relationism lack observational consequences.¹⁷

According to Slavov (2016), the development of special relativity was inspired by relationism. Einstein (1998, 220), in a letter to Schlick (quoted in Slavov 2016, 143), speaks of the influence that Mach's philosophy had on his own work, which is most likely a reference to Mach's relationism.¹⁸ This is borne out in the theory itself. Time, in special relativity, is not an absolute container in which events happen. Rather, time—and certainly the simultaneity of events—is a fundamentally relational notion. Of course, relationism is ultimately discarded in Einstein's general theory of relativity, due to conceptual difficulties with reconciling relativity with relationism. As a consequence, Einstein applies a substantivalist model of spacetime to produce a total theory. But here too we see the influence of the substantivalist/relationist debate on Einstein's work (Romero, 2017, 143).

I have no doubt that there are other cases of applied metaphysics waiting to be brought to light. The history of physics is a good place to look, given the historical relationship between physics and metaphysics. The apparent lack of such cases is, I submit, a reflection of the fact that we, by and large, haven't been looking for them, rather than a reflection of the uselessness of metaphysics.

4.4. OBSERVATIONAL CONSEQUENCES

Fourth objection: in my discussion of applied metaphysics I claimed that theories without observational consequences are being applied within science. For instance, theories of the extended mind are being applied within linguistics. Presumably, however, if those scientific theories in which metaphysics is being applied are any good, then those theories have observational consequences. But then it is wrong to say that the metaphysical theories themselves lack observational consequences. Those theories must have observational consequences in order to serve as any part of a theory that has observational consequences.

This objection rests on a mistake. To see the mistake, consider the mathematical case again. The pure mathematical core of a given scientific theory does not have ob-

¹⁷Newton's chief objection against relationism was that it cannot account for certain observations (this is what Newton's bucket argument is supposed to show). Mach offers a version of relationism, however, that is immune to Newton's criticisms. Thus, if we compare Newton's substantivalism and Mach's relationism, then we appear to have two distinct models that lack observational consequences, at least with respect to the observational data that was available at the turn of the twentieth century.

¹⁸Norton (2004) makes a related point.

servational consequences on its own. Situate it within a scientific theory, however, and the entire theory does often have observational consequences. What happens is that the mathematics in combination with other components of a scientific theory produces an overall empirical picture that has an observational upshot, an upshot that the mathematics alone does not have. The same is true with metaphysics. The metaphysical component of a scientific theory alone may have no observational consequences, but the total theory in which it is situated may, and it may come to have those consequences precisely because of the metaphysical component that is being added.

4.5. INDISPENSABILITY

The fifth objection focuses on the traffic between pure and applied mathematics. I have suggested that one reason to value research in pure mathematics is that such research crosses over into the applied domain. I went on to argue that the same is true of metaphysics. But, one might respond, there is a difference between the mathematical and metaphysical cases. In the mathematical case, mathematics is applied in an indispensable way within science.¹⁹ The same is not true of metaphysics. The difference matters. It is only because some mathematical research comes to be indispensably applied within science that research into pure mathematics is thereby imbued with value. Because metaphysics is not indispensably applied within science, pure research in that area does not inherit the value of applied research.

There are two responses to this line of criticism. The first is to deny that metaphysics is dispensable to science. The use of causal modelling as a tool for conducting scientific investigation is now an indispensable component of research across a number of fields, including social science research. Similarly, the extended mind thesis has come to play an indispensable explanatory role in science, by helping to explain phenomena such as language acquisition and group cognition.

The second response is to deny that mathematics must be indispensable to science for pure research in mathematics to be valuable. Even if it could be shown that mathematics is ultimately dispensable to science because our scientific theories can be formulated in non-mathematical terms (a very big ‘if’, though see Field, 1980), this would do little to undermine the usefulness of mathematics to science. Mathematics may be a dispensable framework, but it can nonetheless help us to drive science forward, which is sufficient to confer value onto research within pure mathematics. In short, the indispensability of mathematics to science is orthogonal to the question of whether pure research in mathematics is valuable. What matters is the usefulness of the application. So too in the metaphysical case. Even if metaphysics is only applied within science in a dispensable manner, if its application yields real benefits then value will be conferred onto pure metaphysical research nonetheless.

5. CONCLUSION

It is time to take stock. I have argued that research into metaphysics is valuable. Such research is valuable for the same reason that research in pure mathematics is valuable.

¹⁹See Colyvan (2001) for a defense of the indispensability of mathematics to science. See Field (1980), Maddy (1992) for responses.

Metaphysical knowledge is as secure as mathematical knowledge, it is also significant and it feeds into applied work in a similar way to mathematical knowledge. Complaints against metaphysics, then, should be filed against the mathematics department as well. The question remains as to just how valuable metaphysics is. Arguably, it is not as valuable as pure mathematics, but is it valuable enough to justify the substantial distribution of intellectual and economic resources in a resource-poor landscape? I believe so. Defending that claim, however, is a topic for another paper.

ACKNOWLEDGEMENTS

I would like to thank two anonymous referees and an anonymous editor at this journal for their extremely useful feedback on this paper. I would also like to thank audiences at the 2015 Australasian Association of Philosophy Conference and Swarthmore College for their comments on earlier versions of the paper. Finally, I would like to thank Suzy Kilmister and David Ripley for extensive discussion of this paper, and for their hospitality in the woods of Storrs, Connecticut where early drafts of this paper were written. The final stages of research on this paper were partly funded by two ARC grants: DE180100414 and DP180100105.

REFERENCES

- Atkinson, Dwight (2010). Extended, Embodied Cognition and Second Language Acquisition. *Applied Linguistics*, 31(5), 599–622.
- Baker, Alan (2005). Are There Genuine Mathematical Explanations of Physical Phenomena? *Mind*, 114(454), 223–238.
- Barnier, Amanda J., John Sutton, Celia B. Harris, and Robert A. Wilson (2008). A Conceptual and Empirical Framework for the Social Distribution of Cognition: The Case of Memory. *Cognitive Systems Research*, 9(1–2), 33–51.
- Bennett, Karen (2011). By Our Bootstraps. *Philosophical Perspectives*, 25(1), 327–353.
- Bernstein, Dorothy L. (1979). The Role of Applications in Pure Mathematics. *The American Mathematical Monthly*, 86(4), 245–253.
- Bollen, Kenneth and Judea Pearl (2013). Eight Myths About Causality and Structural Equation Models. In Stephen L. Morgan (Ed.), *Handbook of Causal Analysis for Social Research* (301–328). Springer.
- Casati, Roberto and Achille Varzi (1999). *Parts and Places: The Structures of Spatial Representation*. MIT Press.
- Cheng, Xiaoyun, Yu Cao, Ken Huang, and Yuejiao Wang (2018). Modeling the Satisfaction of Bus Traffic Transfer Service Quality at a High-Speed Railway Station. *Journal of Advanced Transportation*, 12. <https://doi.org/10.1155/2018/7051789>
- Clark, Andy and David J. Chalmers (1998). The Extended Mind. *Analysis*, 58(1), 7–19.

- Clarke-Doane, Justin (2017). Modal Objectivity. *Noûs*. Advance online publication. <https://doi.org/10.1111/nous.12205>
- Colyvan, Mark (2001). *The Indispensability of Mathematics*. Oxford University Press.
- Colyvan, Mark (2007). Mathematical Recreation versus Mathematical Knowledge. In Mary Leng, Alexander Paseau, and Michael D. Potter (Eds.), *Mathematical Knowledge* (109–122). Oxford University Press.
- Einstein, Albert (1998). *The Collected Papers of Albert Einstein, Volume 8: The Berlin Years: Correspondence, 1914-1918* (Ann M. Hentschel, Trans.). Princeton University Press.
- Field, Hartry (1980). *Science Without Numbers*. Oxford University Press.
- Fine, Kit (2012). A Guide to Ground. In Fabrice Correia and Benjamin Schneider (Eds.), *Metaphysical Grounding: Understanding the Structure of Reality* (37–80). Cambridge University Press.
- Geert, Paul Van (2008). The Dynamic Systems Approach in the Study of L1 and L2 Acquisition: An Introduction. *The Modern Language Journal*, 92(2), 179–199.
- Gibbons, Gary W. and Martin Vyska (2012). The Application of the Weierstrass Elliptic Functions to Schwarzschild Null Geodesics. *Classical and Quantum Gravity*, 29(6), 065016.
- Golob, Thomas F. (2003). Structural Equation Modeling for Travel Behaviour Research. *Transportation Research Part B: Methodological*, 37(1), 1–25.
- Gowers, Timothy W. The Importance of Mathematics
- Greenland, Sander and Babette Brumback (2002). An Overview of Relations Among Causal Modelling Methods. *International Journal of Epidemiology*, 31(5), 1030–1037.
- Greenland, Sander, Judea Pearl, and James M. Robbins (1999). Causal Diagrams for Epidemiological Research. *Epidemiology*, 10(1), 37–48.
- Hamming, Richard W. (1980). The Unreasonable Effectiveness of Mathematics. *American Mathematical Monthly*, 87(2), 81–90.
- Hardy, Godfrey H. (1969). *A Mathematician's Apology*. Cambridge University Press.
- Hawley, Katherine (2016). Applied Metaphysics. In K. Lippert-Rasmussen, K. Brownlee, and D. Coady (Eds.), *A Companion to Applied Philosophy* (165–179). Wiley.
- Hitchcock, Christopher and Joshua Knobe (2009). Cause and Norm. *Journal of Philosophy*, 106(11), 587–612.
- Huemer, Michael (2007). Epistemic Possibility. *Synthese*, 156(1), 119–142.
- Jeshion, Robin (2000). On the Obvious. *Philosophy and Phenomenological Research*, 60(2), 333–355.

- Jones, Vaughan F. R. (1985). A Polynomial Invariant for Knots via Von Neumann Algebras. *Bulletin of the American Mathematical Society*, 12(1), 103–111.
- Kluve, Jochen (2004). On the Role of Counterfactuals in Inferring Causal Effects of Treatments. *Foundations of Science*, 9(1), 65–101.
- Knobe, Joshua (2009). Folk Judgements of Causation. *Studies in the History and Philosophy of Science*, 40(2), 238–242.
- Knobe, Joshua and Ben Fraser (2008). Causal Judgement and Moral Judgement: Two Experiments. In Walter Sinnott-Armstrong (Ed.), *Moral Psychology* (441–448). MIT Press.
- Ladyman, James and Don Ross (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press.
- Lange, Marc (2013). What Makes a Scientific Explanation Distinctively Mathematical? *British Journal for the Philosophy of Science*, 64(3), 485–511.
- Lange, Marc (2016). *Because without Cause*. Oxford University Press.
- Lewis, David (1973). Causation. *Journal of Philosophy*, 70(17), 556–567.
- Lewis, David (1991). *Parts of Classes*. Wiley-Blackwell.
- Lewis, David K. (1986). *On the Plurality of Worlds*. Blackwell.
- Lombrozo, Tania (2010). Causal-explanatory Pluralism: How Intentions, Functions and Mechanisms Influence Causal Ascriptions. *Cognitive Psychology*, 61(4), 303–332.
- Love, Peter E. D., David J. Edwards, and Jim Smith (2016). Rework Causation: Emergent Theoretical Insights and Implications for Research. *Journal of Construction Engineering and Management*, 142(6), 04016010.
- Lowe, Edward Jonathan (1998). *The Possibility of Metaphysics: Substance, Identity, and Time*. Clarendon Press.
- Maclaurin, James and Heather Dyke (2012). What is Analytic Metaphysics Good For? *Australasian Journal of Philosophy*, 90(2), 291–306.
- Maddy, Penelope (1988a). Believing the Axioms I. *The Journal of Symbolic Logic*, 53(2), 481–511.
- Maddy, Penelope (1988b). Believing the Axioms II. *The Journal of Symbolic Logic*, 53(3), 736–764.
- Maddy, Penelope (1992). Indispensability and Practice. *Journal of Philosophy*, 89(6), 275–289.
- Maddy, Penelope (1997). *Naturalism in Mathematics*. Oxford University Press.

- Maddy, Penelope (2000). Naturalism and the A Priori. In Paul Boghossian and Christopher Peacocke (Eds.), *New Essays on the A Priori* (92–116). Oxford University Press.
- Mitra, Subhro (2016). Analysis of Truck Accidents Using Structural Equation Models. *Transportation Journal*, 55(4), 382–399.
- Ney, Alyssa (2012). Neo-positivist Metaphysics. *Philosophical Studies*, 160(1), 53–78.
- Norton, John D. (2004). How Hume and Mach Helped Einstein Find Special Relativity. In Mary Domski and Michael Dickson (Eds.), *Discourse on a New Method: Reinvigorating the Marriage of History and Philosophy of Science* (359–386). Open Court.
- Oakes, John Michael (2004). The (Mis)estimation of Neighborhood Effects: Causal Inference for a Practicable Social Epidemiology. *Social Science and Medicine*, 58(10), 1929–1952.
- Parfit, Derek (1984). *Reasons and Persons*. Oxford University Press.
- Paul, Laurie A. (2012). Metaphysics as Modeling: The Handmaiden’s Tale. *Philosophical Studies*, 160(1), 1–29.
- Pearl, Judea (1995). Causation, Action, and Counterfactuals. In A. Gammerman (Ed.), *Computational Learning and Probabilistic Reasoning* (235–255). Wiley.
- Pearl, Judea (1997). Graphical Models for Probabilistic and Causal Reasoning. In A. B. Tucker Jr. (Ed.), *The Computer Science and Engineering Handbook* (697–714). CRC Press.
- Pearl, Judea (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- Price, Huw (2009). Metaphysics After Carnap: The Ghost Who Walks? In David Chalmers, David Manley, and Ryan Wasserman (Eds.), *Metametaphysics: New Essays on the Foundations of Ontology* (320–346). Clarendon Press.
- Prior, Arthur (1957). *Time and Modality*. Oxford University Press.
- Rivest, Ronald L., Adi Shamir, and Leonard Adelman (1978). A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. *Communications of the ACM*, 21(2), 120–126.
- Romero, Gustavo E. (2017). On the Ontology of Spacetime: Substantivalism, Relationism, Eternalism, and Emergence. *Foundations of Science*, 22(1), 141–159.
- Rosen, Gideon (2010). Metaphysical Dependence: Grounding and Reduction. In B. Hale and A. Hoffman (Eds.), *Modality: Metaphysics, Logic and Epistemology* (109–135). Oxford University Press.
- Russell, Bertrand (1912). *The Problems of Philosophy*. Oxford University Press.

- Saatsi, Juha (2011). The Enhanced Indispensability Argument: Representational vs. Explanatory Role of Mathematics in Science. *British Journal for the Philosophy of Science*, 62(1), 143–154.
- Schaffer, Jonathan (2009). On What Grounds What. In David Manley, David J. Chalmers, and Ryan Wasserman (Eds.), *Metametaphysics* (347–383). Oxford University Press.
- Schaffer, Jonathan (2016). Grounding in the Image of Causation. *Philosophical Studies*, 173(1), 49–100.
- Shapiro, Stewart (2009). We Hold These Truths to Be Self-Evident: But What Do We Mean By That? *The Review of Symbolic Logic*, 2(1), 175–207.
- Slavov, Matias (2016). Empiricism and Relationism Intertwined: Hume and Einstein’s Special Theory of Relativity. *Theoria*, 31(2), 247–263.
- Stephan, Simon and Michael R. Waldmann (2018). Preemption in Singular Causation Judgments: A Computational Model. *Topics in Cognitive Science*, 10(1), 242–257.
- Theiner, Georg, Colin Allen, and Robert L. Goldstone (2010). Recognizing Group Cognition. *Cognitive Systems Research*, 11(4), 378–395.
- Turner, Jason (2016). *The Facts in Logical Space*. Oxford University Press.
- VanderWeele, Tyler J. (2016). Commentary: On Causes, Causal Inference and Potential Outcomes. *International Journal of Epidemiology*, 45(6), 1809–1816.
- Wigner, Eugene P. (1960). The Unreasonable Effectiveness of Mathematics in the Natural Sciences. *Communications on Pure and Applied Mathematics*, 13(1), 1–14.
- Witten, Edward (1988). 2 + 1 Dimensional Gravity as an Exactly Soluble System. *Nuclear Physics B*, 311(1), 46–78.
- Witten, Edward (2012). Quantum Mechanics of Black Holes. *Science*, 337(6094), 538–539.