of logical notions is to justify our belief in the consistency of the models provided in model theory for other theories. But the fictionalist has various strategies for defending such a belief which, while not conclusive, at least put the fictionalist's claim to knowledge of the consistency of our mathematical theories on a better footing than the platonist's claim to knowledge that those theories are true. The fictionalist can thus account for all the logical knowledge she claims to have.

Mathematical recreation versus mathematical knowledge

MARK COLYVAN

1 Empiricism in the philosophy of mathematics

Empiricism in the philosophy of mathematics has a chequered history. Mill defended a version of empiricism according to which the laws of arithmetic were highly general laws of nature. Mathematical truths such as $2+3=5$ were thought by Mill to be empirical in that they tell us that if we were to take two logic books, say, and three ethics books, say, we'd have five philosophy books. But Mill's somewhat naive empiricism found itself on the receiving end of a stinging attack from Frege. This attack, I might add, was considered by many to be decisive. Frege had many complaints, but the most significant was that Mill had confused applications of arithmetic with arithmetic itself.[1]

But empiricism about mathematics arose again in a more subtle form in the work of W. V. Quine. According to Quine's version of empiricism, mathematics is empirical in the sense that the truth of mathematics is confirmed by its applications in empirical science. More precisely, Quine argues that when we empirically confirm a scientific theory, we empirically confirm the whole theory, including whatever mathematics is used. Quine is not vulnerable to Frege's attack on Mill because Quine is not confusing mathematics with its applications. Rather, Quine is invoking the applications as a reason for taking the mathematics to be true.[2] Moreover, according to this Quinean picture, mathematics is taken at face value—it's about mathematical entities such as numbers, functions, sets and the like[3]—and these entities are taken to exist because of the indispensable role they play in our best scientific theories. This argument has become known as *the indispensability argument*.

I won't defend this Quinean indispensability argument here. I've done that elsewhere (Colyvan 2001). Instead, I will highlight some of the attractive features

[1] See Mill (1947) for Mill's empiricist philosophy of mathematics, Frege (1974: §23) for Frege's attack, and Dummett (1991) for a good discussion of this exchange.

[2] See Putnam (1971) and Quine (1981*a*) for articulations and defences of this view. Interestingly, this view can be traced back to Frege (1970: 187).

[3] Although Quine's Ockhamist tendencies drive him to the view that only sets are really needed, so that's all we are ultimately committed to. There are substantial issues here though. Do we reduce the natural numbers, say, to particular sets, or will any set-theoretic construction of the natural numbers do just as well? In either case (though especially the latter) can we still be said to be employing the standard semantics? It might be argued that reductions of mathematics to set theory involve some reinterpretation of mathematical language. Thanks to Mary Leng for this point.

of the kind of empiricism that emerges from it. In particular, I'll discuss how the Benacerraf epistemological problem for mathematical realism does not have any purchase on this empiricist mathematical realism. I'll then consider, in some detail, one feature of this view that has recently come under attack.

2 *An empiricist account of mathematical knowledge*

In 1973, in a now famous paper, 'Mathematical Truth', Paul Benacerraf put voice to an epistemological concern about mathematical realism that had no doubt been around for a very long time. The concern is quite simple. If mathematical entities exist but lack causal powers, it is inexplicable how we could come to know about them. Benacerraf explicitly invoked the causal theory of knowledge as a major premise in the argument but this epistemology fell out of favour not long after the publication of Benacerraf's paper.[4] But still there is something seductive about this argument. W. D. Hart puts the point thus:

> It is a crime against the intellect to try to mask the problem of naturalizing the epistemology of mathematics with philosophical razzle-dazzle. Superficial worries about the intellectual hygiene of causal theories of knowledge are irrelevant to and misleading from this problem, for the problem is not so much about causality as about the very possibility of natural knowledge of abstract objects. (Hart 1977: 125–6)

But what then is the worry about abstract objects? What is it about abstract objects that suggests that it's impossible to have knowledge about them? In my view, the most cogent post-causal-theory-of-knowledge version of this argument is due to Hartry Field. He captures the essence of the Benacerraf argument when he puts the point in terms of explaining the reliability of mathematical beliefs:

> Benacerraf's challenge—or at least, the challenge which his paper suggests to me—is to provide an account of the mechanisms that explain how our beliefs about these remote entities can so well reflect the facts about them. The idea is that *if it appears in principle impossible to explain this*, then that tends to *undermine* the belief in mathematical entities, *despite* whatever reasons we might have for believing in them (Field 1989: 26, emphasis in the original).

This challenge is usually understood to be to account for the reliability of the inference from 'mathematicians believe that P' (where P is some proposition about some mathematical object(s) to 'P', while explicitly detailing the role that the mathematical entities in question play in this reliable process. But stated thus, we see that a substantial question is being begged against Quine and other epistemic holists. Epistemic holists hold that we do not justify beliefs one at a time. Rather, we justify packages of beliefs. How large that package is depends on how radical is the holism. Quine wavered a little on this issue, at times suggesting that it was the whole system of beliefs that was justified, at other times, he more reasonably allowed for (largish) proper subsets of our beliefs to be justified.

[4] Mark Steiner (1975) was one who took issue with the causal theory of knowledge.

So from the epistemic holist point of view, this interpretation of the Benacerraf–Field challenge is simply question begging—it assumes that mathematical beliefs are justified one at a time.

To avoid begging questions, let's take a more charitable reading of the Field version of the epistemic challenge, according to which the challenge is to explain the reliability of our systems of beliefs and to explicitly articulate the role the world plays in this reliable process. Note that we can't ask after the roles of individual objects any longer. Since we are interested in the justification of whole systems of beliefs, the best we can do is to ask after the joint roles played by collections of objects in reliable belief acquisition. In some cases this collection might be so large as to include the whole world (including, of course, whatever abstract objects there are).

Once the challenge is put this way, we see that Quine has already answered it: we justify our *system of beliefs* by testing it against *bodies of empirical evidence*. No distinction is made between mathematical beliefs and other beliefs. Our beliefs form a package that performs well against the usual standards of theory choice and that's all that matters. Any challenge to provide an account of only the mathematical beliefs is illegitimate. According to the holist, mathematical beliefs are justified in exactly the same way as other beliefs: by their role in our best scientific theories and these, in turn, are justified by appeal to the usual criteria of theory choice (empirical adequacy, simplicity, explanatory power, and so on).[5]

Still, it might be thought that all this talk of holism is beside the point. The point is that mathematical objects don't seem to contribute to the success of theories in the same way as, say, electrons, and *this* is what is in need of explanation. But again this way of stating the problem requires singling out the mathematical entities and asking after them. The thorough-going holist would deny that it makes any sense to do this, and they would thus reject the assumption underwriting the challenge. This is not dodging the issue or introducing philosophical 'razzle-dazzle'. Holism cannot be bracketed for the purposes of getting an objection up against the holist. There may be something to the epistemological challenge, but until it is formulated in a way that has some bite against the Quinean empiricist, I'm inclined to suggest that the burden lies squarely with those who believe there to be a problem here for the Quinean.

Finally, a couple of points of clarification. Quine is interested in *justification* and the Benacerraf–Field challenge is seeking an explanation of the *reliability* of the belief-forming process. Have I missed the point by focusing on Quinean holistic justification? I don't think so. In the current context, at least, we are assuming that the methods of current science are reliable methods for forming beliefs. Indeed, we are justified in believing our scientific theories, in large part because we believe that these theories were arrived at by reliable methods. Of course this assumption of reliability might be mistaken, but to push in that direction would be to mount a more general sceptical challenge, and I take it that the Benacerraf–

[5] This line of thought is advanced in Rosen (1992: ch. 3) and Colyvan (2006).

Field challenge is supposed to be a particular challenge for mathematical knowledge. Moreover, it should be noted that such general sceptical challenges typically ask us to step outside our best scientific theories and explain the reliability of (or otherwise justify) those theories from some external vantage point. This, according to the Quinean naturalist, is not possible, so any general challenge thus formulated is again question begging (but this time against naturalism). If the request for justification is in terms of intra-scientific justification, then, as Quine has argued in various places (e.g. in (1974: 3)), it is a challenge that can be met by invoking our best science (theories of optics, psychology and so on). We may be stuck in Neurath's boat, but there are some powerful tools available to us on that boat.

If we accept all this, and we admit that the Quinean epistemic holist has a good reply to the Benacerraf–Field epistemic challenge, a serious issue arises about those portions of mathematics left unapplied. After all, on the view under consideration, it's only the mathematics that finds itself indispensable to our best scientific theories that is justified. The rest, it would seem, must have a different status. I address the issue of the status of unapplied mathematics in the next section.

3 *Unapplied mathematics as mathematical recreation*

Unapplied mathematics is something of a nuisance for Quine. It can't be justified by the same means as applied mathematics, since it's precisely the applications that provide the justification. Moreover, Quine's empiricism won't allow other (non-empirical) means of justification, so it seems that unapplied mathematics does not have the same status as applied mathematics.[6] Applied mathematics is treated realistically—its propositions are believed to be true and the objects quantified over are treated as real—while unapplied mathematics, it would seem, must be (at best) treated agnostically. Charles Parsons (1983) pushes precisely this point and in reply Quine argues that it is reasonable to treat realistically a bit more than the mathematics that does in fact find itself indispensable in applications. We should include whatever mathematics is required for 'rounding out' that which is applied. The latter includes a great deal of set theory, since set theory is usually thought to underwrite most contemporary mathematics, both applied and unapplied. But how much set theory enjoys the exalted position of 'justified'? And what is the status of the rest? In response to the first question, Quine's Ockhamist sympathies come to the fore and he draws the line at the constructible sets: $V=L$. According to Quine, the demand of simplificatory rounding out of applied mathematics may be thought to extend only so far as the constructible sets. As for the second question, Quine bites the bullet Parsons offers and admits that

[6]I'm using 'applied' (and 'unapplied') here in the intuitive sense. In the mouths of mathematicians, 'applied mathematics' corresponds (roughly) to numerical methods (as opposed to pure, analytic methods).

'magnitudes in excess of such demands, e.g. $\beth_\omega$ or inaccessible numbers' should be looked upon as 'mathematical recreation and without ontological rights (1986: 400).[7]

Although I want to defend the distinction between that mathematics which we treat realistically and recreational mathematics, I will part company with Quine on a couple of issues here. First, I think that Quine makes it sound as though there are two quite different kinds of justification at work here. Lower mathematics (the lower reaches of set theory, analysis, and the like) is justified by the indispensable role it plays in our best scientific theories; the upper reaches of constructible set theory (transfinite arithmetic and so on) is justified by quite different means. The latter is justified by something akin to an act of charity: it is justified by simply being close enough to the mathematics that is applied. That is, according to Quine, accepting the upper reaches of constructible set theory is the most natural and simple way to round out the mathematics that is applied. But I think we can do better than this. I note that indispensability is transitive. If a nail gun is indispensable to building houses and building houses is indispensable to building suburbs, then a nail gun is indispensable to building suburbs. Similarly for mathematics. If transfinite set theory is indispensable for analysis and analysis is indispensable for physics, then I say transfinite set theory is indispensable for physics. Perhaps this is what Quine had in mind with his notion of 'simplificatory rounding out'. In any case, this is the justification for the higher reaches of set theory that I endorse. Understood this way, there is only one mode of justification: playing an indispensable role (either directly or indirectly) in our best scientific theories.

But is indispensability really transitive? Gideon Rosen (private communication) has questioned this claim. Rosen suggests that although large cardinals, say, might be indispensable for our best theory of real numbers, and real numbers might be indispensable for our best theories of space–time, it need not follow that large cardinals are indispensable for the physics of space–time. Physicists might

[7]Quine later refined his position on the higher reaches of set theory and other parts of mathematics not applicable to natural science:

> They are couched in the same vocabulary and grammar as applicable mathematics, so we cannot simply dismiss them as gibberish, unless by imposing an absurdly awkward gerrymandering of our grammar. Tolerating them, then, we are faced with the question of their truth or falsehood. Many of these sentences can be dealt with by the laws that hold for applicable mathematics. Cases arise, however (notably the axiom of choice and the continuum hypothesis), that are demonstrably independent of prior theory. It seems natural at this point to follow the same maxim that natural scientists habitually follow in framing new hypotheses, namely, simplicity: economy of structure and ontology (Quine 1995: 56).

And after considering the possibility of declaring such sentences meaningful but truthvalueless, he suggests:

> I see nothing for it but to make our peace with this situation. We may simply concede that every statement in our language is true or false, but recognize that in these cases the choice between truth and falsity is indifferent both to our working conceptual apparatus and to nature as reflected in observation categoricals (Quine 1995: 57).

look for different things in their theories than does the mathematician. Similarly, housing developers may look for different things in a suburb than the carpenter building the houses. While houses are indispensable to suburbs, houses built with nail guns may not be. The carpenter might be interested in strength of construction while the developer is interested in speed of construction, for example. In response, I suggest that there is an equivocation here involving the word 'best'. While it seems right that the best suburb (in the developer's sense of 'best') need not be built from the best houses (in the carpenter's sense of 'best'). But if we insist on the same sense of 'best' throughout Rosen's concern is laid to rest and transitivity is restored. The question is whether, in the scientific examples at issue, we can insist on the same sense of 'best'. Rosen seems to be on firm ground here, for surely set theorists and physicists look for quite different virtues in their best theories. Indeed, Penelope Maddy (1997) argues convincingly that set theorists do not seem to value parsimony as a virtue at all. Set theorists want as many different structures as possible. Physicists, on the other hand do seem to value parsimony.

In response, I suggest that issues concerning disciplinary expertise save the transitivity of indispensability (in the scientific context, at least). Physicists might value parsimony in their physical theories but when it comes to deciding what the best theory of the real numbers is, that's a mathematical question and it is decided by mathematical standards. Sure these standards are different from those of the physicist, but it's the mathematicians who decide what the best theory of the reals is. If the mathematicians decide that a large cardinal axiom is indispensable for this theory, then so be it. The physicists do not get to apply their standards here and they do not have the relevant expertise to do so. Now, if the best theory of space-time requires the real numbers, then whether the physicists like it or not, large cardinals are indispensable to real number theory (or so we are assuming for the point of the example) and so large cardinals are also indispensable to theories of space-time. So even if different theoretical virtues are respected in different parts of science, issues concerning disciplinary expertise ensure the transitivity of indispensability.[8]

A couple of points to note in relation to this somewhat more liberal understanding of indispensability. It may turn out that very little, if any, mathematics is unapplied on this account. After all, on this account, for a branch of mathematics to be unapplied, it must be totally isolated from the main body of mathematical theory; it must not find applications in any chain of applications that bottoms out with applications in empirical science.[9] Also, on this account, it is not so clear that one can draw the line at constructible set theory. The debates in set theory over the various large cardinal axioms, for instance, seem to be about the most natural way to extend ZFC so as to have pleasing and intuitive consequences for

[8]And if you accept Quine's view that all mathematics is set theory, then the transitivity of indispensability might be established more directly: real numbers are sets, so set theory is indispensable to space time theory.

[9]Or at least, it must be dispensable to the main body of mathematical theory. More on this shortly.

lower set theory and higher set theory alike. So even the most abstract reaches of set theory may yet turn out to be applicable, in this extended sense of applicable.[10] At the very least, the view of applications I'm endorsing here does not ignore the higher reaches of set theory (as Penelope Maddy (1992) once complained of Quine's philosophy of mathematics). Finally, I note that what is indispensable now may be dispensable tomorrow. Just as nail guns replaced hammers in building houses, we might find replacements for some of the mathematics that we now think of as indispensable.

The second way I depart from Quine on the issues under consideration follows from this. As I've already noted, it is Quine's Ockhamist sympathies that compel him to keep his ontological commitments to a minimum. While I too have such sympathies in some areas of metaphysics, it's not clear that Ockhamist considerations are appropriate here. After all, Quine is already committed to a very large infinity of abstract objects, so why baulk at a few more?[11] Of course the very nature of the Quinean argument invoked to establish the existence of mathematical entities restricts discussion to those mathematical entities that are indispensable for science. The real issue then is how much set theory is needed for science. (After all, Ockham's razor implores us not to multiply entities beyond *necessity*.) If what I suggested in the previous paragraph is correct, much more than constructible set theory is necessary, so even Ockhamists like Quine can countenance more than just the constructible sets. For the record, my position on this is to side with the majority of set theorists and accept that set theory really does need more than the constructible sets. I thus reject $V=L$. How much more? I take it that the jury is still out on that issue. But I certainly don't have in-principle objections to set theory extended by some a large cardinal axiom such as MC ('there exists a measurable cardinal').[12]

These may seem like major departures from Quine's position, but I think not. On the first issue, the way I justify the higher reaches of set theory is only superficially different from Quine's, if at all. Although Quine never emphasized the chains of applications, this may well be what he had in mind when he suggested justifying set theory up to $V=L$. On the second issue, the disagreement is more substantial. Quine is very restrictive about how much set theory we should treat realistically. I, on the other hand, am endorsing a much less restrictive attitude. But even this difference is not as significant as it might at first seem. I take it that there's nothing in the core Quinean doctrines that drives him to accept $V=L$. He needs to draw the line between applied and unapplied mathematics in a neat and

[10]Though see Feferman (1992) for a defence of the view that not very much mathematics is required for empirical science.

[11]See Burgess and Rosen (1997) for a very nice discussion of Ockham's razor in the context of mathematics.

[12]Also, I do not share Quine's view that a bivalent logic (presumably classical first-order logic) applying to all sentences of the language is simpler than some of the alternatives. (See n. 7 above, second quotation from Quine.) I will resist the temptation to take up this interesting issue here, since it's somewhat tangential to my main purpose.

convenient way (and, as I've already noted, in accord with his Ockhamist sympathies). I too have to draw the line somewhere; it's just unclear to me where that somewhere is, and I'm inclined to draw it a little further along than Quine.

Let's now return to the points on which Quine and I agree. We both accept that mathematics is justified by the indispensable role it plays in our best physical theories. We both accept that such justification does not extend to all contemporary mathematics. At least we both agree that it is conceivable that some portions of contemporary mathematics are without this kind of justification. As I've already noted, Quine takes a fairly hard line with regard to such mathematics and gives it the status of mathematical recreation. And on this too we agree. But it is important to note that in calling it 'mathematical recreation' Quine is not dismissing it. Mathematical recreation remains an important part of mathematical practice. It should not be thought of as mathematicians merely having a good time and engaging in a pastime quite distinct from their normal practice. Mathematicians engaged in mathematical recreation are much like theoretical physicists exploring different possible physical theories. Physicists, for instance, studying the Schwarzchild solution to Einstein's equation or Newtonian celestial mechanics might be thought to be engaged in 'recreational physics'. They are most certainly not studying anything real—we simply do not live in a Schwarzchild or a Newtonian universe. Nor are these physicists just having a good time and leaving behind standard practice. Investigating such non-actual solutions is an important part of standard scientific practice.

What is the point of engaging in recreational physics and recreational mathematics? There are many reasons for pursuing such activity. By studying the non-actual, we often come to a better understanding of the actual (by, for instance, coming to a better understanding of the underlying laws). We might be deliberately making simplifying assumptions because the actual situation is too complicated. We might not be sure of what is actual and so taking a pluralistic attitude means that all bases are covered, so to speak. Or it might be simply intellectual curiosity. The bottom line is that mathematical recreation, like other forms of theoretical scientific investigation, should not be thought of as second class or *mere* recreation.[13]

To sum up my position. I accept that there is a distinction between unapplied mathematics and applied mathematics—even given my very liberal sense of application via chains. I accept that applied mathematics should be treated realistically and with unapplied mathematics we have no reason to treat it this way. Unapplied mathematics is akin to theoretical investigations elsewhere in science and, as such, is an important part of mathematical practice. It is also important to note that while many branches of mathematics are at least initially pursued as recreational, they nonetheless end up being applied. Mathematics can thus change

[13] Indeed, the phrase 'mathematical recreation' is a little unfortunate. Perhaps 'theoretical' or 'speculative mathematics' would have been better, but the phrase 'mathematical recreation' is already in the literature, so I'll stick with it.

its status with regard to the recreational–non-recreational divide. While there remains something of a mystery as to how mathematics pursued by apparently a priori means and without regard to applications can end up being applied,[14] there is no doubt that this happens. On the empirical account of mathematics I'm defending here, applications make all the difference. Once a branch of mathematics finds an application, it should be treated realistically.

4 *Is all mathematics recreation?*

Mary Leng has argued that allowing some mathematics to be treated as recreation and without ontological rights, leads to a slippery slope to all mathematics being recreational. First she voices a general worry to soften us up for the argument to follow. She notes that there's nothing in mathematical practice that distinguishes between recreational mathematics and literally true mathematics. I agree. Leng then goes on to suggest that:

> Considered in this light, Colyvan's distinction between literally true mathematics and merely recreational mathematics begins to look like a distinction without a difference. The literal truth of a mathematical theory will make no difference to how a mathematician goes about working in that theory (2002: 408).

But just because there's no distinction to be found in mathematical practice, does not mean that this is 'a distinction without a difference' (Leng 2002: 408). As I've already pointed out, I take it that there's no methodological difference that cleanly marks the boundary of recreational physics from other parts of theoretical physics. Leng is right that this is not a methodological distinction, but that does not mean that it's not a distinction at all. Still her main point is correct: mathematical methodology does not recognize the recreational–non-recreational distinction. That distinction is extraneous to mathematical practice. It is determined by which parts of mathematics find indispensable applications in physical science. Applied mathematics is where the action is. On that Leng and I agree.

Now to Leng's main argument. Crucial to her argument is what Leng calls 'the modelling picture' of mathematical applications.[15] According to this picture, mathematics is never assumed to be literally true in any applications; it is judged to be adequate or inadequate for a particular application and that's the end of it. The role of the mathematics is to represent particular features of the physical system under investigation and it may do this well or poorly. According to the modelling picture, mathematics can perform this representational function irrespective of the truth of the mathematics in question. Indeed, on the modelling view, mathematics is not confirmed or disconfirmed at all. At best, the adequacy

[14] This feature of mathematical practice is often referred to as 'the unreasonable effectiveness of mathematics' (Wigner 1960).

[15] A similar account of mathematics in applications has been advanced by Christopher Pincock (2004*b*; 2004*a*) and criticized in Bueno and Colyvan (n.d.).

of the representation is confirmed. So, for example, when we find space-time is not Euclidean, we do not claim to have disconfirmed Euclidean geometry *as a mathematical theory*. Rather, we claim to have disconfirmed the adequacy of Euclidean geometry as a suitable representation of space-time, and the latter is quite different.[16]

Using this modelling picture of mathematics in applications, Leng then goes on to argue that all mathematics is recreational.

> If Colyvan is right (and I think he is) that mathematics that is not assumed by science to be true should be seen as recreational (and given some important status as such), then it follows from the modelling picture of the relationship between mathematics and science that all mathematics is recreational. (2002: 412)

This argument can be spelled out thus:

Premise 1. Empiricism holds that mathematics with no empirical confirmation should be viewed as merely recreational.
Premise 2. The modelling view of applications has it that when we use mathematics to represent (or model) non-mathematical phenomena, all that is confirmed is that the mathematics allows for a good representation, not that it is true.
Premise 3. The modelling view accounts for all applications of mathematics.
Conclusion: Therefore, all mathematics is recreational.

This is a very interesting argument. Although I will argue that it is ultimately flawed, I think Leng's argument raises important issues that cut right to the heart of the indispensability argument and the subsequent debate. As I've already indicated, I accept premise 1 (with the earlier provisos about how the more remote reaches of mathematics might be confirmed indirectly). Here I'll take issue with premises 2 and 3.

First, let's look at premise 2 and, in particular, the important role played by Sober's (1993*a*) argument against mathematics accruing confirmational support. Sober argues that the truth of mathematics is never placed on the line—if a mathematicized physical theory such as Newtonian mechanics turns out to conflict with experience, then the mathematics employed is never thought to be shown to be false. At worst, the mathematics is simply thought to be an inappropriate way to represent the theory in question. But this, Sober suggests, shows that mathematics is not really being empirically tested at all. So, in particular, there is no reason to think that mathematics employed in a successful empirical theory enjoys whatever confirmational support the theory accrues.

In Sober's paper (1993*a*), the argument is cast in terms of the contrastive empiricist theory of confirmation.[17] Sober goes on to argue that the main point

[16]Here Leng invokes Elliott Sober's (1993*a*) criticisms of the indispensability argument and the holistic picture of confirmation. More on this shortly.

[17]This theory compares likelihoods, $\Pr(E|H_1)$ and $\Pr(E|H_2)$, of two competing hypotheses H_1 and H_2 in the light of some evidence E. Contrastive empiricism suggests that H_1 receives greater confirmational support from the evidence E if $\Pr(E|H_1) > \Pr(E|H_2)$.

against the indispensability argument is independent of this particular theory of confirmation, and I take it that this is why Leng doesn't address the issue of the plausibility of contrastive empiricism. Be that is it may, Sober's argument is not independent of separatist confirmation theory. That is, he assumes that we can confirm or disconfirm hypotheses one at a time. But as we've already seen, this is a point that Quine denies. So Sober's objection is question begging.[18] Indeed, this can be seen from the fact that other, clearly empirical hypotheses, are never called into question when a theory confronts recalcitrant data. As Michael Resnik points out (1997: 168) various conservation laws seem immune from revision and yet it is unreasonable to deny empirical content to such principles. What's going on here is that some parts of the theory (such as mathematical principles and conservational laws) play a rather central structuring role in the scientific theories in which they appear. Sober is right that they rarely get called into question when the theory encounters recalcitrant data. But this is because, according to the holist, at least, the theory as a whole is untenable. But it is a mistake to conclude from this that every part must share the blame equally. Typically, when a theory conflicts with evidence it is only a small part of the theory that needs to be revised. There's still a substantial issue as to why it is never (or at least *almost* never) mathematics that is revised.[19] But the fact that mathematics rarely takes the fall is no reason to conclude that mathematics should not take at least some of the credit in successful theories.

Now to Premise 3. Crucial to Leng's main argument is the assumption that the *only* role mathematics plays in science is representational (hence the 'modelling picture' of mathematical applications). The central idea of this view of the mathematics–science relationship is that we have some physical system such as a population of organisms, we represent the number of organisms by a mathematical function such as the logistic function—or more commonly, we describe some features of the function in question by stating the appropriate differential equation).[20] If we then notice that the mathematics produces the wrong answers, we say that the mathematics in question was not appropriate. We do not reject the theory of differential equations, say. On this account of the relationship between mathematics and science, mathematics provides nothing more than a convenient set of representational tools. But such an account seems to seriously understate the role of mathematics in science. I've argued elsewhere (2001: ch. 3) that mathematics may contribute directly to explanations in science. If this is right, then

[18]Of course, both Sober and Leng are attacking confirmational holism so it's also question begging for the Quinean to invoke confirmational holism as a response to them. At best they have achieved an unsatisfying stand off. I also think that there are other problems with Sober's argument, but since I've dealt with these problems elsewhere (Colyvan 2001: ch. 6), I won't revisit them here.

[19]See Resnik (1997: ch. 7) for more on why this should be so.

[20]The logistic equation, for instance, is usually represented as a first-order differential equation: $dN/dt = rN(1 - N/K)$, where N is the population abundance, t is time, r is the growth rate, and K is the carrying capacity.

mathematics is more than a *mere* representational tool and the modelling picture is wrong. After all, if mathematics is contributing directly to explanations, it is hard to see how anyone who accepts inference to the best explanation can accept the explanations yet deny the truth of the mathematics.

I now present a few examples where I take it that the mathematics in question is doing more than merely representing; it is also explaining. These examples, thus undermine the plausibility of premise 3 of Leng's argument.

Example 1. Consider the ancient problem of squaring the circle: using only compass and straight-edge, construct a square with the same area as a given circle. Here we can represent the various physical activities (marking off lengths with the compass and drawing lines with the straight-edge) mathematically. Leng is right in suggesting that the mathematics is modelling the physical activities. But she is wrong in suggesting that that's all the mathematics does. For the construction in question, as we now know, is impossible, and the explanation of why it is impossible is that π is transcendental. The mathematics, it would seem, is not only modelling but also *explaining* the impossibility of certain physical activities.

Example 2. A mountaineer sets out at 6.00 am from base camp with a load of supplies and arrives at the top camp later that same day. The following day the mountaineer returns to base camp, again leaving at 6.00 am. It turns out that there will be a point on the mountain that the mountaineer will pass at the same time on both days. Why should there be such a point on the mountain? If we represent the physical situation in the obvious mathematical way, a fixed point theorem then guarantees that there will be such a point on the mountain. Again it seems that the mathematics in question is doing more than merely modelling; it is also explaining the existence of a physical event, namely, the location of the mountaineer at the same height on the mountain at the same time on the two days in question. This case is interesting because although the fixed-point theorem guarantees that there will be some such point on the mountain, it doesn't explain why it's any particular point. An explanation of *that* fact will presumably proceed via a detailed causal story of the mountaineer's slog up and down the mountain. But this does not change the fact that the explanation of why there should be any such point is a topological explanation.

Example 3. The previous example, in particular, is a little artificial, so let me provide a real example of a phenomenon that scientists feel is in need of explanation. Evolutionary biologists are puzzled by the presence of apparently maladaptive traits, such as altruism. As Elliott Sober (1993*b*: 98–102) points out, altruistic individuals are less fit than non-altruists in a given population, so we would expect natural selection to force a decline in the relative frequency of altruism. But altruism is alive and well. How can this be? One crucial piece of the explanatory story may well be purely mathematical in nature, and involves nothing other than simple facts about inequalities, addition and division. It is common sense that if a trait is declining in relative frequency in every group, then it is declining in relative frequency overall. But for all its intuitive plausibility, this piece of reasoning is fallacious. Simpson's paradox (Malinas and Bigelow 2004) shows how a trait can be less fit relative to each of a number of groups, yet fitter relative to the ensemble of groups.[21] To take a simple example, suppose that there are two groups, A and B. In group A altruists outnumber non-altruists by 200:50. In group B there are 100 of each. After selection we find that in group A there are 150 altruists and 45 non-altruists, and in group B, there are 15 altruists and 20 non-altruists. So the fitnesses of altruists in groups A and B are 0.75 and 0.15 respectively. The fitnesses of the non-altruists are 0.9 and 0.2 respectively. As you would expect, in each group, the non-altruists are fitter. But look what happens in the combined population $A+B$. Here the fitness of the altruists is 0.55, whereas the fitness of the non-altruists is 0.43. The explanation for this peculiarity is simple and it's mathematical: although $a/b > c/d$ and $e/f > g/h$ it does not follow that $(a+e)/(b+f) > (c+g)/(d+h)$.[22]

As seductive as the modelling picture of the relationship between mathematics and science is, it ignores important aspects of this relationship.[23] To be sure, there are many cases where mathematics is used to represent and nothing more. Leng's example (2002: 411) of population dynamics may be one such case. Indeed, Ginzburg and I have suggested (2004: 31–3) that ecologists quite rightly resist mathematical explanations of ecological facts—they hold out for ecological explanation. Such examples of mathematics in science suit Leng well. But since she is offering a general account of the mathematics–science relationship, she needs to argue that in *all* applications, mathematics merely represents. In particular, she needs to give an account of cases like those above (and others I present in (2001; 2002)) where mathematics contributes to scientific explanation. Until such an account is forthcoming, we have good grounds to reject the modelling picture of the mathematics–science relationship. At least it cannot be the whole story and so premise 3 of Leng's argument should be rejected. And with that premise goes Leng's conclusion that all mathematics is recreational.

An important issue emerges from this debate though. A great deal of the early literature on the realism–anti-realism debate in the philosophy of mathematics focused on the mere fact that mathematics has applications in science. Leng is right to follow Maddy (1997) and others to look more carefully at the details of those applications. But the relationship between science and mathematics is complex and multifaceted. I don't think that the modelling picture does justice to the variety of applications and the complexity of the relationship between science and mathematics, though I'm not offering any account in its place here. I'm

[21] For present purposes, we take fitness of a group to be the ratio of the number in the group before selection to the number after selection.

[22] See Colyvan (2001) and Lyon and Colyvan (2008) for other examples of mathematical explanation and also Alan Baker's (2005) very nice example of some elementary prime number theory explaining facts about Cicada life cycles.

[23] Here I've focused on one aspect of what the modelling picture ignores: explanation. But elsewhere (2001; 2002), I've suggested that mathematics can contribute to other scientific virtues such as unification and even novel predictions. Although, see Melia (2002) for disagreement on the unification claim and Leng (2005*a*) and Leng (2008) for disagreement about the philosophical significance of mathematical explanation.

inclined to think that a great deal more work needs to be done on this issue, with detailed case studies on particular applications.[24] At this stage I'm rather sceptical that any systematic philosophical account of mathematics in applications will be forthcoming. The best we may ever be able to do is understand particular applications on a somewhat ad hoc and case-by-case basis. But this, of course, is mere speculation.

5 *Empiricism revisited*

So with Leng's argument that all mathematics is recreation dispensed with, we are able to maintain the recreational–non-recreational distinction. This distinction is important for the kind of empiricism I'm advocating here. Even though it might turn out that there is not a great deal of recreational mathematics (if any), there must be room for such activity. For otherwise the empiricism is rather vacuous. We do not want mathematics to be justified simply because some mathematicians study the area in question.[25] That would not be empiricism at all.

I have a great deal of sympathy with the idea that mathematics should be justified on purely mathematical grounds. After all, mathematics is the queen of the sciences and as such might be thought to occupy a privileged position, not in need of any further (external) justification. This view, however, leads to problems. Taking the lead from mathematics, practitioners in other areas might seek justification for their beliefs in terms of their own practices. We might find a push to justify religious beliefs because they belong to a system studied by some religious group or other. Or perhaps an attempt to justify beliefs about extraterrestrial abductions because some UFO cult takes such abductions seriously and claims to study them. Clearly mathematics enjoys a higher status and is much more reputable than either religion or alien abduction theory, but what is it that gives mathematics such status? Empiricism gives a clear answer to this question (at least for all mathematics that's applied): it is justified by its direct and indirect applications in empirical science. Indeed, according to this version of empiricism *all* beliefs must ultimately be answerable to empirical evidence. We are thus able to provide a satisfying account of mathematical knowledge, where mathematics is respected, but it earns this respect by the work it does in empirical science. There is no room for free riders or self-indulgent queens. Everyone pays their way in this version of empiricism—even royalty like mathematics.[26]

[24]See Bueno and Colyvan (n.d.) for some tentative steps towards such an account.

[25]Maddy (1992) suggests extending Quinean naturalism to pay due respect to mathematical practice along such lines.

[26]I'd like to thank Alan Baker, Daniel Isaacson, Mary Leng, Aidan Lyon, Gideon Rosen, and Crispin Wright for very helpful conversations on the topic of this chapter, and Mary Leng for her comments on an earlier draft. I'd also like to thank others in the audience at the Mathematical Knowledge Conference held at the University of Cambridge in June–July 2004 for their comments and criticisms. Finally, I'd like to thank the organizers, Mary Leng, Alexander Paseau, and Michael Potter, for putting together such a terrific conference and for inviting me to speak at it. Work on this chapter was funded by an Australian Research Council Discovery Grant (grant number DP0209896).