Structure and Applied Mathematics

Travis McKenna[†]
University of Pittsburgh

Abstract

'Mapping accounts' of applied mathematics hold that the application of mathematics in physical science is best understood in terms of 'mappings' between mathematical structures and physical structures. In this paper, I suggest that mapping accounts rely on the assumption that the mathematics relevant to any application of mathematics in empirical science can be captured in an appropriate mathematical structure. If we are interested in assessing the plausibility of mapping accounts, we must ask ourselves: how plausible is this assumption as a working hypothesis about applied mathematics? In order to do so, we examine the role played by mathematics in the multiscalar modelling of sea ice melting behaviour and examine whether we can indeed capture the mathematics involved in the kind of mathematical structure employed by the mapping account. Along the way, we note that the cases of applied mathematics that appear in discussions of mapping accounts almost exclusively involve the employment of a single clearly circumscribed mathematical field or domain (e.g. 'the use of arithmetic in counting physical objects'). While the core assumption of mapping accounts may appear plausible in such situations, we ultimately suggest that the mapping account is not able to handle the important added complexities involved in our sea ice case study. In particular, the notion of mathematical structure around which such accounts are framed does not seem to be able to capture the way in which some applications of mathematics require that very different pieces of mathematics be related to one another on the basis of *both* mathematical and empirical information.

1 Introduction

In recent years, appeals to various notions of 'structure' have become commonplace in discussions of the application of mathematics in empirical science.

[†] I would like to thank Bob Batterman, Mark Wilson, Erica Shumener, Marco Maggiani, Stephen Mackereth and James Shaw for their helpful comments on various iterations of this paper. I would also especially like to thank two anonymous reviewers for their patient and insightful feedback.

According to 'mapping accounts' of applied mathematics, the application of mathematics is best understood in terms of mappings between physical structures on the one hand and mathematical structures on the other. Proponents of mappings accounts have spent considerable time attending to complexities surrounding, amongst other things, the kind of mappings that such a view ought to countenance and the best way to understand the (often mathematised) 'physical structures' in which empirical science deals.

Very little attention, however, has been paid to the role played by the notion of a 'mathematical structure' in such proposals. While the notions of 'mapping' and 'physical structure' are taken by proponents of such accounts to stand in need of at least some clarification, the notion of a 'mathematical structure' is treated as relatively self-explanatory. Bueno and Colyvan, for instance, simply declare that "a structure is usually taken to be a set of objects (or nodes or positions) and a set of relations on these" (Bueno and Colyvan, 2011, 347) before noting that they will adopt this standard account and discussing the difficulties that might arise in attempting to identify such 'structures' in the physical world (or at least in the realms of physical science).

If we are interested in assessing the capacity of mapping accounts to provide us with a plausible picture of applied mathematics, however, we must pay closer attention to the mathematical side of the ledger. In particular, mapping accounts seem to rely on the assumption that the mathematics relevant to any particular application in empirical science can be captured in a single self-standing 'mathematical structure'. Given that the strategies that applied mathematicians develop for tackling problems in empirical science often involve stitching together a diverse assemblage of tools from disparate areas of mathematics, we must ask ourselves: how plausible is this assumption as a working hypothesis about applied mathematics? Posing and answering this question will be the focus of this paper.

There are, however, two issues we must get clear on before we are in a position to answer this question. First, we must consider what kind of cases of applied mathematics are most relevant to testing the plausibility of this assumption. We will suggest that discussions of mapping accounts have thus far focussed on examples from empirical science in which there is already an obvious candidate for the 'mathematical structure' in question. If we are to test the plausibility of the mapping account's key assumption, then we must see how this assumption fares in light of more complex examples. Second, we must think about what it is for a mathematical structure to be 'of the right kind', given some application. At various points, proponents of the mapping account have specified in formal terms what the mathematical structures in question must look like. If we are to consider more complex examples, however, we should ask: what aspects of the role played by mathematics in some empirical scientific situation should we expect the structure to capture? We will put forward two requirements that any specification of the mathematical structure at hand should meet. These are that the structure should at least capture (i) all of the pieces mathematics relevant to the application and (ii) the *connections* between those pieces of mathematics that are required for the application.

Having clarified these two issues, we will find ourselves in a position to address our original question: is the assumption relied upon by mapping accounts a plausible working hypothesis about applied mathematics? We do so simply by presenting a case of the relevant kind and asking whether we can provide a corresponding mathematical structure that meets the two requirements indicated above. We will suggest that the notion of 'mathematical structure' employed by the proponents of mapping accounts makes it difficult to see how such a structure can meet our two requirements given some aspects of the case we present. When we venture into more complicated territory, the assumption on which the mapping account relies appears more like wishful speculation than a plausible working hypothesis.

This paper, then, will proceed as follows. In §2, I will outline the details of the mapping account as well as a key assumption on which it relies, which I call the MATHEMATICAL STRUCTURE ASSUMPTION. In §3, I will suggest that discussions of mapping accounts in the literature typically focus on a restricted class of examples in which this MATHEMATICAL STRUCTURE ASSUMPTION is easily satisfied, and that considering the viability of the mapping account will require that we venture beyond this restricted class of examples. In §4, I outline two requirements that should be met by attempts to accommodate more complex cases on behalf of the mapping account. In §5, I present the central case study of the paper, examining the way in which climate scientists and applied mathematicians model the permeability of sea ice as it features in large scale climate models. In §6, I argue that this more complex case presents a serious challenge for the MATHEMATICAL STRUCTURE ASSUMPTION, and in §7 I consider various ways a proponent of a mapping account might try to meet this challenge. Finally, in §8 I offer some concluding remarks.

2 Mapping Accounts and the Mathematical Structure Assumption

The central claim of the mapping account is that the utility of mathematics in application can be explained by the existence of some kind of mapping between the physical world and the 'mathematical structure' employed.[1] On this view, mathematics furnishes us with a rich variety of structures, some of which by chance or design mirror structural features of the world. These structural similarities allow us to employ mappings in order to investigate physical systems by working with corresponding mathematical structures. As long as the mappings preserve structure in the right way, we can simply read facts about our empirical system off the mathematics. According to mapping accounts, then, "the explanation of the utility of mathematics is no different from explaining the utility of street maps" (Bueno and Colyvan, 2011, 346).

The mapping account in this form was first put forward by Pincock (2004a,b). In this original formulation, Pincock did not specify any particular constraints on the mappings involved, suggesting that different applications may require different kinds of mappings. Bueno and Colyvan (2011) in turn

[1]For ease of expression, I will occasionally speak generically of 'the mapping account' instead of 'mapping accounts', but this is merely stylistic. There are, of course, a variety of mapping accounts in the literature, united by their commitment to the claim articulated in this section. These include Bueno and Colyvan (2011); Bueno (2016); Bueno and French (2018); Pincock (2004a,b, 2007, 2011); Nguyen and Frigg (2017)

suggest that this is a shortcoming of Pincock's account, and that "until the central notion of *mapping* is clarified, the account is little more than a gesture at an account" (Bueno and Colyvan, 2011, 348). That is to say that while Bueno and Colyvan agree that Pincock's mapping account as outlined above is "right as far as it goes," they believe that it will remain incomplete until it is augmented by more detailed consideration of the mappings that might be involved. Their own 'inferential conception of applied mathematics' attempts to do exactly that, by clarifying the central notion of mapping and addressing some of the specific issues that Pincock's account leaves open. Bueno and Colyvan explicitly state that their account should be seen as "an extension of the latter [Pincock's] account, in that it agrees that mappings of a variety of sorts are crucial to applied mathematics." (Bueno and Colyvan, 2011, 352)

There are two reasons that we will examine Bueno and Colyvan's account in particular detail and treat it for our purposes as a helpful prototype representative of mapping accounts in general. This augmented inferential conception has come to occupy a prominent place in discussions of mapping accounts in the contemporary literature. It thus forms the backbone, for instance, of the account provided more recently by Bueno and French (2018), as well as the extension offered by Nguyen and Frigg (2017). By and large, discussion of mappings accounts in the recent philosophical literature revolve around some form or another of Bueno and Colyvan's inferential conception. The second reason is that it provides an explicit account of how the mappings between physical systems and mathematical structures are supposed to work. This means that it is possible to consider in more detail how exactly a mapping account might deal with the complexities of some given case of applied mathematics. This will be important later when we want to ask whether the mapping account has any resources by way of which it can defend its central claim. In short, the inferential conception is both the most prominent mapping account on the market and the one that provides the most in the way of explicit suggestions on how more complex cases of applied mathematics might be construed in terms of mappings, and thus will occupy a central place in the argument to come.

Rather than invoking single mappings back and forth between some mathematical structure and physical structure, the inferential conception suggests that the interplay between mathematical structure and physical structure takes place across several steps. We first establish a mapping from some empirical set up to our mathematical structure (the *immersion* step). Having immersed our empirical set up in some mathematical structure, we then draw consequences from our mathematical formalism (the *derivation* step). With these consequences in hand, we then interpret our results in terms of our empirical set up (the *interpretation* step). This interpretation takes the form of a mapping back from our mathematical structure to our physical structure, and importantly this mapping need not simply be the inverse of the mapping that featured originally in the immersion step.

Roughly, then, the application of mathematics in empirical scientific contexts is to be explained by way of this process of immersing some empirical set up, broadly construed, into a mathematical structure where results can be obtained deductively with more ease, at which point we interpret these results

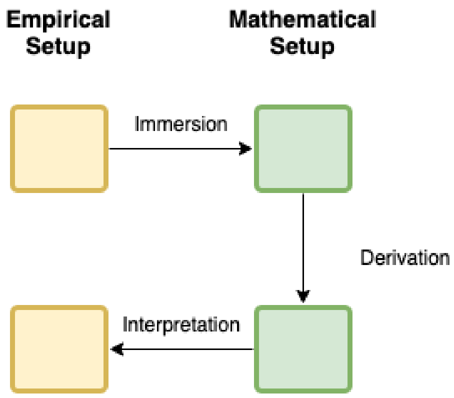

Figure 1: The Inferential Conception of Applied Mathematics.

by mapping back into our empirical set up. It is important to ask: what kind of mathematical structures are at play here? In framing their own amendment to the inferential conception, Nguyen and Frigg offer a helpful statement of the kinds of structure at play in the mapping account. They write:

> "[...] it is important to be clear on the nature of the objects and relations that make up a set-theoretic structure. The important point is that it does not matter what the objects are intrinsically. The only thing that matters from a structural point of view is that there are so and so many of them. All we need are dummies or placeholders. Likewise for relations. It is irrelevant what the relation "in itself" is. All that matters from a structural point of view is between which objects it holds. For this reason, a relation is specified purely extensionally, that is, as class of ordered n-tuples and the relation is assumed to be nothing over and above this class. Thus understood, relations have no properties other than those that derive from this extensional characterization, such as transitivity, reflexivity, and symmetry." (Nguyen and Frigg, 2017, 5)

The mathematical structures employed by the mapping account are nothing more than the extensionally-specified set-theoretic structures of model theory. The results obtainable within some mathematical domain will simply follow from facts about the relations that hold over the elements of that domain.

For the most part, discussions of mapping accounts in the literature have focussed on issues related either to the nature of the mappings involved or the kind of physical structure required. Natural questions arise regarding the kinds of mappings that the account requires (isomorphisms? homomorphisms? monomorphisms? isomorphic embeddings? etc.), what kind of structure the account assumes the physical world to have, the role played by scientific theories in generating that structure, and so on (Nguyen and Frigg, 2017; Bueno and Colyvan, 2011; Bueno, 2016; Bueno and French, 2018; Pincock, 2004a,b, 2007, 2011; Stemeroff, 2021).[2]

[2]Exceptions to this trend include Kasirzadeh (2021), Batterman (2010) and Rizza (2013), though they do not pursue the line of criticism we develop here.

If we are interested in assessing the plausibility of the mapping account as an account of applied mathematics, however, there is another issue that we must consider. The mapping account relies on the following assumption:

> MATHEMATICAL STRUCTURE ASSUMPTION: Given any application of mathematics in empirical science, there is a mathematical structure of the right kind that captures the mathematics relevant to that application.

If the reason that we are able to usefully apply mathematics to physical systems is that mappings obtain between the world and some mathematical structure, then there must be some appropriate mathematical structure for each useful application of mathematics in empirical science. In short, for the mapping account to work, we require a mathematical structure for each application.[3]

Why is it that proponents of mapping accounts are committed to the MATHEMATICAL STRUCTURE ASSUMPTION? Recall that the central claim of the mapping account, affirmed both by Pincock's original account and the inferential conception that has come to dominate the recent landscape, is that the utility of mathematics in application can be explained by the existence of some kind of mapping between the physical world and some appropriate mathematical structure. Suppose, then, that the MATHEMATICAL STRUCTURE ASSUMPTION were false. Then there would be at least one case of applied mathematics for which there is no structure of the right kind to capture the mathematics relevant to the application. Such a state of affairs would clearly threaten the central claim of the mapping account – after all, how can the utility of mathematics in this application be explained by mappings between the physical world and some appropriate mathematical structure if no such appropriate mathematical structure exists? In other words, the reason that mapping accounts rely on the mathematical structure assumption outlined above is that it is required by the central claim to which they are all committed despite their differences.

It may seem at this stage as though I am being uncharitable to proponents of the mapping account: surely they need not hold that all the mathematics involved be crammed into a *single* structure? Why can't multiple mathematical structures play a role without needing to be lumped together into a single behemoth? There are two things worth saying here. The first is that proponents of mappings accounts speak overwhelmingly in terms of 'the structure' into which our physical system is mapped. For instance, Bueno and Colyvan's inferential conception clearly conceives of applications in terms of a *single* empirical set-up and a *single* mathematical structure. The second is that when proponents of mapping accounts *do* talk about multiple structures being employed, they typically opt to construe the broader application in terms of several sub-applications. that is, what might on the face of it look like a single application of mathematics is decomposed into several applications, and in

[3]Although the mapping account falls pretty naturally out of a broadly structuralist philosophy of mathematics, its proponents do not intend for it to depend on any particular view about the foundations of mathematics. In fact, both Bueno and Colyvan (2011) and Nguyen and Frigg (2017) suggest that it is a strength of their mapping accounts that they do not depend on any particular philosophy of mathematics.

each case there is a *single structure* into which we map. In short, the main strategy by way of which mapping accounts bring more than one mathematical structure into play is by construing some case as a *series* of applications, of which the MATHEMATICAL STRUCTURE ASSUMPTION is all the while true. We will discuss such strategies in §7. For now, it is enough simply to note that such strategies do not represent an abandonment of the MATHEMATICAL STRUCTURE ASSUMPTION or an indication that it is too simplistically rendered, but rather an attempt to *save* the assumption by construing one broad application in terms of several separate ones.

We might also ask: given that applications of mathematics appear to serve a variety of purposes in the realm of empirical science, what purposes in particular are relevant to our assessment of the viability of the mapping account?[4] There seems on the face of it to be a difference, for instance, between using mathematics to *represent* some aspect of a system and using mathematics to run some numerical algorithm which might *predict* the future value of some parameter in a system. Be that as it may, proponents of the mapping account do not delineate a particular set of purposes for which they take their accounts to be primarily responsible. Rather, they market their account as capable of handling more or less any application of mathematics in the investigation of the physical world. Pincock, for instance, writes that his account aims at the task of articulating in general "what connection there is between the physical world and mathematics that can explain the successful application of mathematics in scientific reasoning" (Pincock, 2004b, 137). Bueno and Colyvan note that mathematics plays a variety of roles in empirical science, from unifying scientific theories to providing novel predictions to facilitating explanations, before declaring that "all of these roles, however, are ultimately tied to the ability to establish inferential relations between empirical phenomena and mathematical structures, or among mathematical structures themselves" (Bueno and Colyvan, 2011, 352).

The upshot of the this appears to be that the MATHEMATICAL STRUCTURE ASSUMPTION should be construed in an analogously general fashion. If, for instance, it were the case that the mapping account was intended to capture only the application of mathematics in representative contexts, then the MATHEMATICAL STRUCTURE ASSUMPTION would only need to hold in those contexts as well, and we would have to take care to ensure that any purported counterexample did not stray into applicational contexts the mapping account was not designed to cover. As it happens, the stated aim of the mapping account means that we need not be quite so concerned.[5] That is to say that if we find a case that the mapping account does not seem capable of construing in the way it requires, then it is no less a counterexample to the view for being, for instance, a case of mathematics used for predictive purposes rather than explanatory purposes (or representational purposes, or whichever).[6]

[4]Thank you very much to anonymous reviewer for raising this point.

[5]This is helpful, all things considered, since once we begin to think about more complex uses of mathematics than those that have heretofore featured in discussions of mapping accounts, we might find that it is more and more difficult to clearly distinguish exactly what single purpose some particular piece of mathematics is serving at some particular point in time.

[6]It may well be open to the proponent of the mapping account to respond to a purported counterexample by narrowing their account such that it only applies properly to some particular

3 Parcels of Mathematical Structure

If we are interested in the broader viability of the mapping account of applied mathematics, then we must ask ourselves: how plausible is the MATHEMATICAL STRUCTURE ASSUMPTION? One natural way to go about answering this question would be to consider how well the assumption holds up in light of various real life scientific applications of mathematics. Of course, in defending and articulating their accounts, proponents of mapping acounts have considered a variety of such real life cases. Indeed, the literature on mapping accounts features discussions of examples of applied mathematics ranging from the simple (counting rabbits, the Bridges of Königsberg) to the more complex (group theory and quantum mechanics).

There is, however, one sense in which proponents of mapping accounts appear to have confined themselves to the consideration of a restricted class of examples of applied mathematics. Almost exclusively, the examples that appear in discussions of mapping accounts are scientific situations that involve (or appear to involve at first glance) the employment of self-contained fields of mathematics in isolation. That is to say that we are very often asked to "consider the use of [mathematical field X] in [scientific theory Y]." In outlining the way in which the inferential conception (along with some additional formalism surrounding 'partial structures') allows us to understand the real world application of mathematics, Bueno and French choose as their flagship case study "the introduction of group theory into quantum mechanics" (Bueno and French, 2018, 73). (Bueno and Colyvan, 2011, 359) cite the use of analysis in neo-classical economics, while in discussing applications of mathematics in his original articulation of structuralism Shapiro himself considers "the use of real analysis or geometry in physics" (Shapiro, 1997, 252).[7] Pincock (2004a) primarily discusses the use of arithmetic in counting physical objects and discussing their magnitude, while in a later paper (Pincock, 2007) he mentions the application of graph theory to the bridges of Königsberg problem. In detailing the application of their extended account, Nguyen and Frigg consider the use of arithmetic in counting populations of rabbits and representational measurement theory in assessing the length of various metal rods (Nguyen and Frigg, 2017, 16).

There are, however, two important respects in which such cases are not representative of the applications of mathematics that feature in scientific practice. First, the practical requirements of scientific investigation often ensure that the mathematics involved 'overflows' the domain that it might

restricted class of uses. There are, it seems to me, two reasons that this is unlikely to represent an attractive option to proponents of the mapping account. The first is that, as mentioned in the previous footnote, complex cases of applied mathematics may involve questions of prediction, explanation and representation dovetailing in ways that might make the question of what single purpose some mathematical construction is serving at some particular time somewhat difficult to untangle. Perhaps more importantly, the case we will discuss later on at the very least includes mathematics used both for the purpose of representation and the purpose of facilitating predictions. Amongst the purposes for which mathematics is used in empirical science these seem particularly central, and it seems unlikely that proponents of mapping accounts would want to abandon the claim that their account is intended to cover them.

[7] Of course, Shapiro is not offering a full-blown account of applied mathematics in the way that Bueno, Colyvan, French and others are.

at first glance have appeared restricted to. We might, for instance, observe a situation in which a scientist exploits their knowledge of some physical system in order to generate some partial differential equations (PDEs) describing its behaviour. It may then seem clear that the mathematical structure at play in such a case is, as Bueno and Colyvan conclude in the example we saw above, "certain mathematical structures of analysis" (Bueno and Colyvan, 2011, 359).

The problem is this: writing down a set of PDEs is one thing, and solving them is often another thing entirely. Analytic solutions are often difficult if not impossible to find, in which event scientists must rely upon a variety of numerical techniques in order to extract workable information about their system from their set of equations.[8] These techniques are often classified in quasi-experimental terms according to their computational performance, probabilistic error bounds, general dynamic behaviour, and so on.[9] It is not clear that such techniques will always lie within the reaches of "the resources of analysis" (to use Bueno and Colyvan's phrasing). The lesson: just because we can *frame* (i.e. we can write down the equations for) some model of our system in terms of some neat mathematical structure does not always mean that all the mathematics involved in that application can be confined to that structure.

A brief analogy may help to illustrate what I mean by this distinction between the mathematics required to frame some model of our system and the mathematics required to work with the model. Building the foundation and frame of a house might only require wood and cement, plus perhaps some nails and other bits and pieces. With the frame in place, we can see the structure of the house, how the rooms fit together, where and how it is secured to the foundation, and so on. But this does not mean that *constructing the whole house* only requires wood and cement. Turning the frame and foundation into something that people can actually *live in* requires tiles, plumbing, insulation, gyprock, and so on. There is a marked difference between the 'resources' we need to build the central structural frame of the house and those required to turn this frame into a livable domicile. The point of the analogy here is that the central equations around which we organise our model of our system in this context may be something like the wooden frame of our house. We may only need a relatively simple array of mathematical resources to make sense of this central frame, but actually extracting information from our model (building a house someone can live in) may demand an entirely different set of resources.

One might wonder at this point whether framing a model in this sense and extracting helpful information from it might represent two *separate* purposes

[8] As an example, Fourier's Law for heat conduction can be written using relatively simple differential operators:

$$\frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2} + \frac{\partial^2 T}{\partial z^2} = \frac{\rho c_p}{k}\frac{\partial T}{\partial t}.$$

However this PDE can only be solved directly in the most simple cases. As such, solutions are typically approximated by powerful software packages that employ numerical methods and algorithms which are not properly part of analysis itself. See Bergman et al. (2011).

[9] A helpful presentation of the technical details of some of these methods can be found in Süli (2015). Discussion of some of the philosophical and conceptual issues arising out of the increasing prevalence of such computation methods can be found it Fillion and Corless (2014); Fillion (2019)

for which mathematics might be used in application. Either way, as we examine our case study later on there is a sense in which it will be important to keep track of the difference between the mathematics required to write down the equations for a model and the mathematics required to actually work with that model, and a sense in which it will not be. Insofar as proponents of the mapping account, as discussed in §2, do not distinguish between the different purposes for which mathematics might be used in application, we do not necessarily need to keep track of whether the case we present relates to the purpose of 'framing a model' or 'working with a model.' On the other hand, it will be important to pay attention to the distinction between framing and working with a model in the sense that although we can *write down* the key equations of our case study in a way that might lead us to believe that we only need 'the resources of analysis' to appear in the relevant mathematical structure, actually *working with* the model as applied mathematicians do requires many more tools than this.[10]

The second respect in which the class of cases described is unrepresentative of broader scientific practice is this: scientific application of mathematics often involve not just the kind of 'overflow' noted above but also the conscious stitching together of various dissonant pieces of mathematics. It is often the case that the mathematics appropriate for investigating the behaviour of some part of our physical system is very different from that required for dealing with some other part of the system. As Batterman, Wilson and others have emphasized, scientists must commonly embrace exactly this dynamic in dealing with the behaviour of physical systems at different scales (Wilson, 2017; Batterman, 2013; Green and Batterman, 2017).

Suppose we find that some mathematical representation $\mathcal{R}_1$ is appropriate for describing the behaviour of our system at some length scale ℓ_1. We might find, however, that the system's behaviour on the scale of ℓ_1 depends on some aspect of its behaviour at some lower length scale, ℓ_2, for which some completely different mathematical representation, $\mathcal{R}_2$, is appropriate. If we are very lucky, our $\mathcal{R}_2$ might capture the lower scale behaviour in such a way that the relevant information can be directly 'plugged into' our larger scale $\mathcal{R}_1$. More often, however, we are not so lucky, and finding a way to get our two dissonant representations to harmonise fruitfully is a difficult problem. Solving this problem may well require that we rely on additional physical information not properly relevant to either $\mathcal{R}_1$ or $\mathcal{R}_2$. Processing this additional physical information may well call for fresh mathematical techniques. In any event, the information that we finally extract may be the result of various different mathematical tools which are made to work together as the result of both empirical and mathematical considerations.

Laying out the ways in which applications of mathematics in physical science routinely differ from those commonly considered by proponents of mapping accounts helps us to see the following: it is *precisely* in the kind of case typically considered in the literature that the MATHEMATICAL STRUCTURE ASSUMPTION looks most plausible. Given that discussions of mapping accounts in the philosophical literature feature plenty of examples drawn from different areas of empirical science, it might *appear* as though the account has

[10] Thanks to an anonymous reviewer for pushing me on this.

been tested against a selection of cases that is representative of the complexity of modern scientific methodology. What I have argued is that this appearance is misleading. Although the cases that feature in the literature do indeed come from diverse areas of empirical science, they otherwise have tended to cluster within a particular part of the applied mathematical landscape. Moreover, the part of the landscape within which these cases cluster is one in which the terrain is *particularly* favourable to the mapping account's core assumption.

If we are interested in whether it is true that we can find a mathematical structure of the kind required given any applications of mathematics, then we are not going to learn very much if we focus on cases in which there is a clear and obvious candidate for said 'mathematical structure'. The foregoing discussion demonstrates that if we want to investigate the plausibility of mappings accounts of applied mathematics, then we must examine the MATHEMATICAL STRUCTURE ASSUMPTION in light of a case of applied mathematics of a more complicated character. That is, we require a case in which there is not a clear and obvious candidate for the relevant mathematical structure. Before we can present such a case, however, we must address one final issue.

4 Two Requirements

The proposal we have so far developed is that we should not consider the MATHEMATICAL STRUCTURE ASSUMPTION to have been properly tested until we expand our gaze beyond the restricted class of examples typically considered by proponents of the account. The natural next step, then, is to present a case of applied mathematics from somewhere outside of this restricted class and see how the assumption fares. Since the key assumption holds that there is a mathematical structure of the right kind given any application of mathematics, we simply need to see if we can find such a structure given our more complicated case.

But what exactly is a "mathematical structure of the right kind"? As we have seen, proponents of mapping accounts tend to fill in the 'mathematical structure' side of the ledger by citing the broad field or subfield in which the salient mathematics appears to fall. Given that we want to consider precisely those cases in which the salient mathematics cannot be contained in a neat mathematical fiefdom, it will help to be clear about exactly what we expect the 'structure' on the mathematical side of the ledger to capture. For instance, could the proponent of a mapping account very simply circumvent the problem by declaring the relevant mathematical structure to be some model-theoretic surrogate for the entirety of mathematics? In short, we want to think what constraints should apply to our hypothetical mapping account proponent's attempts to fill in the 'mathematical structure' side of the ledger.

A natural first thought is that at a minimum the mathematical structure provided must contain all of the mathematics employed in the application under consideration. We should note here that, as outlined earlier, there is often a marked difference between the mathematics required to *frame* some set of equations or model of a system and the mathematics that we must eventually employ to extract reliable information about our target system from that model. Given that the mapping account attempts to explain why mathematics is so useful to us in application, our mathematical structure should include the full gamut of mathematical tools that we actually use to investigate the behaviour of our system.

This requirement, however, does not rule out that a mapping account proponent might respond by simply suggesting that the relevant structure is some model-theoretic surrogate for the entirety of mathematics, as we mentioned above. Consider the kind of multiscalar modelling situation that we described in §4. Suppose that we can capture all of the mathematics required to frame and work with our larger scale model in some structure $\mathcal{S}_1$, and similarly for our lower scale model and some structure, $\mathcal{S}_2$. Can we not simply say that the relevant overall mathematical structure is something like $\mathcal{S}_1 \cup \mathcal{S}_2$? The problem here is that we sometimes require both additional mathematics and additional empirical information in order to understand how our upper and lower scale model can co-operate to provide a picture of the system's overall behaviour. Our models often need to be 'stitched' together, and understanding the mathematical and empirical character of this stitching is key to understanding what mathematics helps us to achieve in such applications.

Such multiscalar challenges suggest another requirement: the mathematical structure singled out must capture the connections between the various pieces of mathematics that allow them to work together in application. Applied mathematicians take great care to ensure that the threads that run through their various modelling tools are woven in such a way as to produce reliable results.[11] If we respond to the challenge of ensuring that our structure contains all of the mathematics employed in application by gesturing at a structure that does not allow us to articulate the complicated relationships between those mathematical tools, then we have met the first requirement at the expense of overlooking a vital part of the application in question. If mapping accounts are to succeed in explaining the utility of mathematics beyond a limited range of examples, then the structures to which they appeal must be able to capture these more complicated dynamics.

The above considerations put us in a position to suggest two requirements that the mathematical structures referenced in the MATHEMATICAL STRUCTURE ASSUMPTION must meet.

SCOPE: The structure captures all of the mathematics relevant to the application.

STITCHING: The structure captures the relevant relationships between the pieces of mathematics employed in the application.

In short, the SCOPE requirement states that all of the mathematics should be found *somewhere* in the structure. The STITCHING requirement emerges from the fact that the mapping account claims to explain the utility of mathematics in actual scientific practice. The weaving together of various pieces of mathe-

[11] This point has been made in different ways elsewhere by several philosophers of science. These include Batterman (2013, 2021); Green and Batterman (2017); Bursten (2018, 2021); Wilson (2017); Winsberg (2006, 2010).

matical machinery is core to many applications of mathematics in scientific practice, and so our structure must be able to handle this dynamic.[12]

Let's recap. If the mapping account is right, then we are able to employ mathematics in empirical scientific contexts because there exist mappings between the physical world and some mathematical structure. This claim relies on the assumption that there is a mathematical structure of the right kind for each application of mathematics. If we want to evaluate this assumption properly, we must consider a case in which there is not a clear and obvious candidate for such a structure and see if one can nonetheless be provided. Not just any structure will do, however. If the mapping account is truly to explain the utility of applied mathematics, then the structure put forward must be able to handle the fact that applications of mathematics can involve a wide and open-ended variety of mathematical tools (SCOPE) that are woven together in complicated ways to secure results (STITCHING). Now that we know what we are looking for, we can turn our attention to a case of the relevant kind.

5 Sea Ice Permeability and Climate Models

The interaction between large scale climate models and the lower scale microstructure of sea ice provides a vivid example of the kinds of complexities we have so far mentioned. In rough outline, the values of some parameters that feature in large scale models of the Earth's climate depend on the fluid permeability of sea ice (that is, more or less, how easily pockets of fluid can move within large blocks of ice). This fluid permeability is in turn influenced largely by the structure and arrangement of liquid brine inclusions on a scale much smaller than that of the original climate models. This presents two challenges: characterising the lower scale behaviour of the ice-brine microstructure, and finding a way to incorporate this information into our large-scale climate model. In examining the way in which applied mathematicians go about tackling these challenges, we will pay special attention to exactly what pieces of mathematics are relevant to the application and how these pieces are related to one another.

5.1 Climate Models and Sea Ice

Suppose we are interested in modelling large-scale fluctuations in the Earth's climate.[13] As it happens, an important part of this process involves understanding the way that sea ice behaves at various scales. On the whole, solar radiation that hits sea water tends to be absorbed, while solar radiation that hits sea ice tends to be reflected. The ratio of reflected sunlight to incident sunlight is known as *albedo*, and so we might rephrase the above by saying that the albedo of sea ice (0.8) is much larger than that of water (< 0.1). Sea

[12]It is worth noting that the proponents of mapping accounts that we have considered appear to take themselves to be bound by requirements like these, although they do not consider the question as explicitly as we have done. In particular, both (Pincock, 2011, 212) and (Bueno, 2016, 2597) consider the expectations we might apply to mapping accounts of applied mathematics. The standards they put forward seem to me closely related the two requirements outlined here.

[13]The presentation in this section draws largely from Golden (2015, 2009, 1997).

ice thus both prevents solar radiation from warming the water beneath and prevents ocean heat from escaping and contributing to atmospheric warming. As warming temperatures melt sea ice, fewer surfaces are available to reflect sunlight and more ocean heat escapes to warm the atmosphere, which causes more ice to melt. This leads to a decrease in the albedo of the polar oceans, which leads to more solar absorption and warming, which leads to further decreases in albedo, and so on and so forth. This dynamic is called *ice-albedo feedback* and it plays an important role in the global temperature system.

The above is to say that understanding the dynamics of ice pack albedo is a key ingredient in any attempt to predict the future trajectory of the Earth's climate. In turn, the albedo of an ice pack is determined by the presence of melt ponds which form on the surface of the ice. Whether these melt ponds spread, deepen, or drain is for the most part determined by how easily pockets of fluid can move through the ice below. This is known as the ice's *fluid permeability*. Sea ice differs from glacial ice in some important respects. When salt water freezes, the resulting material is a composite of pure ice with inclusions of liquid brine, air pockets and solid salts. As the temperature increases, so does the volume fraction of brine (ϕ). Below a certain critical volume fraction, the sea ice is more or less impermeable to fluid flow. Above this critical volume fraction, the ice's fluid permeability depends not only on the fraction of brine inclusions, but also on the *structure* of those inclusions.

The upshot of the above is that the task of understanding and predicting the large scale behaviour of sea ice relevant to climate models is connected to our ability to understand and characterise the much smaller scale structure of (amongst other things) liquid brine inclusions. In other words, our ability to model sea ice on the scale of hundreds of kilometres depends on our ability to effectively characterise behaviour that takes place on the scale of centimetres and metres (the size of the brine inclusions). The challenge, then, is to work out how we can model the ice-brine microstructure in a way that allows us to fill in the details required by our large scale climate models.

5.2 Homogenization

On the whole, global climate models are made up of systems of partial differential equations. These PDEs are solved using extremely powerful computers that split the Earth's surface into three dimensional grids with horizontal grid sizes in the order of tens of kilometres. Several of these equations contain parameters whose behaviour is influenced by facts about sea ice systems, such as the ice thickness distribution. The equation in which the behaviour of our brine flow plays an important role is the one describing the temperature field ($T(x,t)$) inside our sea ice. Coupling the sea ice to the ocean above and the atmosphere below through the right kind of boundary condition yields:

$$\frac{\partial T}{\partial t} = \nabla \cdot (D(T)\nabla T) - \boldsymbol{v} \cdot \nabla T, \tag{1}$$

where D is the thermal diffusivity of the sea ice and $\boldsymbol{v}$ is an averaged brine velocity field in the sea ice.

The difficulty here is that the velocity of our brine flow oscillates wildly as we move through the sea ice, given that the flow is going to be (more or less) zero at any point that is ice and not a brine inclusion. Moreover, these velocities are going to depend in some way on the geometrical properties of the brine microstructure. As a result, there is not really any straightforward way to calculate the value of the parameter v that features in equation (1).

In cases like this, it is sometimes possible to approximate the value of the parameter in question by way of a method known as *homogenization*. When we are faced with some kind of heterogeneous medium which exhibits important microscale structure, such as our ice-brine composite, homogenization techniques allow us to:

> "take into account the structure of the heterogeneous medium by calculating an equivalent homogeneous 'effective medium' and to use the equivalent medium in further calculations." (McPhedran, 2015, 500)

The basic idea is that we want to replace the oscillating coefficient (our velocity field) with an *effective parameter* that both takes into account the microstructure of the heterogeneous medium and behaves in a more uniform fashion. If we are able to find such an effective parameter, then we can replace our original equation (or set of equations) with one that treats the material as though it is homogeneous rather than heterogeneous. In short, these *homogenized equations* describe a fictitious 'effective medium', the behaviour of which is determined by our effective parameter and thus mirrors the behaviour of our heterogeneous system.[14]

Unless our medium is nicely periodic,[15] there is for the most part no general method for arriving at these homogeneous equations (Santosa, 2015, 103). Applied mathematicians have instead developed a wide variety of treatments and techniques appropriate to different materials, most of which require detailed empirical information about the system's behaviour.[16] In the case of our ice-brine composite, the application of such homogenization techniques relies on empirically measured properties such as the statistical distribution of the inhomogeneities, the Reynolds number of the flow, and the identification of a fast and slow time scale in the oscillations of our velocity field. In addition, researchers must verify that a variety of other conditions of the geometry of the medium are met (Golden, 1997; Torquato, 2002; McPhedran, 2015). With all of this in hand, Golden (1997) was able to exploit several limiting relationships (the *homogenization limit*) to show that the averaged velocity $v(x)$ and the pressure $p(x)$ satisfy

[14] Batterman (2021) contains several helpful discussions and illustrations of the way that homogenizations techniques allow us to extract information about the behaviour of heterogeneous composites by connecting continuum and mesoscale level descriptions. In this respect, our sea ice example falls into a broader category of material scientific problems concerned with heterogeneous materials, including that of understanding the behaviour of a mixed conductor, for instance.

[15] A *periodic* medium is one that, although not homogeneous and thus composed of at least two different kinds of material, is arranged in such a way it exhibits some periodic internal structure (e.g. some characteristic or property of the materials involved varies periodically as you traverse some path through the medium).

[16] A helpful overview of these can be found in McPhedran (2015)

$$v = -\frac{1}{\eta} \boldsymbol{k} \nabla p, \quad \nabla \cdot \boldsymbol{v} = 0 \tag{2}$$

Here $\boldsymbol{k}$ is called our *permeability tensor* and is our *effective parameter*, and η is the viscosity of the brine inclusions.

There are three points worth stressing at this stage. The first is that obtaining our homogenized equations does not yet mean that we understand the behaviour of the effective parameter that features therein. The process of homogenization simply allows us to replace an intractable task (calculating the velocity field over our entire system) with a more tractable one (investigating the effective fluid permeability of ice-brine microstructures). Dealing with this more tractable task will still involve constructing a model of the processes taking place on the smaller scale and using a variety of techniques to understand the behaviour of the effective parameter.

The second point is that the homogenization process then is best thought of as what allows us to stitch together our large scale climate model and the smaller scale model of the ice-brine microstructure that we will meet shortly. Our climate model demands information about our sea ice that we simply cannot calculate on the length scale required. Our smaller scale models of sea ice, however, can provide us with computationally tractable information about some other piece of information. Homogenization techniques serve to mediate this mismatch, relying on a combination of empirical information and mathematical manipulation to relate the information provided at the lower scale to that required at the larger scale.

The final point worth emphasizing is that this process of homogenization is not simply a purely mathematical operation that we perform on our original set of equations. The homogenization process relies crucially both on mathematical properties of the original PDEs and a variety of empirical information about the system as a whole. We can only be sure that the relevant homogenization limits exist, for instance, once we have collected additional data about the structure of our ice-brine composite. The problem faced by applied mathematicians in our sea ice case (and many others) is that the behaviour captured by our large scale model is influenced by much smaller scale processes that cannot be tractably modelled in the way required by the larger scale model. The process of homogenization allows us to relate the requirements of the larger scale model to information we can extract from lower scale models by way of an intermediary effective parameter. If we'd like, we can in some broad sense think of this as allowing us to construct a relation between two mathematical structures (i.e. our upper and lower scale models). In this case, we can put the key point this way: the 'relation' that emerges from such homogenization techniques is not purely mathematical, but rather involves a mixture of mathematical properties of the respective models and empirical features of the system at various scales.

5.3 Predicting Sea Ice Permeability

In the previous section, we noted that the homogenization process alone does not always provide us with an understanding of the behaviour of the relevant effective parameter (in our case, $\boldsymbol{k}$). Before we move on from our sea ice case, it will be helpful to examine the details of some of the techniques employed by applied mathematicians in order to obtain qualitative information on the effective fluid permeability of our sea ice. When a full characterisation of the behaviour of our effective parameter proves elusive, applied mathematicians often focus instead on more restricted aspects of that behaviour, such as values near criticality or upper bounds. Focussing on different aspects of our system's behaviour may require that we employ radically different mathematical techniques. In what follows we will examine the way that lattice models are used to investigate the value of our effective permeability parameter close to observed critical volume fractions.

A *lattice percolation model* consists of a d-dimensional integer lattice, $\mathbb{Z}^d$, along with a network of 'bonds' which join the lattice sites to their nearest neighbours.[17] The bonds are assigned 'conductivities' $\sigma_o > 0$ (the bond is open) or 0 (the bond is closed) according to some probability p. In other words, the probability that any particular bond is open is p and the probability that it is closed is $1 - p$. These probability-driven assignments of conductivity are intended to capture the ability of some particular bond to transport some quantity between lattice sites. Groups of connected open bonds are called *open clusters*.[18]

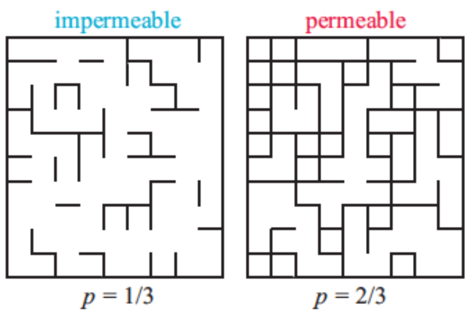

Figure 2: Examples of two-dimensional square lattices for probabilities on either side of the percolation threshold (which, for two dimensional square lattices is $p_c = \frac{1}{2}$).

As it happens, there is a choice of our probability p above which the average size of these open clusters diverges and an infinite cluster (that is, a cluster containing infinitely many lattice sites) first appears. This critical probability (p_c) is known as the *percolation threshold*. When p is below this percolation threshold, the probability that the origin (or any point) is contained in an infinite cluster (denoted $p_\infty(p)$) is 0, while once $p > p_c$ we have $p_\infty(p) > 0$ (Golden et al., 2006). That is, once we reach the percolation threshold there is a non-zero probability that any given point of our lattice will sit within a

[17] An *integer lattice* in d dimensions is one in which the lattice sites are labelled according to d-tuples of integers (for example, $(1, 2, 3)$ is a site in a 3-dimensional integer lattice).

[18] Such lattice models were originaly developed in order to investigate the effective electrical conductivity of composite materials, but may be repurposed to help us understand the fluid permeability of our ice-brine composite. In fact, there are deep analogies between the way that problems of electrical transport and fluid permeability are tackled by way of homogenization techniques (McPhedran, 2015, 502).

cluster of infinite size.

If we think of the bonds that connect the various sites of our lattice as representing the path followed by the various channels of brine through our sea ice, then it follows very naturally that the appearance of infinite clusters should coincide with our ice-brine composite exhibiting some fluid permeability. From this point of view, our infinite cluster density $p_\infty(p)$ plays the same role as our effective permeability $\boldsymbol{k}$. As such, Golden et al. (2006) were able to show that the vertical fluid permeability of our 3 dimensional lattice behaves like

$$\boldsymbol{k}^*(p) \sim k_0(p - p_c)^e, \tag{3}$$

Where $e \approx 2.0$ is also a universal exponent and k_o is called the *permeability scaling factor* and depends on the velocity of the liquid brine and the radius and connectedness of the channels (i.e. the brine microstructure) (Golden, 2009).

The result of the above is that we appear to have a workable model of the lower scale structure that influences the behaviour of at least some aspect of our homogenized parameter. Moreover, it seems like we can extract helpful qualitative information about our homogenized parameter from this model. We might be tempted on this basis to declare that our lattice model represents a 'mathematical structure' that allows us, given the right mapping, to derive the information contained in equation (3) in a purely mathematical fashion.

Things, unfortunately, are not so simple. Where our 3-dimensional lattice model reaches its percolation threshold at $p_c \approx .25$, empirical measurement indicates that the corresponding critical volume fraction of sea ice is $\phi_c \approx 0.05$. As a result of this discrepancy, researchers concluded that it was "apparent that key features of the geometry of the brine microstructure in sea ice were being missed by lattices" (Golden, 2015, 700). In other words, while the lattice percolation model provides a helpful framework for obtaining information about the behaviour of our ice-brine composite, we require some additional information about the microstructure involved before we can employ the result in (3) in application. In particular, the mismatch in critical percolation threshold prevents us from calculating the permeability scaling factor k_0 in a straightforward way.

Researchers have been able to identify a different set of percolation models originally developed for compressed powders that exhibited the same percolation threshold as that measured in our sea ice. The microstructure associated with such compressed powder models also seems to resemble that associated with our ice-brine microstructure in some topological respects. While these compressed powders models are percolation models in the broad sense, they differ from those we saw just above in that they do not employ a lattice structure. In this slightly different model, however, no straightforward derivation analogous to (3) is possible. As it happens, aspects of the lattice structure are integral to our ability to capture the permeability behaviour in some tractable form (Torquato, 2002).

The challenge, then, is this: how can we use the more specific microstructural characterisation of the compressed powder model to fill in the details required to actually extract applicable information from (3)? In particular, how can the compressed powder model help us to calculate the permeability scaling factor, k_0? The form of this problem is the same as the one we discussed in more abstract terms in §3: how do we incorporate information captured by a model using a certain mathematical machinery into a model built using very different mathematical machinery? In short: how can we stitch together these two models in order to obtain the information required for our application?

As in the case of our homogenization techniques, our two models cannot be stitched together in purely mathematical terms. Relatively recent advancements in techniques and technology for the detailed imaging of porous media have allowed researchers to employ high resolution x-ray computed tomographical scans of sea ice samples in addressing this challenge. Powerful software packages allow researchers to obtain highly detailed quantitative characterisations of the brine microstructure of sea ice from these tomographical scans. Since the compressed powder model captures more specifically the microstructural features of the ice-brine composite, we are able to combine this extensive computationally-derived data with the compressed powders model in a way that the lattice model does not allow (Golden, 2015). This high resolution tomographical data allows us to do two things: (1) compute approximate experimentally-derived relationships between the parameters that appear in the compressed powders model and those that appear in the lattice model, and then (2) exploit these relationships in using the compressed powders model to numerically compute the permeability scaling factor that appears in the lattice model. Researchers have achieved these tasks by applying techniques known as *finite-size scaling* and *critical path analysis* to the tomographically-augmented compressed powder model. With our k_0 thus determined, we are finally in a position to understand (3) as telling us something concrete about the permeability of our sea ice near criticality.

To summarise: the lattice percolation model allows us to obtain key information about the effective parameter that appears in our homogenized equations as long as we can calculate the permeability scaling factor. Mismatches between the measured percolation threshold of sea ice samples and that of the the 3D lattice mean that we cannot use the lattice model on its own to calculate this k_0. A different model allows us to calculate this k_0 but only insofar as we are able to combine the model with highly detailed computer-generated scans that allow us to approximate the relationships between the parameters that feature in the two models and then determine the k_0 by more or less brute force algorithmic computation. The chart in Figure 3 lays out visually how these elements all combine to help us understand the behaviour of our effective parameter, $\boldsymbol{k}$. The arrows from one box to another indicate roughly that the former provides input or is required to generate the latter.

The upshot here is that even when dealing with behaviour that takes place at roughly a single scale, scientists must nonetheless sometimes stitch a variety of mathematical tools together in order to obtain workable results. Moreover, just as with our homogenization techniques, this stitching is not purely mathematical. When we extract the result in (3) from our lattice model and discover that we need information about the structure captured by the compressed powders model to actually apply it, we do not simply construct some mathematical bridge between the two structures. Rather, we exploit the fact that the compressed powders model is better placed to interface with the X-ray computed tomographic data and thus that we can employ numerical and computational algorithms to extract the information required. The lattice model allows us to derive (3) and the compressed powder model allows us to capture some of the microstructure missed by the lattices, but the *connections* between the two required for the application are forged largely using empirical data rather than purely mathematical relations.

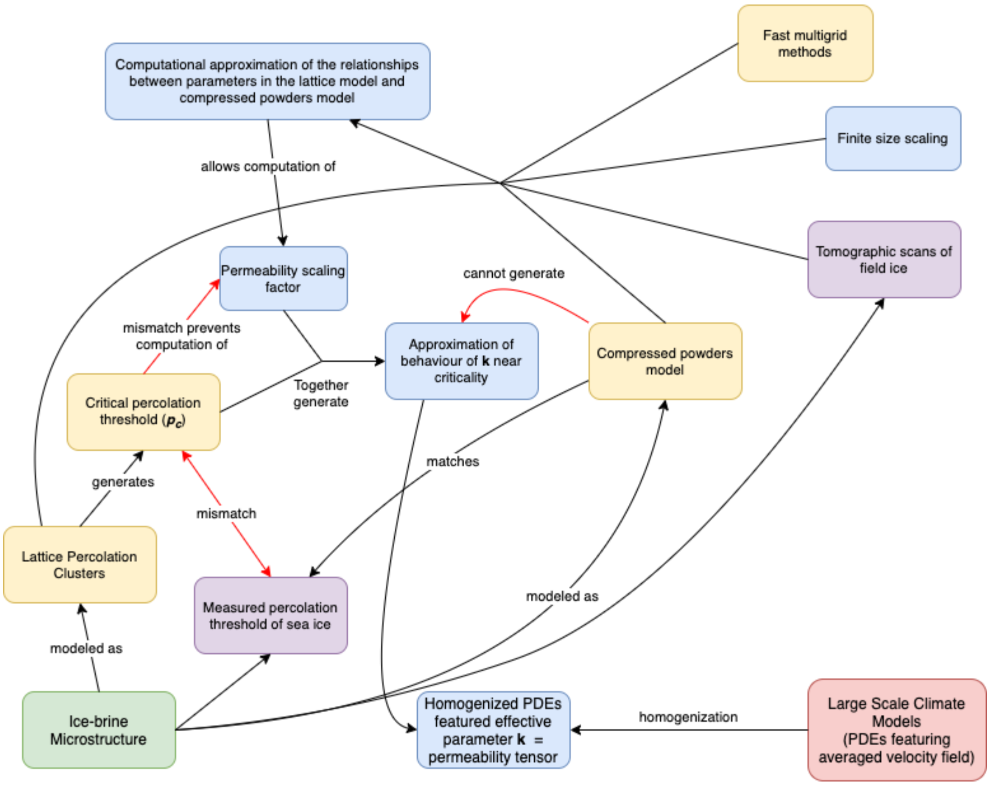

Figure 3: A rough illustration of the final lay of the land. Blue boxes indicate results and methods that involve a mixture of mathematical technique and empirical data. Yellow boxes indicate mathematical representations, methods, and results. Purple boxes indicate directly empirical information. Red arrows indicate that certain elements cannot be connected in ta straightforward way. An arrow pointing from multiple elements to a single element indicates that all of the elements are required to generate the result.

6 Evaluating the Mathematical Structure Assumption

We are now in a position to ask: how plausible does the key assumption of the mapping account seem in cases such as the one we have just examined? Recall that the MATHEMATICAL STRUCTURE ASSUMPTION holds that for any application of mathematics in empirical science there exists a mathematical structure of the appropriate kind. We suggested in §4 that for a structure to be 'of the appropriate kind' it must allow us to capture both the range of mathematical machinery involved and the relationships between those pieces of mathematics that underpin the application in question. As we have seen, applied mathematicians sometimes obtain results by forging complicated connections between different pieces of mathematics on the basis of both empirical and mathematical considerations.[19] If the mapping account is to succeed in explaining the utility of mathematics in application, then there must be appropriate structures for such applications. In short: if we take seriously the suggestion that the requisite mathematical structure meet our SCOPE and STITCHING requirements, then that structure must be capable of handling the complexities of the application we have just met.

We must then ask: can the mathematics employed in our sea ice case be captured in the appropriate kind of structure? There are two particular difficulties that such cases present. The first is that, as we mentioned earlier, the mathematics required to *frame* a model of some system's behaviour is not the same as the mathematics required to *apply* or *work with* that model.[20] Although we may only need 'some structures of analysis' (in Bueno and Colyvan's words) to write down the PDEs that appear in our large scale climate models, we saw that applying these equations to our sea ice system required appeal to geometrical properties of our system, probabilistic correlation functions, numerical and computational algorithms, and so on. Similarly, although we can *frame* the result obtained in (3) in terms of the 'structure' of the lattice model, we saw that applying this result required some additional computational machinery and a new percolation model.

The problem here is that the additional mathematical machinery we must employ to apply such models does not always seem as easy to capture in a neat 'structure' as the mathematics used to *frame* the result or model applied. If we restrict our attention to the mathematics required to frame our result and thus by and large to fields that admit of axiomatic treatment such as arithmetic, geometry, analysis, group theory, graph theory and so on, then we might not see any problem with the suggestion that there is a mathematical structure out there for every piece of mathematics that features. If we want our mathematical structure to account in some sense for all of the mathematics actually *used* in the application, however, then we may have cause for concern. Many numerical and computational methods for dealing with complex PDEs appeal as needed to a seemingly open-ended variety of mathematical notions from different fields, as well as a variety of empirical and quasi-experimental

[19] To clarify, by 'mathematical' here I mean related to the formal relational properties of the elements of some purported mathematical structure, whereas by 'empirical' I mean related to some physical feature of the system in question.

[20] At this point one still might wonder: what *kind* of application of mathematics have we just met? The mathematics in our sea ice case seems to serve a variety of purposes at once, such a prediction, representation, and so on. In particular, is the mathematics in our case study used for the purpose of working with the model involved or simply framing that model by way of a set of equations? I am skeptical that there will be a definitive and clear way to assign a single 'purpose' to applications of mathematics that exhibit the kind of complexity we see in the case just covered. Nonetheless, the point made in §2 bears reiterating: since proponents of mapping accounts take their account to cover *any* application of mathematics, we need not worry about whether the purpose for which the mathematics is employed in our case study is primarily related to framing a model or working with it. The point is that whatever the various purposes involved might be, the application that features in our case requires that we *both* frame a model using some mathematical tools and then investigate the behaviour of that model in certain tractable cases. If the mapping account is to handle such a case then we require a 'mathematical structure' that can capture the mathematics relevant to *both* of these related tasks.

notions from neighbouring disciplines such as computer science (for instance classifications of computational performance). It is one thing to suggest that we can easily capture things like symmetry groups, graphs, or number lines in a 'domain and relations' structure, and quite another thing entirely to suggest that we can do the same for things like spectral methods for dealing with systems of PDEs or homogenization techniques for treating porous media.

Suppose, however, that this difficulty proves surmountable. That is, suppose that we manage to catalogue all of the pieces of mathematics employed in our application and associate each with a 'domain and relations'-type structure. What is the 'mathematical structure' into which we eventually map? Can we simply frame a structure by taking something like the union of all these individual structures and mapping from some construal of our target system into that? Recall that one of the constant refrains of our presentation of the ice permeability case was that the utility or viability of applications of mathematics often turns on our ability to forge meaningful connections between the various pieces of mathematics we employ. It is not enough to know *how* to model our system using the lattice percolation model on the one hand and the compressed powders model on the other. We must understand how it is that the latter allows us to apply a result that we can only obtain within the framework provided by the former. That is to say that we need to be understand the relationships between the various structures in our catalogue.

The second difficulty presented by cases such as our ice permeability model is that, as we stressed in §5, the relations that allow us to employ various pieces of mathematics in concert are often framed in empirical terms and require empirical input. We cannot always construe in purely mathematical terms the way that one model serves to refine, correct, or fill-in another. Our lower scale lattice model is only able to relay information to our larger scale climate model by way of a process of homogenization which appeals to both mathematical features of our original equations and empirical features of the system. We can use our compressed powders model to fill in details of the lattice model because we can apply sophistical computational techniques to high resolution tomographical scans and thus approximate the relationships between the parameters that appear in either model. It is only in this way that the compressed powder model can help us to determine the approximate value of the parameter k_0 which appears in the lattice model.

It is difficult to see how the kinds of 'mathematical structures' to which the mapping account appeals can handle these kinds of combined empirical-mathematical connections between the pieces of mathematics involved. Recall that Nguyen and Frigg suggest that the only kind of relationships between pieces of mathematics that such a notion of structure allows into the picture are those that arise out of the structural arrangements specified by the extensions of the various relations on our domain. Such an extensional characterisation might allow us to consider relations such as "is greater than" or "is a factor of", but as we have repeatedly emphasised there is an important difference between these and relations such as "equation (2) is an effect homogenization of equation (1)" and "parameter μ can be computationally aligned with parameter λ for the purpose of calculating the value of κ". While we can determine whether relations of the first kind hold by looking simply at the way that the elements of our mathematical structure are arranged, whether relations of the second kind hold depends *both* on the mathematical features of the relata *and* additional empirical features of our target system. There does not seem to be any room for relations of the second kind within the kind of structure countenanced by proponents of the mapping account.[21]

The upshot of all of this is that the key assumption of the mapping account does not look particularly plausible in light of the more complex example we have examined. The problem, at base, seems to be the notion of 'mathematical structure' that lies at the heart of the mapping account. If we we want our structure to meet the requirements we laid out in §4, then it will need to in some way account for the more complex dynamics that we saw in our sea ice case. Given these requirements, the MATHEMATICAL STRUCTURE ASSUMPTION claims that there exists a mathematical structure that can capture (a) all of the mathematics relevant to the application and (b) the relationships between these pieces of mathematics.

In our sea ice case, (a) requires that the structure must capture not just the mathematics used to frame the models involved but also the computational and quasi-experimental methods used in working with those models, while (b) requires that the structure be able to capture mixed mathematical-empirical relationships between the pieces of mathematics that are integral to the success of the application. It does not seem that a mathematical structure of the bare 'set and relations' variety outlined by Nguyen and Frigg is capable of tackling (a) and (b). Recall that the MATHEMATICAL STRUCTURE ASSUMPTION claims that the 'right kind of structure' exists given any application of mathematics in empirical science. What emerges from the above then is that if 'right kind of structure' means that the structure is of the bare 'domain and relations' variety described by proponents of the mapping account *and* meets the two requirements outlined in §4, then such an assumption does not look particularly plausible in light of cases such as our ice permeability model.

7 The Dynamics of the Mapping Account

It is worth considering a particular kind of reply to the line of argument developed above. Recall that Bueno and Colyvan's inferential conception suggested that the application of mathematics in empirical science involved several different steps, underpinned by two separate mappings. The proponent of a

[21]The point we develop here suggests a way of responding to the way that Bueno and French (2018) dismiss the criticism of mapping accounts found in Batterman (2010). Batterman suggests that scientific explanations often rely on asymptotic features of the mathematical description of a system for which there is no physical analog. He further suggests that the mathematical operations that allow us to understand these asymptotic features are not the sort of thing that can be captured in a 'mathematical structure'. Bueno and French simply reply that insofar as these operations are run-of-the-mill mathematical operations, they "can be characterized set-theoretically and hence represented within our framework" (Bueno and French, 2018, 187). One way of paraphrasing Batterman's point in the terms of this paper is that the limiting relationships on which such asymptotic relationships rely are not *purely mathematical* relationships. It is not because they are *operations* that Batterman suggests they cannot be captured in the kind of structure required by the mapping acount, but rather because we require both mathematical and empirical input to establish that the limiting relationship exists.

mapping account might suggest that the kind of difficulties we have discussed can be handled by a more complex series of mapping steps. Rather than expecting the bare mathematical structure itself to account for the thornier aspects of some applications, we can instead construe such applications in terms of a more complicated arrangement of immersion, derivation and interpretation steps. In short, such a response argues that there is nothing wrong with the MATHEMATICAL STRUCTURE ASSUMPTION, and that the key to understanding cases such as the one we have presented lies in construing them as a *series* of applications, each employing a different structure. In this section, we will examine whether various ways of complicating the dynamics of the mapping account can address the problems raised in the previous section.

7.1 Iterated Mappings

One way of complicating the dynamics of the mapping account might be suggested by a remark made by Bueno and Colyvan in their original paper. Bueno and Colyvan suggest that sometimes, our applications might require multiple immersion steps. That is, we might immerse our empirical setup in some mathematical structure, and then in turn immerse that mathematical structure in another mathematical structure. To illustrate this, Bueno and Colyvan consider a case in population ecology. We might map from some population of rabbits to the natural numbers in order to make statements about the size of the population at some particular time. However in investigating this behaviour we might want to make use of some results about differential equations, and so we may map our natural number structure into the real numbers. Having availed ourselves of this second structure, we can easily map back into the natural numbers (by rounding off in some way) and interpret the result naturally from there.

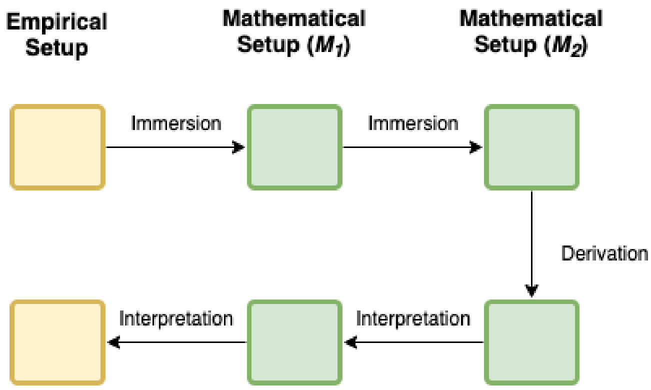

Figure 4: Iterated Mappings.

Of course, not all situations are so simple. We may find that there is no obvious way to interpret results found in further immersion steps as physically significant. This is no matter, however. All we require, Bueno and Colyvan suggest, is "an invertible mapping that is a *conservative extension* of the mapping used at the immersion stage" (Bueno and Colyvan, 2011, 369). Since such a mapping will agree with the immersion mapping on all cases, we will be able to rely on the interpretation mapping from the first domain to our physical system to provide us with a physical interpretation of any non-problematic cases.

The proponent of a mapping account might then suggest that the complexity involved in our sea ice case can be dealt with in a similar way. We start by mapping our system to some set of PDEs and then we map from there into some other domain, and so on and so forth until we have all the information we need. We then simply interpret our way back out through the various mathematical domains we passed on the way in (since the immersion mappings are each conservative extensions in turn) until we can interpret finally from our original PDEs to our physical system once more. In our terms, we might instead accommodate some case by providing a *series* of structures, connected by iterated immersion steps. This *series* of structures would satisfy SCOPE if all the relevant mathematics appeared in at least *some* structure and STITCHING if the relevant relationships between those pieces of mathematics were captured either within some structure or in the mappings between structures.

Promising as this approach may seem, it does not really succeed in avoiding the difficulties faced by the simpler picture. In particular, we must still find some structure for each piece of mathematics and we must still account for the empirical connections established between the pieces of mathematics employed.

Let's begin with the first challenge. Supposed that we take the first immersion step to somehow map our sea ice to the continuum scale PDEs that govern the temperature field inside the sea ice. Recall that the next step in the general approach involved employing homogenization techniques in order to replace the parameter for averaged velocity with an effective fluid permeability tensor. What is the mathematical structure to which this second immersion step might take us? Homogenization theory does not quite exist as its own self-contained fiefdom of mathematics, but rather combines limiting operations with computational methods specific to the nature of the medium. In this sense it is more like a loose collection of methods and results sensitive to the context of application (for instance, opportunities provided by the low Reynolds number flow of our ice-brine composite). This allows us to, if we are careful and lucky, reliably approximate rapidly varying parameters with a better behaved effective parameter. It is not clear that there is an appropriate way to represent homogenization methods so employed as a structure in the way required.

Suppose that in spite of these difficulties, we succeed in finding a singular relevant structure for the second immersion step to map into. We are still left with the lower scale analog of STITCHING: saying something about how these pieces fit together. In setting up some partial differential equation model we almost always at first ignore a variety of physical processes that we must later take into account somehow. Often these physical processes take place at a variety of scales and so understanding the behaviour of our PDE in light of them is a complicated matter. In practice, when we introduce new mathematical tools into our model we are often correcting for precisely these initially overlooked processes and so require not simply a mathematical bridge between our tools but (as we have stressed) new physical information to make sense of this connection.

While our lattice percolation model provided a helpful way of characterising our brine flows, experimental discrepancies alerted us to the fact that "key features of the geometry of the brine microstructure were being missed by lattices" (Golden, 2015, 700). Taking account of these key features involved making use of cutting edge imaging techniques to examine the microstructure of the sea ice in such a way that we could more or less brute force the scaling factor required in the lattice model by a combination of numerical techniques. Once again, the task of constructing a thread through our assembled mathematical tools demands that we not only exploit their mathematical relations but import new physical information as well.

We can see how sharp a challenge this dynamic poses to the mapping account if we recall Bueno and Colyvan's suggestion that all we require for further immersion stages is "an invertible mapping that is a *conservative extension* of the mapping used at the immersion stage" (Bueno and Colyvan, 2011, 369). On such a picture, we immerse our physical system in some mathematical structure and until it comes time to interpret our results we are engaged solely in the task of exploiting mathematical relations of various kinds. Within this picture, we can parse the conservative extension requirement as the demand that any new mathematical tools introduced not strictly speaking tell us anything new relative to the original bridge we established between our physical system and mathematical domain.

This would appear to leave no room, however, for precisely the dynamic that we have described. If we require, as we have seen in our ice melt case, that we import fresh physical information in order to understand the connection between newly introduced mathematical tools, then in fact we *require* that the new machinery tell us about processes we have previously ignored. We exploit numerical methods and imaging techniques applied to the compressed powders model to allow us to correct for the microscale processes ignored by our lattice-related tools. The model-theoretical demand that further immersion be conservative extensions is not in any clear sense met in cases like these, and that is precisely the point!

7.2 Sequential Mappings

If one of the problems with the iterative approach is that it makes it difficult to see how fresh empirical information can enter the picture, the structuralist might instead try to capture the kinds of processes described in our case study in terms of a *sequence* of mappings. In this sense, we might think that our case is better described as a sequence of applications of mathematics. According to this response, our scientists perform the three steps outline by Bueno and Colyvan, employing the appropriate mappings, and thereby learn *something* about the physical system in question. With this augmented understanding of our physical system, we may then map once again from our physical system into some (possibly different) mathematical domain, carrying with us some additional information if required. We may then follow the same process as before, continuing to learn about our system.

That is to say that the task of characterising the behaviour of our system need not be confined to one set of mappings. We can progressively learn more and more about our system by way of mappings between its physical structure and the mathematical machinery appropriate to different scales and dynamics. Eventually, so the response goes, we will begin to develop a picture of the behaviour of our system as a whole.

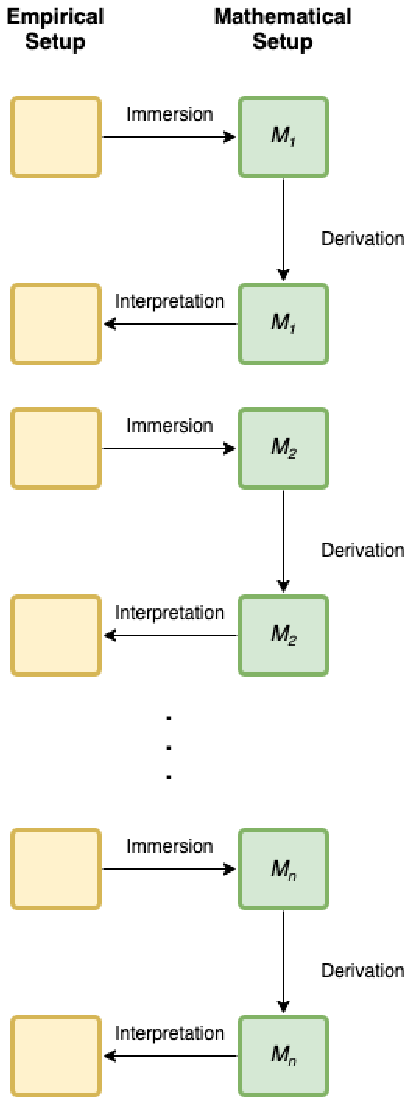

Figure 5: Sequential Mappings.

The difficulty here is that in many cases, we are able to develop an improved picture of our system's behaviour only by stitching together the heterogeneous mathematical representations that we employ at different levels. In the terms of our sequential strategy, we need to link together not just the physical information that we extracts but also features of our various pieces of mathematics which otherwise live in different sequential mapping situations. For instance, using the compressed powder model to fill in details overlooked by the lattice model, we do not simply map into some geometric domain and back in order to obtain some physical information p_1 (equation (3)) and then map into some other domain and back in order to obtain some other pieces of physical information p_2 (the permeability scaling factor). Our ability to extract this p_2 relies not just on p_1 considered as a physical fact about our system interpreted in some new mathematical domain but rather on intricate details to do with the mathematical process by way of which we were able to obtain p_1 in the first place.

Similar considerations mean that such a sequential mapping approach is ill-suited to capturing the interactions between the behaviours our system displays at different scales. We might think that we can simply represent our models at different scales as separate mapping applications, from which we extract pieces of physical information that we can combine in order to provide some more general picture of our system's behaviour. But as above, the problem is that a proper understanding of how systems behave across scales rarely involves the passing of purely physical information between different models. In practice, we can fruitfully combine mathematical representations at different scales often only by by exploiting mathematical relationships that obtain between them and by using fresh empirical data to facilitate communication between these representations.

Where the iterative strategy made it very difficult to see how the required fresh physical information is able to enter into the picture at key moments, the sequential strategy makes it hard to see how our heterogeneous mathematical domains can interact with one another. If we wipe the mathematical slate clean with every new mapping, how can our models interact with one another *mathematically*? Leaving aside the issues regarding the SCOPE requirement which affect the sequential approach as much as the iterative one, it is hard to see how this kind of response is able to meet our STITCHING requirement given that the various mathematical tools we employ are not really made to work together in any meaningful sense.

7.3 A Hybrid Approach

To summarise the above, we might say: iterated mappings may (with serious caveats regarding SCOPE) help us to get more mathematics into the picture as we need it, but we face problems bringing fresh empirical information into the picture. On the other hand, sequential mappings allow for fresh empirical information to be brought to bear but close off our various mathematical tools from interacting with one another. The natural next thought might be: why wouldn't some *hybrid* approach work? We could accommodate our sea ice application in terms of a mix of iterated and sequential mappings, utilizing mathematical connections between tools and empirical connections between levels of description as appropriate. When we need fresh mathematical machinery, we iterate, and when we need fresh empirical information, we begin a new sequence.

Yet as we saw above, our various mathematical tools are stitched together in ways that rely on *both* mathematical connections and empirical information. One of the constant refrains of this paper has been that the investigation of complex scientific problems often demands that mathematical and empirical concerns be intertwined in intricate ways. The connections between our various mathematical tools often only make sense in light of fresh empirical detail, evidenced by the fact that we required the use of tomographical data and computational techniques to forge the right kind of connection between our compressed powders model and our lattice model. Additionally, the connections between empirical details of our model must often be mediated according to the mathematical tools employed in securing them, as in the fact that the information we obtain about our brine microstructure can only be accommodated in our climate models by way of the relevant *homogenization limit*. A hybrid approach allows us a limited capacity to bring new mathematics into play and the ability to return to the drawing board in search of fresh physical information, but it does not appear to allow these processes to be intertwined in any meaningful way. Leaving aside the challenge of finding some single structure within which we can house our various iterated or sequential mappings (i.e. SCOPE), any approach which leaves no room for the way that mathematics and empirical data combine to secure reliable scientific results fails to satisfy our STITCHING requirement.

That is all to say that even this hybrid approach has trouble meeting our STITCHING requirement in more complicated circumstances. The problem lies in the assumption that the relationships between the different tools we assemble after our first foray into some mathematical domain can be captured in purely mathematical terms. As we have seen, a great deal of physical information originally ignored must often be introduced in order to articulate how these tools relate to one another. In fact, such new tools are often introduced precisely *because* discrepancies at one level have demanded that we find some way to account for information previously ignored. In short, it does not look as though we can save the MATHEMATICAL STRUCTURE ASSUMPTION by construing cases such as the one we have examined in terms of a series of sequential, iterated, or hybrid mapping applications.

8 Concluding Remarks

Where does all of this leave the mapping account? The problem, we have suggested, lies in the core notion of 'mathematical structure' around which it is framed. One way of putting the upshot of this paper is this: in more complicated cases, it is not at all clear that mathematical structures of the bare 'set and relations' variety are able to meet two requirements we might plausible like them to, given the explanatory ambitions of the mapping account. Applying mathematics to the physical world involves more than simply finding the right way to associate parts of some physical structure with elements in some mathematical structure and letting a deductive inferential procedure unfold and doing our best to interpret what comes out the other side. The inferential procedures that unfold once we have managed to find some appropriate mathematical model of our physical system often rely on fresh empirical data and mathematical manipulation in equal measure in a way that seems difficult to capture within the bare formal landscape in terms of which the mapping account is framed.

We might ask: if the mapping account does not work in more complex cases, then how ought we to conceive of the application of mathematics in empirical science? Of course, I can hardly present a fully formed account here, but the problems we have outlined may hold some clues to answering this question. The main difficulty for mapping accounts as I have argued is the fact that the complexities of real life application often require that we employ a variety of pieces of mathematics, the relationships between which are mediated by thoroughly empirical considerations. However we are to understand applied mathematics, then, it cannot be in such a way that our mathematical tools are siloed off from the empirical considerations that must inevitably inform their use.

To frame a broad suggestion: perhaps the lesson to learn from the difficulties encountered by the mapping account may be that mathematics plays *other* roles in empirical application than the *inferential* one highlighted by proponents of the mapping account. To be sure, mathematics does often provide us with a domain of structures we can use to represent physical systems at some level of generality and within which inferences can be easier to obtain. It also allows us to develop physically-informed techniques for correlating variables and parameters in different mathematical models, find effective parameters by exploiting patterns in the physical behaviour of a system along with facts about its geometry, develop broad recipes for applying computational methods to certain kinds of empirical systems, and so on. Most importantly, intermingling empirical considerations with mathematical ones is precisely what allows us to understand how all of these individual uses of mathematics can coalesce into a helpful treatment of some system.

It may be true that the various pieces of mathematics that play the roles canvassed above can be captured in isolation as the kind of extensionally-defined structures favoured by proponents of the mapping account. Yet as we have seen this does not necessarily shed light on the *role* that they play in empirical application. If it is true, as I have argued, that mapping accounts cannot capture the intermingling of mathematical and empirical considerations characteristic of more complex scientific cases, then the upshot of this may be that we require a broader appreciation of the various roles played by mathematics in empirical application.[22]

At the very least, I hope to have shown that properly evaluating the plausibility of mapping accounts of applied mathematics will involve venturing beyond the restricted class of examples typically featured in the literature. Modern scientific practice is increasingly characterised by multiscalar modelling, computational and numerical methods and the use of a heterogeneous array of mathematical tools. We should not accept any account of applied mathematics as satisfying before we have given detailed consideration to its ability to handle applications of this kind. In considering one such example, I hope to have shown that the difficulties encountered by such accounts stem not from complexities parochial to each account but rather from the common way in which such accounts invoke the notion of mathematical structure.

[22]In fact it seems to me that both Batterman (2010) and Kasirzadeh (2021) are efforts in more or less this vein.

References

Robert W. Batterman. On the Explanatory Role of Mathematics in Empirical Science. *British Journal for the Philosophy of Science*, 61(1):1–25, 2010. doi: 10.1093/bjps/axp018.

Robert W. Batterman. The Tyranny of Scales. In Robert W. Batterman, editor, *The Oxford handbook of philosophy of physics*, pages 255–286. Oxford University Press, 2013.

Robert W. Batterman. *A Middle Way: A Non-Fundamental Approach to Many-Body Physics*. Oxford: Oxford University Press, 2021.

Theodore Bergman, Frank Incropera, David DeWitt, and Adrienne Lavine. *Fundamentals of Heat and Mass Transfer*. Wiley, 2011.

Otávio Bueno. An Anti-Realist Account of the Application of Mathematics. *Philosophical Studies*, 173(10):2591–2604, 2016. doi: 10.1007/s11098-016-0670-y.

Otávio Bueno and Mark Colyvan. An Inferential Conception of the Application of Mathematics. *Noûs*, 45(2):345–374, 2011.

Otávio Bueno and Steven French. *Applying Mathematics: Immersion, Inference, Interpretation*. Oxford: Oxford University Press, 2018.

Julia Bursten. Conceptual Strategies and Inter-Theory Relations: The Case of Nanoscale Cracks. *Studies in the History and Philosophy of Modern Physics*, 62: 158–165, 2018.

Julia R. S. Bursten. Classifying and Characterizing Active Materials. *Synthese*, 199(1):2007–2026, 2021. doi: 10.1007/s11229-020-02870-2.

Nicolas Fillion. Conceptual and Computational Mathematics†. *Philosophia Mathematica*, 27(2):199–218, 2019. doi: 10.1093/philmat/nkz005.

Nicolas Fillion and Robert M. Corless. On the Epistemological Analysis of Modeling and Computational Error in the Mathematical Sciences. *Synthese*, 191(7):1451–1467, 2014. doi: 10.1007/s11229-013-0339-4.

Kenneth M. Golden. Percolation Models for Porous Media. In Ulrich Hornung, editor, *Homogenization and Porous Media*, pages 27–43. Springer, 1997.

Kenneth M. Golden. Climate Change and the Mathematics of Transport in Sea Ice. *Notices of the American Mathematical Society*, 56(5):562–584, 2009.

Kenneth M. Golden. The Mathematics of Sea Ice. In Nicholas J. Higham, editor, *The Princeton Companion to Applied Mathematics*, pages 694–705. Princeton University Press, 2015.

Kenneth M. Golden, Amy L Heaton, Hajo Eicken, and Victoria I. Lytle. Void bounds for fluid transport in sea ice. *Mechanics of Materials*, 38(8):801–817, 2006.

Sara Green and Robert W. Batterman. Biology Meets Physics: Reductionism and Multi-Scale Modeling of Morphogenesis. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 7161:20–34, 2017. doi: 10.1016/j.shpsc.2016.12.003.

Atoosa Kasirzadeh. A New Role for Mathematics in Empirical Sciences. *Philosophy of Science*, 88(4):686–706, 2021. doi: 10.1086/713906.

Ross C. McPhedran. Effective Medium Theories. In Nicholas J. Higham, editor, *The Princeton Companion to Applied Mathematics*, pages 500–505. Princeton University Press, 2015.

James Nguyen and Roman Frigg. Mathematics is Not the Only Language in the Book of Nature. *Synthese*, pages 1–22, 2017. doi: 10.1007/s11229-017-1526-5.

Christopher Pincock. A New Perspective on the Problem of Applying Mathematics. *Philosophia Mathematica*, 12(2):135–161, 2004a.

Christopher Pincock. A Revealing Flaw in Colyvan's Indispensability Argument. *Philosophy of Science*, 71(1):61–79, 2004b.

Christopher Pincock. A Role for Mathematics in the Physical Sciences. *Noûs*, 41(2):253–275, 2007.

Christopher Pincock. On Batterman's 'On the Explanatory Role of Mathematics in Empirical Science'. *British Journal for the Philosophy of Science*, 62(1): 211–217, 2011. doi: 10.1093/bjps/axq025.

Davide Rizza. The Applicability of Mathematics: Beyond Mapping Accounts. *Philosophy of Science*, 80(3):398–412, 2013. doi: 10.1086/670295.

Fadil Santosa. Homogenization. In Nicholas J. Higham, editor, *The Princeton Companion to Applied Mathematics*, pages 103–105. Princeton University Press, 2015.

Stewart Shapiro. *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press, 1997.

Noah Stemeroff. Structuralism and the Conformity of Mathematics and Nature. *Studies in History and Philosophy of Science Part A*, 86:84–92, 2021. doi: 10.1016/j.shpsa.2021.01.004.

Endre Süli. Numerical Solution of Partial Differential Equations. In Nicholas J. Higham, editor, *The Princeton Companion to Applied Mathematics*, pages 306–319. Princeton University Press, 2015.

Salvatore Torquato. *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Springer, 2002.

Mark Wilson. *Physics Avoidance: And Other Essays in Conceptual Strategy*. Oxford University Press, 2017.

Eric Winsberg. Handshaking Your Way to the Top: Simulation at the Nanoscale. *Philosophy of Science*, 73(5):582–594, 2006. doi: 10.1086/518328.

Eric B. Winsberg. *Science in the Age of Computer Simulation*. University of Chicago Press, 2010.