



UNIVERSITY OF LEEDS

A Robustly Inferential Conception of the Role of Mathematics in Science



Colin Matthew McCullough-Benner

University of Leeds

School of Philosophy, Religion and History of Science

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

December 2022

Intellectual Property Statement

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Colin Matthew McCullough-Benner to be identified as Author of this work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1988.

© 2023 The University of Leeds and Colin Matthew McCullough-Benner.

Previously Published

Chapter 4 is based on McCULLOUGH-BENNER, C. (2019). Representing the world with inconsistent mathematics. *British Journal for the Philosophy of Science*, 71, 1331–58.

Section 5.2 and parts of section 5.1 are based on McCULLOUGH-BENNER, C. (2022a). Applying unrigorous mathematics: Heaviside’s operational calculus. *Studies in History and Philosophy of Science Part A*, 91, 113–24.

Chapter 6 is based on McCULLOUGH-BENNER, C. (2022b). The metarepresentational role of mathematics in scientific explanations. *Philosophy of Science*, 89, 742–60.

Acknowledgements

Thanks to my supervisors Steven and Juha, as well as my former supervisor Marcus Rossberg, for all your work, patience, and support over the years.

Thanks to my colleagues for helpful philosophical discussions—as well as unhelpful, non-philosophical ones.

Thanks to the School of Philosophy, Religion, and History of Science at the University of Leeds for the scholarship that made it possible for me to come to Leeds for a PhD.

Thanks to my family, especially my parents, for your unwavering support (and for never asking “What are you going to do with *that?*”). Lots of parents encourage their children to follow their dreams, but few remain so encouraging when that dream turns out to be a PhD in philosophy.

Special thanks to my wife and partner Rasa, without whom I could not have finished this thesis.

Extra-special thanks to my daughter Eglut , without whom I could have finished this thesis much sooner. I wouldn’t change a thing.

Abstract

In the thesis, I develop and defend a novel account of mathematical scientific representations, the Robustly Inferential Conception (RIC). According to RIC, mathematics places constraints on what a representation's target system must be like by specifying inferences that must preserve truth if the representation is accurate in all respects. A mathematical scientific representation is treated as having three ingredients:

(RIC1) A partial physical interpretation of the language in which the relevant mathematics is expressed,

(RIC2) An initial description of the target system in this physically interpreted mathematical language, and

(RIC3) A set of mathematical inference patterns licensed by the relevant mathematics.

The informational content of the representation, according to RIC, is given by the closure of the statements in RIC2 under the inference patterns in RIC3, under the interpretation RIC1.

I argue that RIC has three significant advantages over its most prominent alternative, the mapping account. First, it can be applied in a wider range of cases. Second, it is more successful as a meta-level representational device to be used by philosophers of science to represent philosophically salient features of scientific practices in which mathematics is applied. Third, it makes fewer assumptions about the nature of mathematics.

I spend most of the thesis substantiating the second of these points, which I take to be the strongest case for RIC. In particular, I show how RIC can be used to shed light on philosophical issues concerning applications of inconsistent and otherwise unrigorous mathematics and the role of mathematics in scientific explanations and scientific understanding.

CONTENTS

1	Introduction	1
1.1	The many problems of the applicability of mathematics	2
1.2	Why focus on representation?	4
1.3	The plan	6
I	Representation	13
2	Scientific Representation and Mathematical Scientific Representations	14
2.1	Accounts of scientific representation	15
2.2	From scientific representation to mathematical scientific representation	21
2.3	How to argue for an account of mathematical scientific representation	24
2.4	Varieties of the mapping account	26
2.5	Arguments against mapping accounts and the need for an alternative	29
2.5.1	General arguments	29
2.5.2	Particular arguments and the need for an alternative to mapping ac- counts	31
3	The Robustly Inferential Conception of Mathematical Scientific Representations	34
3.1	The robustly inferential conception	35
3.1.1	RIC1: Physical interpretation	37

3.1.2	RIC2: Initial description of the target system	38
3.1.3	RIC3: Inference patterns	40
3.1.4	RIC as a generalization of mapping accounts	44
3.2	An initial case for RIC	46
3.2.1	RIC is applicable in a wider range of cases	47
3.2.2	RIC does better as a meta-level representational device	49
3.2.3	RIC makes fewer assumptions about the nature of mathematics	51
3.3	Scientific representation and RIC	53
3.3.1	RIC as the only account in keeping with the motivations of deflationary inferentialism	53
3.3.2	RIC and the motivations behind substantive accounts of scientific representation	57
3.4	Conclusion	62
II	Case Studies	66
4	Applying Inconsistent Mathematics: The Early Calculus	67
4.1	The Early Calculus	68
4.2	Mapping Accounts and the Early Calculus	72
4.2.1	Partial structures	73
4.2.2	Inconsistent structures	80
4.2.3	Related total consistent structures	82
4.3	A Robustly Inferential Account of the Early Calculus in Applications	84
4.4	Beyond Inconsistent Mathematics	90
5	Applying Unrigorous Mathematics: Heaviside's Operational Calculus and Path Integrals in Quantum Physics	94
5.1	Unrigorous mathematics in general	96

5.1.1	Inferentially permissive and inferentially restrictive methodologies . . .	96
5.1.2	Inferentially restrictive methodologies and accounts of applications of mathematics	99
5.2	Heaviside’s operational calculus	108
5.2.1	Heaviside’s operational calculus and resistance operators	110
5.2.2	Failures of rigor and inferential restrictions	116
5.2.3	Failures of rigor and “physical mathematics”: The physical demand for fractional differentiation	122
5.2.4	The Laplace transform in heavy disguise?	127
5.2.5	Conclusion	133
5.3	Path integrals in quantum physics	133
5.3.1	Path integrals in quantum mechanics	136
5.3.2	The path integral beyond quantum mechanics	145
5.3.3	Conclusion	151
5.4	Conclusions: The scope of an account of mathematical scientific representation	152
5.4.1	Formalisms and inferentially restrictive methodologies	152
5.4.2	Are these applications of <i>mathematics</i> ?	155
III Explanation and Understanding		158
6 The Metarepresentational Role of Mathematics: Mathematical Scientific Explanations in the RIC Framework		160
6.1	Introduction	161
6.2	Representation and Metarepresentation	165
6.2.1	Representation	165
6.2.2	Metarepresentation	168
6.3	Explanatory Generality	170

6.3.1	Scope-generality	170
6.3.2	Topic-generality	171
6.4	Example: The Prime Cycles of Periodical Cicadas	173
6.5	Metarepresentation and Scope-Generality	175
6.6	Metarepresentation and Topic-Generality	179
6.7	Explanatory Depth	183
6.8	Conclusion	186
7	Inferentialism, Rigor, and Understanding	187
7.1	Introduction	188
7.2	Understanding and unrigorous mathematics	191
7.2.1	Scientific understanding	191
7.2.2	Unrigorous mathematics	196
7.3	Case studies	198
7.3.1	Unrigorous infinitesimals	198
7.3.2	Heaviside's operational calculus	204
7.4	Implications and further work	210
8	Conclusion	212
8.1	Contributions of the thesis	213
8.2	Directions for further work	214
8.2.1	RIC, computational inference tools, and data-driven science	214
8.2.2	RIC and mathematics as such	221
	References	224

LIST OF FIGURES

5.1	A DC circuit with a resistor and inductor in series, as treated in [Heaviside, 1899 , §283, pp. 129f].	114
8.1	An artificial neural network with two hidden layers.	217

CHAPTER 1

Introduction

1.1 The many problems of the applicability of mathematics

Like many of the most interesting problems, the applicability of mathematics in science is at once mundane and perplexing. On the one hand, mathematical representations of phenomena of interest are the bread and butter of modern science. On the other, an air of mystery surrounds the contributions of mathematics to the remarkable success of much of the science built on these representations. Nobel Prize-winning physicist Eugene Wigner famously went so far as to conclude that “The miracle of the appropriateness of the language of mathematics for the formulation of the laws of physics is a wonderful gift which we neither understand nor deserve” [Wigner, 1960, p. 14].

The miracle, according to Wigner, is that many of the most useful mathematical concepts in physics have arisen through the pursuit of pure mathematics for its own sake. While many basic concepts in mathematics—like those involved in arithmetic or Euclidean geometry—might be understood as generalizing some feature of our experience, the concepts from higher mathematics that played a crucial role in the enormous advances in physics in the twentieth century were developed to solve seemingly unrelated problems in pure mathematics. And the choice of problems and concepts in pure mathematics is driven by aims that are not *prima facie* aligned with those of the other sciences. [cf. Steiner, 1998].

On the opposite end of the spectrum is a range of cases that are at least as mysterious as those Wigner has in mind but that have received far less philosophical attention. These are cases in which physicists have developed and used bespoke mathematical tools that fall well short of the standards of rigor normally applied to pure mathematics, like the development of the early, inconsistent infinitesimal calculus in part for Newtonian physics or the development of wildly unrigorous path integral techniques for quantum field theory. As Urquhart [2008a, p. 410] puts it, “the methods that [physicists] use are frequently so far from normal mathematical practice that it is sometimes not clear that the objects [they appeal to] themselves are even mathematically well-defined.” Given that these methods look in all other ways like normal

1.1 The many problems of the applicability of mathematics

applications of mathematics, just with what seem to be much weaker epistemic credentials, their critical role in some of the most successful historical and contemporary work in physics also has an air of the miraculous.

Beyond these “miraculous” cases, understanding how mathematics contributes to the success of science requires us to answer a number of more fine-grained questions concerning contributions of mathematics to particular features of science, including not just representation, but also confirmation, explanation, understanding, and unification, among others. Answering these questions in turn supports a deeper understanding of the epistemology of science.

Such an understanding of the contributions of mathematics to the success of science also has the potential to shed light on the nature of mathematics itself. Whatever our overarching account of the nature of mathematics might be, a core feature of mathematics is that it can be applied. A good philosophical account of mathematics therefore ought to support our best understanding of how mathematics is applied. This is evident, for example, in the enormous literature on indispensability arguments, which support platonism on the grounds that it falls out of our best understanding of the application of mathematics in science.

To attempt to answer all of these questions in the thesis would be overly ambitious, particularly at the level of grain necessary to do justice to the complexity and diversity of relevant scientific practice. (I suspect this would be a stretch even in a monograph!) Instead, in this thesis I develop and argue for a framework for making sense of episodes in which mathematics is applied based on a novel, inferentialist account of mathematical representations in science. I then show how this framework can be applied to shed light on some of the questions considered so far, including the fruitfulness of inconsistent and otherwise unrigorous mathematical techniques, the role of mathematics in scientific explanations, and the contributions of mathematics to scientific understanding.

1.2 Why focus on representation?

I have chosen to build the thesis around an account of mathematical representation for three reasons. First, mathematics generally appears in science in the context of a mathematically specified scientific representation, either as part of the language specifying the representation, as a device to make inferences using the representation, or as a way of relating different representations to one another—for instance, deriving one from another. So our understanding of how mathematics plays these roles in relation to mathematical scientific representations naturally must be part of the answer we give to any of the questions considered so far.¹ An account of mathematical scientific representation therefore has the potential to significantly constrain—for better and for worse—what one can say in response to a wide range of questions about the applicability of mathematics.

Second, given the centrality in science of representational practices in general and mathematical representational practices in particular, accounts of representation (and of mathematical scientific representations in particular) are well placed to serve as meta-level representational devices to be used by philosophers to represent episodes from scientific practice in a way that brings out their philosophically salient features. For example, the partial structures approach to scientific representation, which will be discussed at length in the next chapter, represents the informational content of scientific representations in terms of partial structures, model-theoretic structures in which not every item in the domain must be in the extension or antiextension of every predicate. The content of a scientific representation is what must be true of the world for an appropriate morphism to exist between some partial structure or collection thereof and the structure of the phenomenon represented. One can

¹In fact, we can say a bit more than this. As [Pincock \[2012, p. 3\]](#) observes, on most philosophical accounts of the aims of science, realist and anti-realist alike, the success of science has something to do with the accuracy of scientific representations; these views typically differ instead in how they specify conditions for success in terms of the features with respect to which representations must be accurate and the precise sense in which the representation must be accurate with respect to those features. For each such view, the question of how mathematics contributes to the success of science in some case or range of cases is at least partly to be answered by determining how the use of mathematics contributes specifically to the *accuracy*—in the relevant sense and with respect to the relevant features—of the relevant representation.

then explain various practices involving a particular representation in terms of features of partial structures. Scientists' inferences can be understood to be licensed by the posit that the relevant morphism exists in the given case. Scientists' uncertainty about some aspects of the representation can be represented in terms of partial structures in which some domain elements are in neither the extension nor the antiextension of some predicates. The tools of model theory, adapted to partial structures, can then be used to reason further about these practices. In this way, competing accounts of representation might have the same expressive capabilities, in the sense that the range of truth-conditional content they can ascribe to scientific representations is the same for both accounts. Such accounts might nonetheless differ significantly in terms of their utility as meta-level devices for representing the philosophically salient features of scientific practice.

Finally, I think the space of possible accounts of mathematical scientific representations has been woefully underexplored. With very few exceptions, accounts appealed to in the literature are versions of what [Pincock \[2004\]](#) christened the *mapping account*. According to the mapping account, mathematics can be used to represent phenomena in the world via the positing of a structure-preserving mapping between the relevant mathematical structure and the structure of the relevant phenomenon. The mapping account certainly is plausible, and it does permit several kinds of variation: in the structures allowed, in the mappings allowed, in whether anything further must be built in (as, e.g., in the inferential conception proposed by [Bueno & Colyvan \[2011\]](#)), and so on. But given the centrality of mathematical representation to other philosophical issues concerning the role of mathematics in science, and given the potential of an account of mathematical scientific representation to serve as a useful meta-level device for representing the philosophically salient features of episodes in which mathematics has been applied, the payoff of exploring alternative accounts of mathematical representation is potentially very high.

1.3 The plan

In chapter 2, I lay the groundwork for the novel account of mathematical scientific representation that forms the core of the thesis. I begin by discussing accounts of scientific representation in general. I pay special attention to the distinction between inferentialist accounts, according to which scientific representations are defined by their ability to support surrogate inference about their target systems, and substantivist accounts, according to which there is a more substantial property common to all scientific representations that explains, among other things, how representations are able to support surrogate inference. This distinction plays an important role in the thesis because it is very closely related to the distinction between the mapping account and my own inferentialist account of mathematical scientific representations.

I move on to discuss mathematical scientific representations in particular, including how they relate to accounts of scientific representation in general in both their structure and their aims, why we need an account of mathematical scientific representation over and above our account of scientific representation in general, and in what respects our account of one of these two kinds constrains our account of the other. Building on this discussion, particularly of the aims of accounts of mathematical scientific representations, I discuss how to adjudicate between competing accounts of mathematical scientific representation, particularly when such accounts agree on the informational content they assign to scientific representations but differ in *how* they spell that content out—and so differ also in their role as meta-level devices for representing philosophically salient features of scientific practice.

In the rest of the chapter, I shift my attention specifically to the mapping account. I present and motivate several influential variations on this view, to which I refer back frequently in the rest of the thesis. I conclude by considering some existing criticisms of mapping accounts.

Chapter 3 is the philosophical core of the thesis. In it, I articulate and defend a novel account of mathematical scientific representation, the robustly inferential conception (RIC).

Unlike the mapping account, RIC represents the informational content of a mathematical scientific representation in terms of the inferences about the target system of the representation that must be truth-preserving if the representation as a whole is accurate in all respects. According to RIC, such a representation should be represented as having three ingredients:

- (RIC1) A (possibly incomplete) *physical interpretation* of the mathematical language used in the representation in terms of the target system. This might but need not be something like a model-theoretic interpretation, as in the the mapping account. What matters is just that such an interpretation suffices to provide at least some expressions in the interpreted mathematical language with physical truth conditions.
- (RIC2) An *initial description of the target system* in the interpreted mathematical language.
- (RIC3) A *class of privileged inference patterns* which is a subset of those licensed by the mathematical theory or theories from which the mathematical vocabulary in the representation is drawn, under its mathematical interpretation.

The informational content of the representation is then given by those claims in the closure of RIC2 under the patterns of inference in RIC3 that are assigned physical truth conditions by the interpretation RIC1.

After spelling out the view in greater detail, I make an initial case for RIC over mapping accounts on three grounds. First, RIC can accommodate a wider range of mathematical representations. Second, RIC does better as a meta-level device to be used by philosophers of science to represent philosophically salient features of scientific practice because it represents those features more perspicuously. Finally, RIC requires fewer assumptions about the nature of mathematics. I take the second of these to make the strongest case for RIC, and the rest of the thesis is largely dedicated to substantiating this point in relation to a range of case studies.

Finally, I turn to the relationship between RIC and accounts of scientific representation in general, arguing that there is good reason to adopt RIC regardless of one's account of scientific

representation. First, I argue that RIC is the best way to extend deflationary inferentialism about scientific representation to an account of mathematical scientific representations in particular while preserving the motivations of the view. Then I discuss two ways in which RIC can be reconciled with the motivations behind substantive accounts of representation. First, one might adopt a pluralist account of scientific representation, which reconciles the motivations of deflationary and substantivist views. Second, one can combine RIC directly with a substantivist account of scientific representation by treating RIC as describing not the target of the representation but rather its source.

In chapter 4, I discuss one sort of case that causes trouble for the mapping account but is usefully represented in terms of RIC: applications of inconsistent mathematics in general and of the early infinitesimal calculus in particular. The immediate problem inconsistent mathematics creates for mapping accounts is that no classical model-theoretic structure satisfies an inconsistent theory. As a result, if a mapping account is to accommodate applications of inconsistent mathematics, it must either incorporate a non-classical notion of structure or specify a plausible way to relate inconsistent theories to classical structures that supports the same explanatory work. I explore in detail how each of these strategies might be deployed and conclude that each leaves something to be desired. I consider two ways in which the notion of structure might be extended: to partial structures [e.g. [da Costa & French, 2003](#)] and to inconsistent structures [[Colyvan, 2009](#), p. 167]. I find that both accounts fail to differentiate all of the possible ways in which scientists might have tried to make physical sense of the inconsistent properties of naïve infinitesimals and that the formal apparatus of each approach obfuscates rather than clarifies the ways in which scientists, through a judicious choice of inference strategy, managed to use the early calculus to make useful inferences without deriving undesirable results. In contrast, RIC naturally captures these features of their practice and allows for a straightforward epistemic explanation of why those practices were reasonable. On the other hand, I also consider an approach that might tempt a proponent of a more conservative version of the mapping account, namely to explain representations using the

early inconsistent calculus in terms of structures picked out by the modern calculus. This approach concedes a great deal to RIC. In effect, it means that the mapping account is limited to providing retrospective explanations of such cases. It certainly can be useful to explain the features or the success of a representational practice in terms of its relation to a theory or practice that is better understood. And, indeed, one can use RIC to provide at least as good an explanation of the success of the early calculus in terms of its agreement with the modern calculus as we can using a mapping account. But cases of inconsistent mathematics raise a number of philosophical questions that cannot be addressed in this retrospective way—particularly the questions of why it was *reasonable at the time* for scientists to appeal to inconsistent mathematics, as well as how we should understand the methodologies they used to manage the risks associated with using an inconsistent theory.

In chapter 5, I argue that the case considered so far is not an outlier, but an instance of a broader range of phenomena that are not particularly uncommon in the history and current practice of physics. Physicists have often developed and used mathematical techniques that fall well short of the standards of rigor to which work in pure mathematics is generally held. Such techniques have nonetheless proven to be enormously fruitful and often persist long after more rigorous alternatives have been developed. Applying such techniques requires the adoption of what Davey [2003] calls an “inferentially restrictive methodology,” one in which some classically valid inferences are disallowed. While this might be understood in terms of the adoption of a sub-classical logic, in almost all cases the relevant restrictions are more local—though the degree to which this is so varies considerably from case to case. As a result, mapping accounts face difficulties substantially similar to those they face in making sense of applications of inconsistent mathematics. Structures, even partial structures, are often not the right tool to represent the relevant inferential practices at the right level of grain. As a result, mapping accounts are ultimately far clumsier than RIC as tools for reasoning about the many epistemic questions raised by the success of these practices. I begin by introducing Davey’s distinction between inferentially restrictive and permissive methodologies and illustrate its

utility by considering the differences between those at work in the application of the early calculus and those Dirac used to apply another piece of inconsistent mathematics, the delta function. I then turn to two much more detailed case studies.

The first concerns Heaviside's use of his operational calculus to solve problems in electrical engineering at the end of the 19th century. Heaviside's operational calculus reduced difficult problems expressed in terms of differential equations to much simpler algebraic ones in a manner so unrigorous that the Royal Society took the unprecedented step of subjecting his work to peer review, just a few years after naming him a Fellow. Nonetheless, Heaviside appears never to have published an incorrect result derived using his operational calculus. There are a number of striking features of the methodological approach Heaviside adopted to make this work. More than in other cases in which unrigorous mathematics has been applied, Heaviside's inferential restrictions were highly local and ad hoc, which makes it even more difficult to use a mapping account to represent them at the most appropriate level of grain. Beyond this, Heaviside used physical reasoning to inform and constrain both individual mathematical inferences and the development of his mathematical tools. I argue that this too resists treatment in terms of the mapping account. Finally, the relationship between Heaviside's operational calculus and its more rigorous successors is complex, and I argue that RIC does better than mapping accounts even at providing a retrospective explanation of the success of the operational calculus in terms of its successors.

The second concerns the use of path integral techniques in quantum physics. Unlike the other techniques considered so far (but like the better part of quantum field theory!), path integral techniques have not yet been put on a mathematically rigorous foundation except in a very limited range of cases that are not very useful to physics. Again, I argue that because RIC can better represent the philosophically salient features of the diverse inferentially restrictive methodologies used to apply these techniques in various contexts, it is better placed to help us answer the most pressing philosophical questions that the use of these techniques raise.

In chapter 6, I shift my focus to the contributions of mathematics to scientific explana-

tions. In the past two decades, this topic has become central to what is perhaps the debate about applied mathematics with the most extensive literature: the debate over indispensability arguments for platonism. Early versions of the argument supported platonism on broadly Quinean grounds; scientists indispensably make use of quantification over mathematical objects, and such quantification brings with it ontological commitment. As the Quinean version of this argument fell out of favor, it was replaced by a version of the argument that appealed instead to the indispensable role of mathematics in our best scientific *explanations*, most influentially in a paper by Baker [2005]. This topic has also become an important part of the burgeoning literature on non-causal explanations, in many putative examples of which mathematics plays a critical role.

In the chapter, I consider a feature of such explanations that is important in both debates: the extremely high degree of generality achieved by many explanations in which mathematics plays a prominent role. Building on the account of representation in the first part of the thesis, I describe a further, metarepresentational role of mathematics. In this role, mathematics allows us to represent and reason about properties of our mathematical representations that remain stable as we vary the mathematics involved or its physical interpretation. I argue that mathematics in this role makes it possible to represent highly abstract, but still physical features of target systems. These features, I argue, are sufficiently general to serve as the explanantia in the highly general explanations at the center of both debates considered above. As a result, such explanations support neither explanatory indispensability arguments nor the existence of mathematical explanations as a *sui generis* variety of scientific explanation.

In chapter 7, I address the related question of the contributions of mathematics to scientific understanding. I show how RIC can be combined with a prominent inferential account of scientific understanding to shed light on the multifaceted relationship between the degree of mathematical rigor observed in a representational practice and the understanding it produces. Doing so makes it possible to distinguish several potentially conflicting ways in which mathematical tools can contribute to scientific understanding. Actual conflict between

these contributions then helps to explain an apparent tension concerning mathematical rigor and understanding. Both unrigorous techniques and their more rigorous successors have been supported—with what seems to be good reason in each case—on the grounds that those techniques promote understanding in a way that alternative (more rigorous or less rigorous) techniques cannot. Distinguishing these contributions makes it possible to understand each side in these controversies as taking a reasonable stance in the given context, as well as why unrigorous techniques have sometimes subsequently fallen out of favor as that context has changed. I conclude by briefly arguing that this work has implications for two broader issues concerning scientific understanding: first, the contributions of mathematics to scientific understanding in general, and, second, the relationship between understanding and a wider range of ways in which representations “get things wrong.”

Finally, in chapter 8, I draw some initial conclusions, and I explore directions for future work. I consider the possibility of extending RIC’s treatment of applied mathematics to a broader range of cases involving “computational inference procedures,” including general-purpose machine-learning models at the heart of data-driven science, as well as computer simulations more generally. I also consider the possible implications of this work for our understanding of mathematics in general, sketching what I take to be a promising approach based on the work in the thesis.

Part I

Representation

CHAPTER 2

Scientific Representation and Mathematical
Scientific Representations

In this chapter, I discuss several preliminaries to the central arguments of the thesis. My purpose here is not to exhaustively review the relevant literature, which is both diverse and extensive, but rather to provide the context needed to frame the arguments that follow. I first discuss the debate between substantivist and inferentialist accounts of scientific representation in general (§2.1). This will be useful in framing the disagreement between mapping accounts and the inferentialist account of mathematical representation that I defend in the next chapter. I then discuss the extent to which this debate at the level of scientific representations in general bears on the more specific question of how to understand mathematical scientific representation (§2.2). I briefly argue that while the former can greatly inform the latter (and vice versa), there is an important gap between these questions. In particular, there is conceptual space—and, in light of my arguments in the rest of the thesis, good reason—for a proponent of a substantivist (e.g., structuralist) account of scientific representation in general to adopt an inferentialist account of mathematical representation of the kind that I propose in chapter 3. I then discuss how to adjudicate between accounts of mathematical scientific representation, particularly when they agree in the informational content they assign to individual representations (§2.3). Finally, I discuss existing versions of the mapping account (§2.4), as well as arguments against them in the existing literature (§2.5).

2.1 Accounts of scientific representation

Accounts of mathematical scientific representations are closely tied to accounts of scientific representation more generally, and the literature about the latter is significantly broader than the literature about the former. Accounts of mathematical scientific representations concern a particular kind of scientific representation, and so they must pick out some more specific way in which the conditions picked out by a general account of scientific representation are realized. It is therefore worth pausing to consider accounts of scientific representation in general.

There are two broad aspects of scientific representation of which we might want to give an account: informational and functional [Chakravartty, 2010]. The informational aspect of representation concerns the objective relationships—such as the existence of morphisms or other similarities—that must hold if a representational vehicle accurately represents its target. For a given representation, we call these objective relationships its *informational content*. The functional aspect of representation concerns the cognitive activities in which representations are used in scientific practice, especially in the use of representational vehicles to make surrogative inferences about their targets. More pithily, an account of the informational aspects of representation concern what a representation *is*, while an account of its functional aspects concerns what scientists *do* with representations. Or, better: informational accounts concern the semantics of scientific representations, while functional accounts concern their pragmatics.

Chakravartty argues persuasively that accounts of these two aspects of scientific representation are complementary, not in direct conflict. To understand scientific representation, we must understand both aspects. A functional account would seem to rely crucially on the informational content of a representation to explain how it can play its functional role in scientific practice; shorn of this, it is hard to see what philosophical work such an account could accomplish. And the same is true of informational accounts; the very point of them would seem to be to describe the features in virtue of which representations are able to play the roles they play in scientific practice. Insensitivity to these features of practice would likewise undermine the very purpose of such an account.

There are two broad kinds of view on the nature of scientific representation: deflationary and substantive views.

Substantive views identify some substantive property or relation common to all scientific representations in virtue of which they are scientific representations.¹ Common versions ap-

¹I take this to allow for a pluralist account of scientific representation, according to which the relevant substantive property is, say, a disjunction of several substantive properties, though I do not know of anyone who explicitly defends such a view.

peal to similarities [e.g., [Giere, 2004, 2010](#)] or structural mappings (morphisms) [e.g., [da Costa & French, 2003](#)] between representational vehicle and target. In their contemporary forms, these views don't take the existence of such a similarity or morphism to be necessary or sufficient for scientific representation. Instead, they typically take the existence of such similarities or morphisms to be necessary for *accurate* representation, while representation simpliciter is achieved when agents *posit* such similarities or morphisms in order to use a representational vehicle to represent its target. In recent years, more sophisticated views have been developed (e.g., Frigg and Nguyen's [2016, 2017] DEKI view), which describe more intricate relations between representing agents, representational vehicles, and representational targets.

In presenting a substantive property (like the existence of a morphism or similarity) as a necessary condition for accurate representation, substantive views thereby give an account of the informational content of scientific representations. But such views are typically also concerned with the functional aspects of representation, explaining these functional aspects in terms of their construal of the informational content of the representation. Indeed, sometimes particular accounts of the informational aspects of representation are chosen precisely for the ways in which they can be used to shed light on functional aspects of representation. For instance, the partial structures approach [[da Costa & French, 2003](#)] is a version of the structuralist account of scientific representation that appeals to more liberal notions of structure and morphism. This is in no small part because, according to proponents of the approach, partial structures can be used to represent scientists' reasoning with inconsistency and uncertainty, theory change over time, and various common heuristics, among other (functional) things. But even the most standard structuralist approaches take scientists' surrogative reasoning about a representational target in terms of its source, a functional aspect of representation *par excellence*, to be licensed by the positing of an appropriate morphism. And discussion in structuralist terms of the heuristic value of surplus structure, another functional aspect of representation, goes back as far as [Redhead \[1975\]](#).

In contrast to substantive views, deflationary views deny that there is a substantive prop-

erty or relation common to all scientific representations in virtue of which they are scientific representations. Instead, they take the concept to be exhausted by platitudes or surface level features of scientific practice. Though earlier forms of the view exist (e.g. perhaps Hughes' [1997] DDI account), Suárez's [2004, 2015] deflationary inferentialist account is the most prominent. It is based on two platitudes about scientific representations, roughly: (i) scientific representations are *about* their targets and (ii) scientific representations can be used to make inferences about their targets. And so, according to Suárez, the following necessary conditions for scientific representation are the best we can do: "A represents B only if (i) the representational force of A points towards B, and (ii) A allows competent and informed agents to draw specific inferences regarding B" [Suárez, 2004, p. 773]. Such views are motivated by concerns that substantive views of scientific representation are overly restrictive and therefore fail to capture at least some instances of scientific representation. Faithfulness to the particularities of scientific practice requires forgoing a substantive account.

At first glance, this seems to rule out a deflationary treatment of the informational aspects of representation, leaving deflationists able only to describe (and not explain) the functional aspects. As Bueno & French [2018, p. 65] put it:

to base one's account of representation on surrogate reasoning or inferences without accepting the underlying formal aspects is, as Chakravartty notes, to engage in a confusion. Indeed, we would go further and insist that such a move makes no sense: it is only within some account of the formal relationships involved that we can understand how the relevant reasoning can be appropriately surrogate.

For similar reasons, Contessa [2007, 2011] formulates a substantive version of Suárez's inferentialism. Contessa is motivated by the thought that a deflationary view fails to answer the central questions we might expect an account of scientific representation to answer, like how it is that scientific representations support successful surrogate inference, essentially

because it fails to give an account of the informational aspects of representation.² Because the deflationist insists that there is nothing to the concept of scientific representation beyond bare platitudes about surrogative inference, they seem to explicitly rule out having anything informative to say about the informational aspects of representation.

But I think this is a bit too hasty. The deflationist need not deny that representation has informational aspects, but only that there is something substantial to be said about those informational aspects in general—at the level of *all* scientific representations. A deflationary inferentialist *can* say *something* about the informational content of scientific representations in general—and quite a lot about the informational content of more specific families of such representations.

At the level of all scientific representations, we can provide a basic deflationary semantics for scientific representations with something like the schema:

- (R) A representation R , as used by competent agent S to represent T , is accurate if and only if the surrogative inferences S is willing to use R to draw about T are appropriate (or preserve truth or accuracy or ...).

While this indeed says very little about the informational content of representations, it says more than nothing. And it leaves room to say more at a lesser degree of generality. Compare the above schema to the T-schema, frequently cited by deflationists about truth as an exhaustive characterization of the concept of truth:

- (T) $\ulcorner p \urcorner$ is true if and only if p .

While this certainly says less about truth than a substantive theory, like the correspondence or coherence theories, it says enough, according to deflationists about truth, to explain how truth plays the functional role it plays in discourse in general. It also leaves room to say more about the objective relations in virtue of which $\ulcorner p \urcorner$ is true for more specific values of p . But

²While I am sympathetic to Contessa's point here, I do have my doubts about his proposed solution, which I suspect collapses into an overly restrictive form of structuralism. Unfortunately, sustained discussion of Contessa's view would go well beyond the scope of this thesis.

these objective features are most usefully discussed at a more particular level because they crucially depend on the content of p .³

The deflationist about scientific representation should say something analogous. The schema (R) is all we can say about the informational content of scientific representations is general. But at lower degrees of generality, the relation it picks out is realized by more concrete relations about which we can usefully say something substantial. This puts the view in direct competition with substantive views like the structuralist, similarity, and DEKI accounts. But it doesn't thereby exclude informational aspects of representation from consideration. Instead, the disagreement concerns the level of generality at which these informational aspects can be usefully be characterized in terms of substantive properties. While Suárez, to my knowledge, doesn't directly address this, something like this does seem to be his own view. For instance, he writes noncommittally: "In every specific context of inquiry, given a putative target and source, some stronger condition will typically[!] be met" [Suárez, 2004, p. 776].⁴

An intriguing consequence of this is that there is conceptual space for a spectrum of deflationary inferentialist views differing in the level of generality at which something substantive can be said about the informational content of a scientific representation. On one end of this spectrum is an extreme particularism according to which the relevant representational relationships must be characterized in a very fine-grained way in reference to highly specific ranges of practices. The problem with this sort of view is that work done in accordance with it runs the risk of turning into an endless line of case studies, with few resources to draw broader philosophical lessons from those case studies. On the other end are views according to which both high- and low-level characterizations of the informational content of repres-

³For a useful summary of the enormous literature on deflationism about truth, see [Armour-Garb *et al.*, 2022].

⁴I suspect he writes "typically" because his view is deflationary in another sense: it purports to describe necessary but not jointly sufficient conditions for scientific representation. While I am sympathetic to many aspects of his view, I cannot follow him in this. Even if some stronger condition is met, as Suárez writes, this needn't mean that the more general conditions are not jointly sufficient but only that, when they are met, some condition from a class of more specific conditions is also met. Indeed, I can't imagine a case in which Suárez's two conditions are met that fails to be a scientific representation.

2.2 From scientific representation to mathematical scientific representation

entations may be called for depending on which functional aspects of practice we wish to explain. Such views have the potential to do much the same philosophical work as substantive accounts by appealing to the morphisms or other similarities that those accounts posit. But they need not appeal to the resources of those accounts in more specific ranges of cases in which those resources are less useful.

While I am sympathetic to the deflationary inferentialist approach, especially in its less particularist forms, for the purposes of the thesis, it is enough to show that both kinds of account are viable options. Ultimately, I will argue that my account of mathematical scientific representations is the best option for proponents of both views, even though it stands in a similar relationship to the mapping account as deflationary inferentialism stands to substantive accounts of scientific representation in general. In the next section, I turn to the more specific question of mathematical scientific representation and what we can bring to it from the substantivist-deflationist debate about representation in general.

2.2 From scientific representation to mathematical scientific representation

Accounts of mathematical scientific representation are generally intended to perform two tasks. The first is to provide an answer to what [Nguyen & Frigg \[2021\]](#) call “the general application problem.” This is the problem of explaining how mathematics can represent target systems in general, with emphasis on the question of how any piece of mathematics could “hook on” to the worldly phenomenon it represents. Given that the most prominent accounts of the nature of mathematics, both platonist and nominalist, take the subject matter of mathematics to be distinct from those of the domains in which it can be applied, the fact that mathematics can be made to represent such domains itself demands an explanation.

The second task that accounts of mathematical scientific representation are generally intended to perform is to serve as a meta-level device to be used by philosophers of science to

2.2 From scientific representation to mathematical scientific representation

represent the philosophically salient features of particular episodes from historical and contemporary scientific practice. For example, [Bueno & French \[2018\]](#) explicitly emphasize the utility of their account—a combination of Bueno and Colyvan’s [\[2011\]](#) inferential conception of applications of mathematics⁵ and the partial structures approach to scientific representation [[da Costa & French, 2003](#)]⁵—as such a device, spending much of the book applying that account to episodes from scientific practice to illustrate philosophical problems that arise in connection with applied mathematics. Similarly, [Pincock \[2012\]](#) starts by setting out a version of the mapping account of applied mathematics, but he spends most of the book using that account as a meta-level device to bring out philosophically salient features of cases in which mathematics is applied in service of his conclusion that the central contribution of mathematics to science is an epistemic one.

These two tasks map nicely onto the informational-functional distinction discussed previously. The first concerns how a mathematical representation could come to have physical content, while the second largely concerns what scientists use such representations to do. In this sense, the aims of an account of mathematical scientific representation parallel those of an account of scientific representation in general.

The substantivist-deflationist distinction can also be naturally extended to the debate over the nature of mathematical scientific representations. Indeed, by far the most prominent approach to mathematical scientific representation for the purpose of fulfilling both of the tasks discussed above is some version of the mapping account, a very natural way of extending structuralist accounts of scientific representation in general to the more particular case of mathematical scientific representation. Such accounts explain the informational content of mathematical scientific representations in terms of morphisms between relevant mathematical structures and the structures of their target systems. This explanation of the informational content in turn helps to fulfill the second task, bringing out philosophically salient features of particular episodes from scientific practice. Likewise, the account of mathematical scientific

⁵As I will discuss in the next section, despite the name, this account is a version of the mapping account rather than an instance of the sort of inferentialist account of mathematics I defend.

2.2 From scientific representation to mathematical scientific representation

representation I put forward in the next chapter can be understood as a very natural way of extending deflationary inferentialism about scientific representation to the more particular case of mathematical scientific representation, offering a basic inferential semantics for such representations and using that inferential semantics to bring out philosophically significant features of scientific practice.

Despite these natural alliances, the substantivist/deflationary inferentialist distinctions in these two domains cut across one another. Perhaps least controversially, for reasons discussed in the previous section, a deflationary inferentialist about scientific representation in general could accept a version of the mapping account (or conceivably some other substantive account) about mathematical scientific representation in particular. This is because deflationary inferentialism allows for a spectrum of views differing in the level of generality at which something substantive can be said about the informational content of a scientific representation. While a highly particularist form of deflationism about scientific representation rules out saying anything substantive about a range of representations as wide as that of mathematical scientific representations, there is more than enough conceptual space for a more moderate deflationism that accepts that a substantive account can be given of the informational content of mathematical scientific representations, while denying that any such account can be given for the full range of scientific representations. Such a view could then allow that, in the particular case of mathematical scientific representations, a mapping account correctly explains the worldly relation that underwrites scientists' surrogative reasoning.

More controversially, I believe a substantivist about scientific representation in general could accept a version of the deflationary inferentialist account of mathematical scientific representations that I propose in the next chapter. In such a case, one would accept that scientific representations in general are underpinned by a substantial worldly relation, but take it that, when present, mathematics serves to pick out such relations in a number of diverse ways, which cannot themselves be captured by a single substantive relation. For example, one might be a structuralist about representation in general, but hold that mathematics

serves to pick out the source structures of those representations in a diverse range of ways; sometimes there is a straightforward morphism between a mathematical structure and the relevant source structure, while in other cases, the mathematics cannot be associated with a well-defined mathematical structure and is associated with a representational source only through the inferential affordances of the formalism used to express it. I discuss this possibility in detail in §3.3.2.

Consequently, the question of whether to accept a substantive or deflationary inferentialist account of mathematical scientific representation cannot be reduced to the more general question of whether to be substantivist or deflationist about scientific representation in general. While each debate can arguably greatly inform the other, as they share a common structure, each is ultimately independent of the other and must be considered individually.

2.3 How to argue for an account of mathematical scientific representation

So how should this debate be adjudicated in the case of mathematical scientific representations? First, it is important to note that, as I will argue in the next chapter, RIC recovers mapping accounts as special cases. Accordingly, adjudicating between the two views comes down to the question of whether RIC's generality has philosophical benefits that outweigh its costs; if RIC had no such benefits, mapping accounts would be preferable on the basis of their greater specificity. What might such benefits be?

The two tasks for an account of mathematical scientific representation are a good place to start. The first of these tasks is explaining how mathematics can in principle be made to represent non-mathematical target systems. The most essential part of this task, as I see it, is to present an adequate approach to the semantics of mathematical scientific representation, so that whatever representational relation one posits can capture the informational content of the full range of mathematical scientific representations. While I will discuss a few reasons

2.3 How to argue for an account of mathematical scientific representation

to think mapping accounts might have trouble with this task in a few corner cases, for most of the thesis I assume for the sake of argument that the mapping account (and so also my own account, which recovers mapping accounts as special cases) can adequately carry out this task or naturally be extended to do so.

It is more difficult to judge the merits of competing accounts with respect to the second task, serving as a meta-level representational device to bring out philosophically salient features of scientific practices in which mathematics is applied. In particular, if the strengths of two competing accounts lie in different areas, it might be that for each account there are domains and purposes for which it is the best tool for bringing out the philosophically salient features of certain scientific practices. In such a case, it is not clear that we would have reason to favor either account over the other as a *general* account of mathematical scientific representations. In addition to this, there is the problem of specifying what makes one account a better meta-level device than another even in a single particular case.

Concerning the first of these questions, a simplifying factor is that RIC recovers mapping accounts as special cases. (I argue for this in §3.1.4.) As a result, any resources available to the mapping account are also available to RIC, provided that they don't require the mapping account to hold of *all* applications of mathematics in science. This means in particular that in the cases in which mapping accounts are best suited to serve as meta-level representational devices for philosophers of science, the proponent of my account can coopt the representational tools that allow the mapping account to do so well in those cases. If I am right about this, then the question of comparing the two views comes down to whether RIC has advantages in other cases sufficient to warrant the move to a less substantive and therefore strictly less informative account. Throughout the thesis, I will consider a number of cases from the history of science in which I take RIC to have such advantages.

But what makes one account a better meta-level device than another in relation to a single case study? This necessarily varies from case to case, as it depends essentially on the philosophical aims of those using the account. However, what we can say in general is that an

account performs better in this role to the extent that it more perspicuously represents the features of scientific practice relevant to the philosophical purposes. Later in the thesis, I will focus on one important way in which one's account of mathematical scientific representations can affect the perspicuity with which these features are represented—viz., the level of grain at which scientific practice is represented. Some questions turn on features of scientific practice that are highly specific to particular cases, while others require a greater degree of generality. I will argue that a benefit of RIC is that it allows for representations of scientific practice with a far broader range of levels of grain, from maximally fine-grained descriptions of individual scientists' inferential behavior, to much more coarse-grained accounts of broader inferential patterns, including those most amenable to treatment in terms of the mapping account. I argue that these benefits of RIC outweigh the cost of being able to say less about mathematical scientific representations in full generality.

2.4 Varieties of the mapping account

Finally, before moving on to my main argument for RIC, it is worth pausing to discuss the existing varieties of the mapping account. These represent the bulk of existing work on the nature of mathematical scientific representation, and I will refer back to them frequently in the rest of the thesis. Again, my aim here is not to be exhaustive, and I will discuss details of these views again as they become relevant to my arguments.

While often simply called “the mapping account”, mapping accounts are a diverse group of positions. Their common core is the idea that mathematical representations represent in virtue of positing a relation between the structure(s) picked out by the mathematics and the structure of their target systems. Such representations represent their target systems as bearing a structural similarity to the structure (or members of the class of structures) picked out by the mathematics, with this similarity cashed out in terms of a structure-preserving mapping between (a substructure of) the structure of the target system and (a substructure of) the

relevant mathematical structure. The posited mapping serves to determine the informational content of the representation, which in turn helps to explain other features of scientific practice related to the representation, like the inferences it allows scientists to make, its role in scientific explanations, and so on.

The mapping account is at least implicit in most philosophical work on applications of mathematics from the past few decades, but often is not quite fully developed [see, e.g., Baker, 2003, Balaguer, 1998, Leng, 2002, Shapiro, 1997]. The best developed versions of the view are the “inferential conception” put forward by Bueno & Colyvan [2011] and refined by Bueno & French [2012, 2018] and Räs & Sauer [2015] and the more classic version of the mapping account defended by Pincock [2012]. But, as the mapping account is the analogue for mathematical representations of the structuralist account of scientific representation in general⁶, it can come in even more forms corresponding to the many versions of the former view. In particular such views can differ in what they take the relevant mathematical source structures to be, the nature of the relevant mappings, as well as in the inclusion of various pragmatic features.

With respect to the relevant structures, these are classical set-theoretic structures by default [e.g., Pincock, 2012]. But proponents of the partial structures program have proposed using partial structures instead (see, e.g., da Costa and French 1990, 2003; Bueno 1997; French & Ladyman 1999; French 2003, 2014; Bueno & Colyvan 2011; Bueno and French 2011, 2012, 2018). These differ from ordinary structures in that they can leave relations and functions undefined for certain arguments in order to represent the uncertainty of scientists about what these structures are like. The notion of a partial structure is a modest generalization of the standard notion of a set-theoretic structure. A classical set-theoretic structure consists of a domain D and a family $\{R^i\}_{i \in I}$ of relations defined on that domain, where these relations are

⁶Again, it is important to remember that, while it is the natural analogue of the structuralist account of representation, the mapping account can be understood in terms of most (if not all) general accounts of scientific representation. In the case of substantive views, the mapping explains how substantive properties like similarity are realized in terms of the relevant sort of mapping. In the case of deflationary accounts, it provides an account of *how* mathematical representations support surrogative reasoning.

understood in the traditional, extensional way, so that an n -ary relation R is the set of n -tuples of items in D that stand in that relation. In a partial structure, relations need not be defined for every n -tuple of items in D . Instead a relation R in a partial structure is represented as a triple $\langle R_1, R_2, R_3 \rangle$, where R_1 is the set of n -tuples of which R holds, R_2 the set of n -tuples of which R does not hold, and R_3 the set of n -tuples for which R is undefined. (And so $R_1 \cup R_2 \cup R_3 = D^n$.) Total consistent structures are special cases of partial structures, in which R_3 is empty for every relation in the structure. A total consistent structure \mathcal{S}' extends a partial structure \mathcal{S} if and only if they have the same domain and for every relation R in \mathcal{S} , there is a relation R' of the same arity in \mathcal{S}' such that $R_1 \subseteq R'$ and $R_2 \cap R' = \emptyset$. A sentence ϕ is then said to be ‘partially true’ in a partial structure \mathcal{S} if and only if there is a total consistent structure \mathcal{S}' extending \mathcal{S} such that ϕ is true in \mathcal{S}' . We might use a similar formal device to make sense of “inconsistent structures” in the context of the mapping account, as suggested by Colyvan [2008a,b, 2009].⁷

Mapping accounts and structuralist accounts of representation alike can also differ with respect to the sort of mapping they require. Some proposals require the posited mapping to be as strong as an isomorphism or isomorphic embedding [van Fraassen, 1980, 2008], while others weaken this requirement by allowing the structural similarity to be cashed out in terms of homomorphism [Bartels, 2006, Lloyd, 1984, Mundy, 1986], a merely partial morphism (proponents of the partial structures program again)⁸, or something even more permissive, like the “more intricate sorts of structural relations” described by Pincock [2012, p. 27].

Finally, such views can differ in how pragmatic features figure into the account. While both Pincock’s [2012] version of the mapping account and Bueno and Colyvan’s [2011] inferential conception make explicit provision for certain pragmatic features like scientific agents’

⁷See chapter 4 and [McCullough-Benner, 2019] for more details.

⁸A *partial homomorphism* between partial structures $\mathcal{A} = \langle D, \{R^i\}_{i \in I} \rangle$ and $\mathcal{B} = \langle D', \{R'^i\}_{i \in I} \rangle$ is a partial function $f : D \rightarrow D'$ such that for every n -ary partial relation R in \mathcal{A} and every $a_1, \dots, a_n \in D$, $R_1 a_1 \dots a_n \Rightarrow R'_1 f(a_1) \dots f(a_n)$ and $R_2 a_1 \dots a_n \Rightarrow R'_2 f(a_1) \dots f(a_n)$. We can formulate partial versions of other kinds of morphism by adding a clause for the R_2 block to their usual definitions in the same way. In each case, the defined partial morphism reduces to the original sort of morphism in the special case that it obtains between two ordinary, set-theoretic structures.

intentions, the inferential conception is so named because it combines a mapping account of the informational content of mathematical scientific representation together with a more elaborate account of how to understand the practices through which such structures support scientists' surrogative inferences. In particular, inspired by Hughes' [1997] DDI model of scientific representation, they posit that applications of mathematics can be broken into three (conceptually, but not necessarily temporally) distinct phases: (1) *immersion*: the establishment of a mapping from a posited structure in the world to some mathematical structure, allowing problems to be framed mathematically; (2) the *derivation* of results about that mathematical structure; (3) *interpretation*: the establishment of a mapping from the mathematical structure back to the target structure (not necessarily the inverse of the first mapping) so that the results of step 2 can be physically interpreted.

2.5 Arguments against mapping accounts and the need for an alternative

A number of objections have already been put forward against mapping accounts. In this section, I briefly consider two types of arguments against the mapping account. First, I consider several influential general arguments against mapping accounts, concluding that each is ultimately unpersuasive. Second, I consider a more promising argument strategy, which concerns shortcomings of mapping accounts in particular kinds of cases, rather than arguments from first principles, and I briefly argue that such arguments are more persuasive when framed in terms of an alternative to the mapping account.

2.5.1 General arguments

Frigg [2006] argues that what structure a target system instantiates depends essentially on how it is described, though Nguyen & Frigg [2021] suggest supplementing mapping accounts with an account of "structure-generating descriptions" as a remedy.

2.5 Arguments against mapping accounts and the need for an alternative

Suárez [2003] presents five arguments against accounts of scientific representation based on structural relations (and so presumably also mapping accounts). Two of these—the “logical” and “non-sufficiency” arguments—seem only to apply to accounts that do not incorporate pragmatic elements. Since most contemporary proponents of mapping accounts (and structuralist accounts more generally) do not oppose the inclusion of pragmatic elements, these can safely be ignored for present purposes.

The first of the remaining arguments, the “argument from variety,” alleges that structural relations are not the means of representation in many cases of scientific representation. Most importantly for present purposes, Suárez argues that even in paradigm cases of mathematical scientific representation, in which the relevant structural relation is present, that structural relation is not the *means* of the representation. For example, in a typical application of a differential equation, scientists don’t explicitly reason about the structural relationship between structures satisfying the equation and the phenomenon under investigation; instead, they reason about the equation itself, looking for solutions given particular boundary conditions and comparing its parameters to features of the phenomenon. While I very much agree with Suárez’s emphasis on how little the proposed structural relation has to do with scientists’ actual activity in applying mathematics, I think this argument is ultimately unpersuasive. In particular, as a proponent of a mapping account would rush to point out, their claim is not that mathematical scientific representations all involve scientists’ *explicitly* reasoning about morphisms. Rather, the mapping account is a rational reconstruction of scientific practice, and the thought is that what licenses the type of surrogative reasoning that Suárez describes is the tacit presupposition that certain structures satisfying the differential equation are appropriately morphic to the system under investigation. If scientists thought no such relationship existed, then we would be at a loss to explain how their reasoning about the equation had anything to do with the phenomenon under investigation.

In the second remaining argument, the “argument from misrepresentation,” Suárez claims that a view based on isomorphism cannot account for misrepresentation. Suárez distinguishes

2.5 Arguments against mapping accounts and the need for an alternative

two kinds of misrepresentation, mistargeting and inaccuracy. Concerning mistargeting, the claim is that a representation may fail to represent a system that it is appropriately morphic to. This again is only a problem for views that do not incorporate a pragmatic element. I suspect Suárez would agree that the reason why something might fail to be a representation of a system it is appropriately morphic to is simply that it is never *used* as such.

Concerning inaccuracy, the thought seems to be that an inaccurate representational source necessarily fails to be isomorphic to its target. It is in this sense that an isomorphism-based view “cannot account for inaccurate representation at all” [Suárez, 2003, p. 235]. Suárez extends this criticism to views based on partial structures, but not those based on homomorphism [Lloyd, 1988] or Swoyer’s [1991] Δ/Ψ morphisms. While I have my doubts about Suárez’s extension of this criticism to partial isomorphism⁹, note that this leaves mapping accounts based on other kinds of morphism—including not just homomorphism and Δ/Ψ morphism, but also partial homomorphism and partial Δ/Ψ morphism—entirely untouched.

The final remaining argument is Suárez’s “non-necessity argument,” according to which the existence of an isomorphism is not necessary for scientific representation. But Suárez’s reason for this is again that isomorphism views cannot accommodate inaccurate representation, and so this argument does not extend to views based on weaker kinds of morphism.

Ultimately, I suspect that the prospects aren’t good for an argument from first principles against the mapping account. Such arguments can tell us something about the form a plausible version of the mapping account should take, but existing arguments are unconvincing when construed as arguments against mapping accounts in general.

2.5.2 Particular arguments and the need for an alternative to mapping accounts

A more promising strategy for arguing against mapping accounts is to point out their shortcomings in concrete cases. For instance, Batterman [2010] argues that mapping accounts fail to capture instances of “asymptotic reasoning,” in which a move to a more faithful model

⁹See e.g. Bueno and French’s response [2018, p. 69].

2.5 Arguments against mapping accounts and the need for an alternative

would lead to a loss in explanatory power. And [Rizza \[2013\]](#) argues that certain applications of mathematics are “qualitatively different” from those that are well understood in terms of the mapping account. In particular, appealing to an example from social choice theory, he argues that in some cases scientists are better understood as borrowing concepts and forms of reasoning from mathematics, rather than structures. These concepts and forms of reasoning can then be directly applied to reason about the formal properties of an empirical target system, rather than indirectly in terms of a mathematical structure.¹⁰ Likewise, in chapters 4, 5, and 7, I present several further cases in which mapping accounts have significant shortcomings.

But it is not clear what we should conclude from such cases in the absence of a viable alternative to the mapping account. Perhaps it is perfectly fine if, for example, we can treat the representations Batterman is interested in in terms of the mapping account, but we must look beyond the resources of the mapping account itself to make sense of asymptotic reasoning. Certainly this will be the right response in at least some cases. For instance, [Suárez & Cartwright \[2008\]](#) criticize the partial structures account on the grounds that important features of episodes of theory change—including scientists’ motivations, available techniques, and background knowledge—cannot be represented in terms of partial structures. Bueno and French respond:

But, of course, to try to represent model-theoretically the relevant scientists’ *motivations* would be an entirely misguided endeavor, akin to the attempt [...] to accommodate sociological factors in theory change by stipulating let S be a set of scientists! [[Bueno & French, 2018](#), p. 232, emphasis in original]

We simply should not expect the formal apparatus of the mapping account—or any other account of mathematical scientific representation for that matter—to represent *every* philosophically salient feature of scientific practice involving mathematical scientific representations.

But which features should and shouldn’t we expect to the formal apparatus of a mapping

¹⁰One way to understand the account I propose in the next chapter is as suggesting that we can think of all applications of mathematics in something like this way, even in cases well treated in terms of mapping accounts.

2.5 Arguments against mapping accounts and the need for an alternative

account to capture? Ultimately, I suspect that this is a question we can't answer confidently by considering mapping accounts in isolation. Instead, in determining what features of practice a good account should capture, we should be doing so in comparison to other accounts. In the unlikely case that some account could capture all of the benefits of the partial structures approach, while also providing a useful tool for representing and reasoning about scientists' motivations, we would have good reason to adopt it instead of the partial structures approach. In a more likely case, an account might make trade-offs in order to represent scientists' motivations, and in that case we would have to holistically weigh the benefits and drawbacks of that account against those of the partial structures account.

In the next chapter, I present what I believe to be the first viable and fully worked out alternative to the mapping account in the literature on mathematical scientific representation, the robustly inferential conception (RIC). Then I present an initial case for RIC over the mapping account. Because RIC recovers mapping accounts as special cases, this only requires me to show that RIC has benefits over the mapping account that outweigh the cost of moving to a more general account. The rest of the thesis is devoted to exploring some of these benefits in connection to cases including applications of inconsistent (chapter 4) and otherwise unrigorous (chapter 5) mathematics and the role of mathematics in scientific explanations (chapter 6) and scientific understanding (chapter 7).

CHAPTER 3

The Robustly Inferential Conception of
Mathematical Scientific Representations

In this chapter I present and defend a novel account of mathematical scientific representation, the robustly inferential conception (RIC). After presenting RIC in detail (§3.1), I present my central arguments for the view (§3.2) and examine its relationship to accounts of scientific representation more generally, arguing that there is good reason to adopt RIC regardless of one's views on scientific representation generally (§3.3).

3.1 The robustly inferential conception

Recall from chapter 2 that accounts of mathematical scientific representation are generally intended to perform two tasks. The first is to explain how mathematics can represent target systems in general, with a focus on how any piece of mathematics could “hook on” to a system it is used to represent. The second is to serve as a meta-level representational device to be used by philosophers of science to bring out philosophically salient aspects of scientific practices in which mathematics is applied.

There we examined one broad set of ways to approach these two tasks: mapping accounts. Mapping accounts address the first task by explaining mathematical representations in terms of structural relations (usually morphisms); mathematics “hooks on” to the non-mathematical world in virtue of the shared structure entailed by the existence of relevant relations between mathematical structures and those of concrete target systems. These structural relations also provide a framework for thinking about applications of mathematics in scientific practice (the second task). Scientists's mathematically mediated inferences are licensed by the existence of such structural relations, and many heuristic moves can be understood in terms of interpreting surplus structure via a more extensive mapping or relating the relevant structures to further structures via further mappings.¹

In this chapter I present an alternative, the Robustly Inferential Conception of Mathematical Scientific Representations (RIC), which approaches these tasks by appealing not to shared

¹See, for example, [Bueno & Colyvan \[2011, pp. 364f\]](#) and [Bueno & French \[2018, pp. 141ff\]](#) on the reasoning that led Dirac to posit the positron.

structure but to shared patterns of inference. The central idea behind RIC is that all we can say *in full generality* in response to the first task is that mathematics is relevant to physical target systems because some of the patterns of inference appropriate for reasoning about the mathematics are also appropriate for reasoning about those target systems. Mathematics places constraints on what the target system of a representation must be like by helping to specify inferences about the target system that preserve truth according to the representation—that is, inferences that would have to preserve truth if the representation were perfectly accurate in all respects. The forms of these inferences are drawn from the patterns of inference that are permissible in reasoning within the relevant mathematical theory. They become physically relevant because, in the context of a mathematical scientific representation, the language in which the mathematical theory is expressed is decoupled from its mathematical interpretation and given a (partial) physical interpretation. Mathematical inferences expressible in this language become purely physical inferences when the premises and conclusion are decoupled from their mathematical interpretation and provided with a physical interpretation instead. The commitments of such a representation then are the physically interpreted statements expressible in the language of the mathematical theory that can be derived from an initial description of the target system through (perhaps repeated) application of the allowed inference patterns.

In other words, such representations have three basic ingredients:

- (RIC1) a *physical interpretation* of the language of the mathematical theory sufficient to provide at least some sentences in this language with physical truth conditions,
- (RIC2) an *initial description of the target system* in the language of the mathematical theory, given this interpretation, and
- (RIC3) a *collection of privileged inference patterns* from those licensed by the original mathematical theory.

The informational content of such a representation is the closure of the claims in RIC2 under

the inference patterns in RIC3, under the interpretation given by RIC1.

These ingredients are not always (or even usually) cleanly separated in practice. For instance, the inference patterns and physical interpretation chosen play a significant role in determining what the initial description of the target system must look like. Even very simple mathematical descriptions of a target system depend very much on which mathematical moves are allowed. As a simple example, consider Newton's second law of motion, $F = ma$, including a standard interpretation according to which F represents the force on some object, m that object's mass, and a that object's acceleration (all in some set of units). Understood in the usual way, the second law of motion commits us not just to claims about force, mass, and acceleration, but, e.g., to claims about the positions and velocities of objects at certain times. But it can do this only if certain inference patterns from calculus belong to our collection of inference patterns. Otherwise, $F = ma$ does not express the law that it is usually meant to express, since the latter does have consequences for the positions and velocities of objects and how these change over time. For instance, a cannot really represent acceleration (or at any rate be a *good* representation of it) if, say, v represents velocity, and it is possible for the magnitude of a to be non-zero, while v remains constant.²

Despite the fact that these three ingredients are frequently intertwined in practice, they are conceptually distinct, and understanding applications of mathematics in terms of them is extremely fruitful. In the rest of this section, I more precisely specify what these ingredients are.

3.1.1 RIC1: Physical interpretation

By physical interpretation, I mean a way of correlating some of the mathematical vocabulary with parts or features of the representation's target system in such a way that we can use it to give a physical content to claims expressed in otherwise purely mathematical language.

²That said, it is not strictly true that any representation with Newton's second law of motion as part of the initial specification of the target system must commit us to claims about positions and velocities. But the only exceptions will be cases in which we fail to represent positions and velocities at all, in which case the second law of motion would fail to do much work in the first place.

The requirements placed on physical interpretations are fairly minimal, so that the relevant physical property or object may be, for instance, very poorly understood, non-fundamental, or even a mere placeholder. For instance, to physically interpret Newton's second law, it suffices to say that values of F represent the total force on a given object and direction of that force, that values of m represent its mass, and that values of a represent its acceleration and the direction of that acceleration, all in an appropriate set of units. Such an interpretation *might* be specified in terms of a mapping between a physical and a mathematical structure, but it need not be. All that is required is that statements in the language of the mathematical theory can be decoupled from their mathematical interpretation and given physical truth conditions.

Also important to note is that such an interpretation need not provide physical content to every statement expressible in the language of the mathematical theory. (It is in this sense that the robustly inferential conception only requires a *partial* physical interpretation of the mathematical vocabulary.) This roughly corresponds to the notion of surplus structure at work in the mapping account. In the mapping account framework, a mathematical structure might be useful for representing some target system even though it contains structure not taken to be present in the target system. This surplus structure is then simply not mapped to any part of the physical structure (and vice versa). The surplus structure is then just an idle part of that representation, though thinking about this extra structure might be a useful heuristic for formulating a new representation that takes aspects of that structure to be present in the target system. Likewise, in the framework of the robustly inferential conception, there are statements in the language of the mathematical theory that are taken to have no appropriate physical analogue—for example, statements about what is the case when a variable interpreted as representing mass takes a negative value.

3.1.2 RIC2: Initial description of the target system

The initial description of the target system is a (possibly empty) set of claims about the target system either in the relevant mathematical language under the physical interpretation

RIC1 or readily translatable into that language. These claims represent the information we have about a target system prior to the mathematically mediated inferences we make about the target system using the relevant mathematical scientific representation. More precisely, it represents the information about the target system that can enter into premises in mathematically mediated inferences about the target system licensed by the representation; this is why it must be stateable in the relevant mathematical language under the physical interpretation RIC1. RIC2 then serves as a base from which further commitments of the representation can be derived by means of the inference patterns in RIC3.

RIC2 might be empty in non-trivial representations provided that RIC3 contains premise-free inference patterns. Perhaps it could even be dispensed with altogether by replacing each claim in RIC2 with a corresponding premise-free inference pattern in RIC3. Nonetheless, I think RIC2 is useful for highlighting the distinction between information about the target system that is *presupposed* by the representation (which appears in RIC2) and information that can be *inferred* by means of the representation.

Something similar arguably must be introduced to supplement mapping accounts for what are ultimately similar reasons. [Nguyen & Frigg \[2021\]](#) argue that mapping accounts on their own fail to provide an adequate account of the structure of the target system. Mapping accounts presuppose that there is such a structure, which can be mapped to mathematical structures, but there are a number of problems in determining just how we should understand the structure of the target system. They argue that there is no such thing as “the” structure of the target system; instead, for any particular (token) physical system, there are many structures that that system instantiates, only some of which are relevant to any given mathematical representation of that target system. So mapping accounts must be supplemented with an account of how a particular structure (or range of structures) is picked out from this broader set of structures instantiated by the target system.

Nguyen and Frigg propose that such structures are picked out by “structure-generating descriptions” of the target system. Such descriptions identify which physical entities and

properties should correspond to the objects and relations of such a structure, which in turn makes it possible to specify a (range of) morphism(s) between this structure and a mathematical structure, allowing the mapping account to get off the ground.

Like RIC2, Nguyen and Frigg's structure-generating descriptions ultimately represent the presuppositions about the target system that must be made for a given mathematical scientific representation to license further, mathematically mediated inferences about that target system. The difference is that the purposes of the two kinds of descriptions reflect the differences in the accounts in which they play a role. In the context of a mapping account, these presuppositions serve to pick out a structure, while in the context of RIC, they serve to pick out information about the target system in a form that can figure into the premises of mathematical inference patterns (RIC3).

3.1.3 RIC3: Inference patterns

To understand the collection of privileged inference patterns (RIC3), it is important to distinguish several related notions.

By inference, I mean a move from supporting evidence (premises) to some (possibly intermediate) conclusion. As such, inferences include both simple logical inferences and complex chains of reasoning from a set of premises to a conclusion. In the mathematical case, this means that inferences include whole proofs, individual steps in proofs, and everything in between. I take these to include inferences justified by features of the system under investigation as well—mathematical or otherwise.

By inference pattern, I mean a pattern of reasoning displayed by an inference or collection of inferences. Roughly, an inference pattern is to an inference what a derivation schema is to a concrete derivation in proof theory.³ In taking oneself to be entitled to an inference pattern, one takes oneself to be permitted to make any inference exhibiting that pattern. As the comparison with derivation schemata suggests, inference patterns not only make it more

³This is also very close to the notion of schematic argument in [Kitcher \[1981\]](#).

convenient to specify which inferences are permissible, but also are helpful in reasoning more generally about whatever one is investigating. For instance, the practice in mathematics of reasoning about an arbitrary member of some class in order to show that every member of that class has some property can fruitfully be understood as specifying an inference pattern whose instances are inferences allowing one to infer that a given member of that class has that property and, taken together, show that each member of the class has that property. For both of these reasons, it will be useful to think about applications of mathematics in terms of mathematical inference patterns rather than individual inferences, with “inference pattern” used in such a way that individual inferences count as a special case.

Importantly, the robustly inferential conception strips these inference patterns of their mathematical interpretation in their role as part of RIC3. As a result, the robustly inferential conception treats instances of these patterns as neither mathematical nor physical inferences in their primary role in scientific representations. They are simply a means to generate more uninterpreted statements in the language of the relevant mathematical theories from an initial collection of such statements. Once (at least some of) these claims are provided with a physical interpretation, this serves to fix the accuracy conditions (or, in the language of [Chakravartty \[2010\]](#), the informational content) of the representation. In this role, they are not instances of premise-conclusion reasoning, and so in a sense it is misleading to call them “inferences” or “inference patterns” at all. They are a constitutive part of the *content* of the representation, while genuine inferences are *justified by* that content. In other words, the uninterpreted inference patterns involved in the robustly inferential conception serve to *specify* the accuracy conditions of scientific representations, while genuine inferences about the target system must be justified by those accuracy conditions.

In order to convey this, I will sometimes describe the relevant collection of inference patterns, when they are stripped of their mathematical interpretation, as the *algorithms* for generating additional statements from those in the initial description of the target system or simply the algorithms at work in the representation. As such, talk of “algorithms” should

not be taken very seriously, since I use it here only as a shorthand for “ways of generating additional statements in the uninterpreted language of the mathematical theory from an initial set of such statements.”

To see how mathematical inference patterns figure into RIC, consider how mapping accounts and the robustly inferential conception explain a very simple representation. Consider a system with one spatial dimension containing a single object of mass $m = 1\text{kg}$ with a force $F = 1\text{N}$ consistently applied to it. We assume this system is governed by Newton’s second law of motion (so that it too is included in our initial description of the target system). It follows from this that the object’s acceleration is 1m/s^2 , as we can tell by plugging these values into Newton’s second law of motion and solving for a .

A mapping account will explain this by pointing to a structure-preserving mapping between the structure of the target system and a mathematical structure satisfying all of $F = ma$, $m = 1$ and $F = 1$ —e.g., the reals with the usual arithmetical operations defined on them together with an interpretation that interprets these arithmetical operations in the usual way and that assigns the denotation 1 to both F and m . Using purely mathematical reasoning about the latter structure, we *infer* that $a = 1$ is true in that structure. We then use this together with the fact that this structure is isomorphic to the target system to *infer* that the acceleration of the object in the target system is 1m/s^2 .

According to the robustly inferential conception, the purely mathematical inferences appealed to in the above explanation will be instances of the inference patterns in RIC3. Indeed, all inference patterns appropriate to reasoning about the reals with the usual arithmetical operations, stripped of their mathematical interpretation, should belong to RIC3. Our interpretation of the vocabulary is simply that real values of F represent the force on the object in newtons, that those of m represent the mass of the object in kilograms, and that those of a represent the instantaneous acceleration of the object in m/s^2 . Our initial specification of the target system includes $F = ma$, $F = 1$, and $m = 1$, with these physically interpreted as above. Now, among the inference patterns is one that corresponds to the mathematical reasoning

from $F = ma$, $F = 1$, and $m = 1$ to $a = 1$. Since every physically interpreted claim resulting from applying these inference patterns to our initial specification of the target system is among the physical commitments of our representation, this means $a = 1$, interpreted as above so that it means that the object's acceleration is 1m/s^2 , is among these commitments.

So the reasoning from $F = ma$, $F = 1$, and $m = 1$ to $a = 1$ is treated by mapping accounts as a purely mathematical inference, one that is part of the longer chain of reasoning from the initial specification of the target system to the conclusion that the acceleration of the object is 1m/s^2 . The features of the mathematical structure justify the purely mathematical intermediate inferences, while the mapping between this structure and the structure of the target system justifies the intermediate target-system-to-mathematics and mathematics-to-target-system inferences. In contrast, the robustly inferential conception treats the inference patterns in the purely mathematical part of this chain of reasoning (in part) as algorithms that contribute to determining what the representation says about the target system (i.e., to fixing the accuracy conditions of the representation). In this role, going from $F = ma$, $F = 1$, and $m = 1$ to $a = 1$ does not constitute *reasoning* or *inference* at all.

One might worry here that this picture leaves too little room for genuine inference. In treating the role of mathematics as “algorithmic” rather than inferential, RIC would seem to rule out mathematics-based inferences altogether. But we can certainly genuinely *infer* that the object's acceleration is 1m/s^2 on the basis of doing the relevant mathematics. Fortunately, RIC does not rule this out. While the algorithms or inference patterns in RIC3 are not themselves inferences, *actually carrying out* the relevant computations can certainly constitute an inference. RIC3 determines in part the accuracy conditions of the representation by generating claims about the target system that must be true for the representation to be accurate, whether or not someone actually applies them to the claims in RIC2. In actually applying one of these inference patterns to a particular set of claims in RIC2, one does something else. From the fact that the result of applying certain of these inference patterns is some physically interpreted claim, one can infer (on the basis of the fact that they appear in RIC3) that that

physically interpreted claim is one that the representation is committed to. So, in carrying out the relevant mathematics, one can genuinely infer that the acceleration of the object must be 1m/s^2 . Where mapping accounts and the robustly inferential conception differ in this respect is not in whether this constitutes a genuine inference, but rather in what underwrites that inference. According to mapping accounts, it is the relevant mapping together with the relevant mathematical structure, while according to the robustly inferential conception, it is the physical interpretation together with the collection of inference patterns included in the representation.

3.1.4 RIC as a generalization of mapping accounts

Important both for understanding RIC and especially for the arguments I make later in this chapter is the observation that RIC is strictly more general than the mapping account. For any given mathematical scientific representation that can be represented in terms of a mapping account, we can produce an equivalent representation in terms of RIC that appeals to the same structure and mapping (and consequently attributes the same informational content to it). As a result, any philosophical explanation of a feature of scientific practice that can be expressed in terms of a mapping account can also be expressed in terms of RIC. But RIC also makes possible representations of mathematical scientific representations that are not possible simply with a mapping account.

This is because the structures and mappings that mapping accounts appeal to are, from the perspective of RIC, one way among many of realizing components RIC1–RIC3. In very broad strokes, the mapping serves to specify the physical interpretation (RIC1), structure-generating descriptions or similar devices for picking out the target structure (as discussed in §3.1.2) serve as the initial description of the target system (RIC2), and the mathematical structure serves to specify the class of mathematical inference patterns (RIC3).

More concretely, consider a representation that relates physical structure \mathcal{S}_P and mathematical structure \mathcal{S}_M with a morphism $f : \mathcal{S}_M \rightarrow \mathcal{S}_P$. RIC1 is determined by f in the following

way. If \mathcal{J}_M is an interpretation of the relevant mathematical language for structure \mathcal{S}_M , then we can construct a (possibly partial) interpretation \mathcal{J}_P of that language for structure \mathcal{S}_P by letting $\mathcal{J}_P(x) = f(\mathcal{J}_M(x))$ for all x for which $f(\mathcal{J}_M(x))$ is defined. In this case, RIC1 assigns physical truth conditions to any sentence ϕ in the mathematical language such that $\mathcal{J}_P(\phi)$ is defined, namely those assigned to it by treating \mathcal{J}_P as an interpretation function for \mathcal{S}_P in the ordinary model-theoretic sense. RIC2 consists of the claims in the structure-generating description, expressed in the relevant mathematical language under the interpretation RIC1. RIC3 consists of all inference rules $\Phi \vdash \psi$, where Φ is a set of sentences and ψ is a sentence in the relevant mathematical language, such that $\mathcal{S}_M \models (\bigwedge \Phi) \rightarrow \psi$.

But nothing about RIC requires us to spell out any of these ingredients in this way. In particular, when it is unclear what mathematical structure to associate with an application (or even whether there is such a structure at all), these ingredients can all be spelled out more directly. For instance, if scientists use highly circumscribed inference strategies to reason with mathematically problematic concepts, as discussed in chapters 4 and 5, the collection of inference patterns RIC3 can be specified through a description of these procedures rather than indirectly in terms of a structure (or more likely a highly complex arrangement of structures and mappings). This can be useful even when the mathematics in question is not problematic, but simply not in the business of bearing a structural resemblance to the target system, as in the recent trend of using data-driven computational models. In such cases, the philosophically salient features of the practice concern not so much the relationship between mathematical and physical structures as the role of mathematics in general-purpose computational strategies. (See §8.2.1.) For this purpose again, it is useful to be able to represent such strategies more directly in the specification of RIC3. And once RIC3 isn't represented in structural terms, there is little point in characterizing RIC1 and RIC2 indirectly in those terms either.

Due to the diversity among mapping accounts, this account of RIC as a generalization of mapping accounts may need to be supplemented to show how RIC can accommodate features distinctive to a particular version of the mapping account. In the appendix to this chapter,

I discuss in detail two such features: the distinction between immersion and interpretation mappings and the possibility of nested mappings, both of which play a prominent role in the presentation of Bueno and Colyvan's [2011] inferential conception.

3.2 An initial case for RIC

Having presented RIC in detail, I will now make an initial case for adopting RIC as an account of mathematical scientific representation. I will substantiate and build upon this case through the rest of the thesis.

Because RIC recovers mapping accounts as special cases, adjudicating between RIC and mapping accounts comes down to the question of whether RIC's generality has philosophical benefits that outweigh its costs. If RIC had no such benefits, mapping accounts would be preferable on the basis of their greater specificity. Even though RIC can recover mapping accounts as special cases, if a version of the mapping account sufficed to do all of the philosophical work done by RIC, then we would be better off adopting that "special case" as our full account of mathematical scientific representations.

There are three broad kinds of benefit that RIC might have over mapping accounts.

1. RIC might do better in the first task for accounts of mathematical scientific representations—that is, in explaining how mathematics can in principle be used to represent non-mathematical target systems—because it can do so in a wider range of cases than mapping accounts can.
2. RIC might do better in the second task for accounts of mathematical scientific representations—that is, as a meta-level device used by philosophers of science to represent philosophically salient features of scientific practice—because it more perspicuously or otherwise more usefully represents those features.
3. Finally, even if RIC didn't have benefits of the first two kinds, it might be thought to

have the benefit of doing the same work as the mapping account while requiring fewer assumptions about the nature of mathematics itself.

In the thesis as a whole, I focus on showing that RIC has the second of these benefits in a wide range of cases, as I think it is the most important of the three. That said, I think there is good reason to think it has benefits of each kind. In the rest of this section, I present an initial case for each.

3.2.1 RIC is applicable in a wider range of cases

The first benefit of RIC is that it can achieve the first task for accounts of mathematical scientific representations—explaining how in principle mathematics can be used to represent physical target systems—in a wider range of cases than mapping accounts. The extent of this benefit depends a great deal on the version of the mapping account in question.

Very simple versions of the mapping account fail to achieve this first task in a wide range of cases in which the mathematics applied is inconsistent, unrigorous, or otherwise cannot be straightforwardly associated with a structure that can straightforwardly be related to the target system. These versions of the mapping account leave unexplained how the mathematical structure ultimately mapped to the structure of the target system is picked out.

In many cases, this is not a problem at all, as the relevant structure is obvious. For instance, from the point of view of a mapping account, applications of rigorous number-theoretic results probably should always be understood in terms of the set \mathbb{N} of natural numbers with the relevant number-theoretic properties and relations defined on it, and applications of rigorous real analysis should probably be understood in terms of \mathbb{R} with the relevant arithmetical and analytical operators, properties, and relations defined on it.

But in a wide range of cases, it is not obvious how to understand the relevant structure and its relation to the relevant mathematical practice. This is most obvious in the case of applications of inconsistent mathematics, as discussed in chapter 4. As no classical structure satisfies an inconsistent theory, either the relationship between the inconsistent theory and

the relevant classical structure stands in need of further explanation or non-classical structures must be admitted, and the relationship between mathematical practice and non-classical structures will stand in need of explanation. The same goes for applications of mathematics that don't meet the standards of rigor of recent pure mathematics because some of the concepts used are incoherent or underspecified, as discussed in chapter 5. In such cases, if the mathematics could be straightforwardly associated with a classical structure in all cases, then these conceptual shortcomings would be resolved. As such, the presence of such concepts in an application of mathematics is a sign that more needs to be done to relate the mathematics in question to a structure appropriate for applications.

In contrast, as I show in chapters 4 and 5, RIC can straightforwardly represent applications of inconsistent and otherwise unrigorous mathematics in terms of the inference strategies used to manage the inconsistent or otherwise problematic concepts in the relevant mathematics. While RIC makes it possible to represent an application in terms of structures and mappings, it also makes it possible to specify the patterns of inference allowed in an application in other terms—including, in the limit, by enumerating these patterns of inference one-by-one. This makes it remarkably flexible in representing applications of mathematics, in particular by allowing it to explain how a given piece of mathematics can be used to represent a physical target system without first associating the mathematics with a well-defined structure.

On the other hand, there are more sophisticated versions of the mapping account that address these issues. For instance, [Pincock \[2012\]](#), rather than taking the mathematical structure for granted, explicitly examines the processes by which the content of a representation is refined through adjustment of the relevant structure and mapping. And [Bueno & Colyvan \[2011\]](#) appeal to partial structures, a type of non-classical structure intended to represent inconsistency, uncertainty, and conceptual change in science (among other things). Such accounts have fewer issues in representing the informational content of unrigorous and inconsistent representations (and so in achieving the task of explaining how in principle mathematics is applicable in those cases).

However, these more sophisticated versions of the mapping account still have some problems in a narrower range of cases, which I discuss in greater detail in chapter 4. In particular, an approach in terms of classical structures (like Pincock's) can't explain cases in which inconsistent mathematical concepts are themselves given physical interpretation and so in particular can't explain the use of inconsistent mathematics to produce inconsistent representations (§4.2.3). An approach in terms of partial structures does better, in that it can represent cases in which inconsistent concepts are given a physical interpretation. However, I argue in §4.2.1 that such accounts still have two problems. First, while they get the accuracy conditions of inconsistent representations right on an *extensional* construal of their content, they get them wrong on a more fine-grained *intensional* construal of their content. Second, they can't represent a case in which a scientist posits a physical correlate to an inconsistent mathematical concept, but doesn't take a stance on many of its properties, so that the physical concept is not inconsistent (because it has so little content). While these are narrow ranges of cases, they are nonetheless important. For instance, there was an extreme diversity of thought after the introduction of the early, inconsistent calculus (the case study of chapter 4), and it's not at all clear that the concept of infinitesimal wasn't thought to have a physical interpretation by at least some early practitioners. Representing this state of affairs in terms of an account of mathematical scientific representations would seem to require more flexibility than even the most flexible forms of mapping accounts allow.

3.2.2 RIC does better as a meta-level representational device

A more significant benefit of RIC over more sophisticated versions of the mapping account is its success in achieving the second task for an account of mathematical scientific representations—serving as a meta-level device for philosophers of science to use to represent philosophically salient aspects of scientific practice.

As discussed in §2.3, this might at first seem unpromising as a way to adjudicate between accounts of mathematical scientific representation, as it might be that an account does better

than all the alternatives in some cases but worse in others. If so, that would be weak evidence (if evidence at all) for any account—or perhaps evidence instead for a pluralist account, as in the literature on pluralism about truth.⁴ However, since mapping accounts come out as special cases of RIC, in which RIC1–RIC3 are spelled out in terms of structures and morphisms, RIC can always directly coopt the tools of a given version of the mapping account in its role as a meta-level device for representing scientific practice. This means that when mapping accounts are particularly well-suited to serve as a meta-level device for a given set of purposes, RIC can play that role equally well by appealing to exactly the same configuration of structures and mappings. Now, if it simply coopted the tools of the mapping account in all cases, this would again simply favor mapping accounts. So the question is whether RIC can do significantly better than mapping accounts in at least some cases. The next part of the thesis is primarily devoted to arguing that it can.

Now, again as noted in §2.3, what it means for an account of mathematical scientific representations to be a better meta-level representational device varies considerably from case to case, as it depends crucially on the philosophical aims of those using the account. What we can say in general is that an account does better in this respect in a given case to the extent that it more perspicuously represents the features of the given scientific practices relevant to the user’s philosophical purposes. Of course, in addition to the kinds of features of practice that are relevant to a given set of purposes, what it means to represent those features perspicuously will also vary considerably from case to case. As a result, the bulk of this argument must wait until I discuss concrete case studies in the next part of the thesis.

That said, we are already in a position to see that this is a promising line of argument. Owing to its flexibility, RIC can be used to represent scientific practice at a remarkably wide range of levels of grain. At one extreme, it can be used to represent individual inferences of individual scientists by enumerating the set of privileged inference patterns (RIC3) one-by-one. At the other end of the spectrum, the inferences licensed by a representation shared by

⁴See [Pedersen & Wright, 2018] for a useful introduction to this extensive literature.

a much larger community or multiple communities can be represented in a highly general way—for instance, by specifying RIC3 in terms of the inferences licensed in reasoning about a particular, well-understood mathematical structure. In the cases considered in the next part of the thesis, an intermediate level of grain is called for. Philosophical features relevant to the practice of particular scientists and scientific communities are most straightforwardly represented in terms of more local inference strategies than the mapping account can represent directly. Even when mapping accounts can represent these strategies indirectly, they do so, as Dirac wrote in an entirely different context, “only in a cumbersome way which would tend to obscure the argument” [1967, §15, p. 59]. This is of course not to say that every representation in terms of RIC is even at least as good in this respect as an equivalent representation in terms of a mapping account. For instance, a representation in terms of RIC that simply enumerates one-by-one all the allowed mathematical inferences will likely never be the most perspicuous representation of an episode of scientific practice. The strength of RIC as a meta-level device is not that it is maximally fine-grained, but that it allows for a wide-range of levels of grain, which can be chosen according to the purposes of the user.

3.2.3 RIC makes fewer assumptions about the nature of mathematics

Finally, provided that RIC does at least as well in these other respects, we have one further reason to prefer RIC: it does the work of an account of mathematical scientific representation without making any substantial assumptions about the nature of mathematics. What RIC requires from mathematics is only that it provides us with a collection of (uninterpreted) inference patterns (RIC3) that we can use to make physical inferences by physically interpreting the language in which they are expressed (RIC1). This requires only that mathematics involves something that *looks like* inference. And this commitment is shared by *all* accounts of mathematics, from the most Platonic platonism, to mathematical error theory, to the most extreme formalism, according to which mathematics is just the manipulation of uninterpreted strings of symbols. This gives us reason to favor RIC over any account that does equally well

in all other respects, but makes any substantial assumption about the nature of mathematics. RIC might still have a bearing on questions about the nature of mathematics by divorcing them from questions about applied mathematics⁵ but is ultimately entirely neutral.

In contrast, the commitments of mapping accounts are fairly contentious. On the one hand, mathematical representation is usually thought (or at least supposed for the sake of argument) to be ontologically innocent. For instance, the debate over indispensability arguments for platonism has largely shifted from representational to explanatory indispensability arguments in part for this reason.⁶ On the other hand, mapping accounts would seem to favor platonism if taken at face value, since they posit the existence of a mathematical structure which is mapped to the structure of the target system. Recently, [Heron \[2020\]](#) has taken this line, using the face-value commitments of mapping accounts to support a new representational indispensability argument for mathematical platonism. And others who do not explicitly endorse such an argument nonetheless have commitments that would make any straightforward understanding of the mapping account ontologically committing. For instance, [Baron \[forthcoming, §5\]](#) argues that his account of mathematical explanation is incompatible with nominalism on the grounds that structural properties, even of physical systems, make indispensable reference to mathematical objects. (Earlier in the paper (§2), he even cashes out the notion of these structural properties in terms of the mapping account.) And even those who take mapping accounts to be ontologically innocent nonetheless make at least some substantial assumptions about the nature of mathematics. For instance, [Pincock \[2012, chapters 9–10\]](#) argues that his version of the mapping account supports truth-value realism (but not ontological realism) about mathematics. Finally, even if it should turn out that mapping accounts are entirely neutral concerning the nature of mathematics, RIC would still be preferable on the grounds that it makes this neutrality obvious, not requiring significant reconstruction or explication of its central concepts to demonstrate this neutrality.

⁵For example, I take the argument in chapter 6 to show among other things that we should do just this.

⁶See, e.g., [\[Baker, 2005\]](#).

3.3 Scientific representation and RIC

Finally, I turn to the relationship between RIC and theories of scientific representation more generally. My aim here is to clarify how RIC is situated in the context of broader debates over scientific representation and in so doing show that there is good reason to adopt RIC regardless of one's take on scientific representation more generally.

In §3.3.1, I lay out the deflationary inferentialist case for RIC and argue that RIC, unlike even nominally “inferentialist” mapping accounts, coheres with the motivations for deflationary inferentialism. Now, because RIC is a natural ally of deflationary inferentialism, it might be thought to be incompatible with the motivations underlying substantive views of scientific representation. In §3.3.2, I present two ways in which to preserve these motivations while adopting RIC. The first is to recognize the conceptual space between deflationary and substantive accounts of representation and understand RIC as a view that occupies that space. The second is to understand RIC as an important supplement to a substantive account of scientific representation.

3.3.1 RIC as the only account in keeping with the motivations of deflationary inferentialism

The account of scientific representation most naturally aligned with RIC is Suárez's [2004, 2015] deflationary inferentialism about epistemic representation. This view is deflationist in the sense that representation is not taken to be a substantive, explanatory relation [Suárez, 2015]. Instead, Suárez takes it to be best accounted for in terms of core, abstract, surface-level features of scientific representations—features that are concretely realized in radically different ways in practice. Suárez's account is based on two platitudes about scientific representations, roughly: (i) scientific representations are *about* their targets and (ii) scientific representations can be used to make inferences about their targets. And so, correspondingly, Suárez describes necessary conditions for (epistemic) representation as follows: “A represents

B only if (i) the representational force of *A* points towards *B*, and (ii) *A* allows competent and informed agents to draw specific inferences regarding *B* [Suárez, 2004, p. 773]. The inferential character of the account then comes from the second platitude, which concerns surrogative inference.

Suárez also takes his view to be deflationary in the sense that it rejects the project of providing individually necessary and jointly sufficient conditions for epistemic representation. Suárez takes (i) and (ii) above to be necessary but not jointly sufficient for epistemic representation. I do not wish to follow him in this. Even if, as Suárez claims, “[i]n every specific context of inquiry, given a putative target and source, some stronger condition will typically be met” [Suárez, 2004, p. 769], this needn’t mean that the conditions he’s already given are not jointly sufficient, but only that, when they are met, some condition from a class of more specific conditions is also met. Indeed, I cannot conceive of a case in which these conditions are met for putative vehicle *V* and target *T* such that *V* fails to be an epistemic representation of *T*. So, unlike Suárez, I do take (i) and (ii) to be jointly sufficient for epistemic representation.

Such an account can accommodate the diversity in the practices of producing and working with scientific representations, but at the cost of being largely uninformative. (This is by design. By definition, a deflationary analysis of a concept won’t “shed explanatory light on our use of the concept” [Suárez, 2015, p. 39].) Nonetheless, it does allow for more informative accounts of narrower ranges of representations, as it can be supplemented by accounts appealing to more concrete properties shared by representations of these narrower kinds. Indeed, since it says relatively little about scientific representations in full generality, it is compatible with an extremely diverse range of accounts of more specific kinds of scientific representation. I believe RIC, as I presented it above, is the best way to provide a more informative account of MSRs while maintaining the philosophical ideas and inclinations motivating deflationary inferentialism.

The core idea behind RIC essentially involves extending deflationary inferentialism as an

account of what makes something an epistemic representation to an account of what makes something a *faithful* or *accurate* epistemic representation. Since, according to deflationary inferentialism, a representational vehicle V represents its target T in part in virtue of allowing users to make inferences about T on the basis of V , it is then natural to say that V is an *accurate* representation of T if and only if (i) V is a representation of T , of course, and (ii) the inferences V licenses about T are good (i.e., truth-preserving) inferences [cf. [Contessa, 2007, 2011](#)]. In this way, Suárez's platitudes about scientific representation give us not just an inferential account of what a scientific representation *is*, but also a basic inferential semantics for scientific representations. A useful way to understand most substantive accounts of scientific representation is as supplementing this basic inferential semantics with an explanation of what must be in place to ensure that these inferences are truth-preserving—e.g., the existence of relevant similarities between V and T (similarity accounts of representation) or the existence of a structure-preserving mapping between structures instantiated by V and T (structuralist accounts of representation)—a move that the deflationary inferentialist ought to reject given the commitments that lead them to accept deflationary inferentialism about what epistemic representations *are*.

It is for this reason that even nominally inferential versions of the mapping account—particularly Bueno and Colyvan's [2011] inferential conception—are inconsistent with the philosophical principles that motivate deflationary inferentialism. While the Bueno-Colyvan inferential conception emphasizes the role of mathematics in facilitating the surrogative inferences characteristic of scientific representations, as understood by the inferentialist, it insists on these inferences' all being underwritten by the same kind of substantive relation (a partial morphism between structures taken to be instantiated by the vehicle and target). RIC, in contrast, makes no such move. In some cases, the relevant inferences might be best understood as underwritten by morphisms, while in other cases—including those I discuss in chapters 4 and 5—the relevant inferences might have a more local justification or, in extreme cases, may be almost entirely ad hoc.

However, while more general than accounts like the Bueno-Colyvan inferential conception—and so more readily able to account for the diversity of uses of mathematics in scientific practice, consistent with the motivations for deflationary inferentialism about scientific representation—RIC can provide a much more informative account of *how* mathematics facilitates surrogate inferences than the deflationary inferentialist view on its own. These inferences will (i) share their form with (some of) the purely mathematical inferences available to someone using the mathematical theory on its own. However, in order for conclusions to be drawn about the target, (ii) they must be given an interpretation in terms of the target. Like the claim that epistemic representations can be used to make inferences about their targets, which is central to the inferential account of scientific representation, these claims are about as close to platitudes as we can find about mathematical scientific representations, and these claims are ultimately what are encoded in the ingredients of such a representation as understood in terms of RIC. The partial physical interpretation of the mathematical vocabulary (RIC1) captures the second of these platitudes, while the set of privileged inference patterns (RIC3) drawn from the relevant mathematical theory capture the first. (The remaining ingredient of a mathematical scientific representation understood in terms of RIC—viz., an initial description of the target system (RIC2) in the language of the mathematical theory—is just there to provide an initial set of premises for these inference patterns to work on.)

Recall that the mapping account can also be understood in terms of these ingredients—i.e., as one way (among others) in which the abstract conditions described by RIC might be realized. The partial physical interpretation of the vocabulary is provided by the posited mapping, the relevant inference patterns are just those that correspond to truth-preserving mathematical inferences about the given mathematical structure, and the initial description of the target system is given by a structure-generating description or similar device. These are all aspects of any plausible formulation of the view. But this, according to RIC, is simply one way in which the more abstract properties fundamental to mathematical scientific representations can be realized. So, while RIC incorporates features specific to mathematical scientific representa-

tions, it is deflationary in the sense that it builds its account of what these representations are, as well as its semantics for them, out of abstract platitudes without proposing to explain them in terms of a single, more concrete, substantive property (though such properties will be there in particular cases). As a result, as I take the case studies in the next part of the thesis to show, it better accounts for the diversity among practices of applying mathematics in science.

3.3.2 RIC and the motivations behind substantive accounts of scientific representation

In the previous subsection I argued that RIC is the view that is most consistent with the motivations behind deflationary accounts of scientific representation and Suárez's deflationary inferentialism in particular. While this makes RIC more attractive to someone who has already accepted a deflationary account of scientific representation, it might initially make the view less attractive to those who have already accepted a substantive account of scientific representation, such as a structuralist or similarity-based view. In this subsection I argue that RIC is consistent with the motivations underlying substantive accounts of scientific representation.

There are two ways in which a substantivist about scientific representation might adopt RIC. The first is to recognize a conceptual space between deflationary and substantive monist accounts of scientific representation, much like the conceptual space that allows for pluralism as an alternative to both deflationary and substantive monist theories of truth. Many of the pluralist accounts of scientific representation that this makes possible successfully reconcile the motivations underlying substantivism with adoption of RIC. The second is to maintain a substantive monist theory and treat RIC as an account not of how mathematics represents physical target systems, but instead of how it can be used to specify and reason about the representational vehicles posited by the relevant theory of scientific representation.

RIC and pluralism about scientific representation

The first way of reconciling adoption of RIC with the motivations underlying substantive accounts of scientific representation is to adopt a form of pluralism about scientific representation. Influential substantive views are monist in that they pick out a single substantive property—like similarity or shared structure—that underlies (and explains) all instances of scientific representation. If RIC as described so far is correct and complete as an account of mathematical scientific representation, then it would seem to be incompatible with such views, as RIC explains a particular type of scientific representation, mathematical scientific representation, without having to appeal to similarities or structural relations between representational vehicles and their targets. One way to respond to this would be to adopt a deflationist account of scientific representation as described above. But this is not the only—and perhaps not even the best—possible response. An alternative, which I sketch here, is to adopt a pluralist account of scientific representation.

Like certain varieties of pluralism about truth (e.g. [Wright \[1992\]](#) and [Lynch \[2009\]](#)), the pluralism I have in mind shares with deflationism the conviction that all we can do to characterize all instances of scientific representation is to characterize its functional role through what are ultimately platitudes about the concept of representation. Where the pluralist disagrees with the deflationist is in holding that we can give a substantive account of narrower classes of scientific representation in terms of substantive properties that realize this functional role in this narrower range of cases.

Start from the deflationary inferentialism of [Suárez \[2004, 2015\]](#), according to which S represents T only if (1) the representational force of S points to T and (2) S allows competent and informed agents to draw specific inferences regarding T . (For the purposes of this discussion, let's even make that a biconditional!) This provides a basic functional characterization of scientific representation.

We can get from this functional characterization to the informational content of the representation via a basic inferential semantics: S (perfectly) accurately represents T if and only

if every inference regarding T a competent and informed agent can make on the basis of S is truth-preserving. The informational content of the representation is what must be true of T for S to (perfectly) accurately represent it. (Cf. the account when a representation is “true” in [Suárez \[2004, p. 776\]](#).)

This would make representation mysterious if the claim were that there was never anything more to say about how S comes to have its informational content about T . But all the functionalist definition above, together with a rejection of substantive monism, commits one to is that there is nothing more to say about how S comes to have its informational content *at a maximal level of generality*. It leaves it open that substantive relations between S and T explain why the inferences regarding T licensed by S are truth-preserving for some more restricted class of representational vehicles and targets. So the difference between the substantivist monist and the deflationary inferentialist is not that the former but not the latter can explain how representations support surrogative inferences, but rather (primarily) that they disagree on the degree of generality at which such an explanation is possible. The substantivist monist thinks that such an explanation is possible in full generality in terms of their favored substantive property. The deflationist thinks that such an explanation is impossible in full generality, but not necessarily that such explanations are impossible for narrower ranges of cases.

Some deflationary inferentialists explicitly endorse this. [Kuorikoski & Ylikoski \[2015, p. 3828\]](#) write, “It is certainly not arbitrary that a specific diagram, set of equations, or physical scale model is more helpful in inferring about a specific target phenomenon than some alternatives. And this depends as much on the intrinsic properties of the inferential apparatus as it does on the cognitive make-up and perceptual abilities of the model user.” Suárez similarly claims that on his account, a representational source must be “inferentially suited” to its target [[Suárez, 2004, p. 778](#)].

Once we recognize this, an underexplored spectrum of philosophical views emerges, differing in the level of generality at which they allow that substantive properties can explain

how scientific representations support surrogate inferences. On one end are substantive monist views, which hold that we can explain this for all scientific representations in terms of their favored substantive property. On the other end is a radical particularism according to which such explanations can be given for only maximally specific contexts—perhaps as specific as a particular inference made by a particular user of a representation. In the middle are a range of pluralistic views that hold that such explanations can be given for certain types of episodes from scientific practice.⁷

The higher the level of generality of a given account on this spectrum, the better it is able to recover substantivists' explanations of functional properties of scientific representations in particular cases. And we should expect RIC to be compatible with very high-level pluralist accounts of this kind. As a result, RIC is compatible with a range of accounts which largely recover the substantive explanations put forward by substantive monist accounts while providing alternative explanations in cases in which the substantive account is lacking. In an important sense, this preserves the motivation behind substantive monist accounts while abandoning the letter of them.

Here it is also worth noting that such pluralist accounts come in both deflationist and substantivist forms. The deflationist's primary disagreement with a substantive monist concerns the level of generality at which explanations of representations' functional properties in terms of substantive properties are possible. But for the "deflationist" label to be appropriate, such views should also deny that the substantive properties that explain the functional properties of representations in narrower ranges of cases themselves *constitute* representation relations. The substantivist pluralist alternative is to take these substantive properties to constitute (a plurality of) representation relation. They do so on this view precisely because

⁷Some deflationary inferentialists downplay the philosophical relevance of these relations. For instance, Kuorikoski & Ylikoski [2015, p. 3828] write "These (perfectly objective) dependencies between the properties of external inferential apparatuses and their possible applications, i.e., the ways in which cognitive agents can perform inferential tasks with different kinds of external aids, are empirical and therefore proper objects of study for cognitive science, not philosophy." But this is much too hasty. While the cognitive activity of human agents is certainly part of the story about how a given source can be used by human agents to make inferences about a target, that does not mean that the story as a whole should be left to cognitive science.

they realize the functional role constitutive of the concept of scientific representation and as a result bear a closer conceptual relationship to scientific representation than the deflationist allows. In each instance, the “representation role” is played by a range of properties of varying degrees of specificity, and which is of philosophical interest depends on the level of generality at which we are working. The result is an analogue of the many-one pluralism about truth [Lynch, 2009], which perhaps even better recovers the motivations of substantivist monist views.

RIC as a component of substantive accounts of scientific representation

While deflationary and substantive forms of pluralism each have a strong claim to recovering many of the motivations of existing substantive accounts of scientific representation, they still amount to a rejection of those accounts. Alternatively, with minor modifications, RIC can be recast in terms of one’s preferred substantivist account. In brief, if an account of scientific representation explains how a representational vehicle V represents a target T , an account of mathematical scientific representation serves to explain how mathematics helps to realize the conditions picked out by the account of scientific representation. And we can understand RIC as doing so by treating the informational content described by RIC as characterizing V rather than T . In practice, this means the “physical” interpretation RIC1 will be in terms of V rather than T .

For instance, even if we already accept structuralism about scientific representation, mapping accounts still have further work to do—despite the fact that mapping accounts are natural analogues of structuralism about scientific representation in the specific case of mathematical scientific representation. Mapping accounts go further by specifying that the relevant source structures are either drawn directly from the mathematics or are appropriately morphic to a structure provided by the mathematics. Now, the cases I use to support RIC are largely cases in which it is not straightforward to associate the mathematics used in a representation with a well-defined structure. In such cases, further work is needed to understand the rela-

tionship between scientists' mathematical reasoning and the source structure of the relevant scientific representation, and mapping accounts are of little help. On the other hand, if we understand the informational content posited by RIC as concerning not the target structure of the representation but instead the source structure, then RIC can explain how mathematics helps to realize the conditions on scientific representation specified by structuralism in a way that mapping accounts cannot. And since RIC recovers mapping accounts as special cases, it can appeal to a morphism between a mathematical structure and a representational vehicle in cases well-handled by mapping accounts. The result, I think, is an enrichment of the structuralist account, rather than an abandonment of its core principles.

The same goes for other substantive views. RIC can be understood as characterizing how mathematics is used to specify and reason about the relevant representational vehicle if we treat RIC1 as concerning that vehicle rather than the target system itself. As a result, RIC can be made to function in the context of any substantive account of scientific representation and so ultimately doesn't require such accounts to be abandoned.

3.4 Conclusion

In this chapter, I have presented a novel account of mathematical scientific representations, the robustly inferential conception. I presented three central arguments for RIC: (1) it can be used in a wider range of cases than mapping accounts; (2) it is more useful as a meta-level device for representing scientific practice than mapping accounts; and (3) it does so in a way that is manifestly neutral regarding the nature of mathematics, including its ontology. In the next part of the thesis, I substantiate the first two of these arguments, with a focus on the second, by discussing a number of detailed case studies.

Appendix: Some complications in understanding mapping accounts as special cases of RIC

In this appendix, I discuss two features of some kinds of mapping account, particularly Bueno and Colyvan's [2011] inferential conception, that complicate the discussion in §3.1.4 of how RIC recovers mapping accounts as special cases: the distinction between immersion and interpretation mappings and the possibility of nested mappings.

Immersion and interpretation mappings

Some versions of the mapping account, particularly the inferential conception [Bueno & Colyvan, 2011], treat mathematical scientific representations in terms of two morphisms. There is an immersion mapping $f : \mathcal{S}_P \rightarrow \mathcal{S}_M$, which serves to mathematize the physical system, facilitating an inferential move from physical claims to mathematical ones. And then there is a separate interpretation mapping $g : \mathcal{S}_M \rightarrow \mathcal{S}_P$, which licenses inference from claims about the mathematical structure to claims about the physical structure. This may require us to complicate our picture of how RIC can capture mapping accounts as special cases.

Where f and g are simply inverse mappings, the discussion in the previous section extends to this version of the mapping account without any changes. The RIC representation of the case only needs to appeal to the interpretation mapping.

Where f and g are not inverse mappings, we have options. One is to continue to prioritize the interpretation mapping in spelling out the physical interpretation RIC1, treating cases in which f and g are not inverse mappings as cases in which mathematical reasoning motivates a change to a new representation with a modified physical interpretation. Indeed, I suspect that this is what is going on in cases that motivate Bueno and Colyvan's decision to allow for cases in which f and g aren't inverse mappings. This again allows the discussion in the previous section to apply unchanged.

A further option requires modest modifications to the discussion above but captures the

full range of possible cases in which f and g are not inverse mappings. In this case, RIC2 consists not of the structure-generating description per se, but of the mathematical expressions related to the claims in the structure-generating description by the mapping. That is, if ϕa is in the structure-generating description and the immersion mapping takes the denotation of ϕ to ϕ_M and the denotation of a to a_M , then the mathematical expression for $\phi_M a_M$ must appear in the initial description of the target system RIC2.

Nested mappings

A further complication is that some mapping accounts—again particularly the inferential conception [Bueno & Colyvan, 2011]—also directly represent various intermediate mappings between various mathematical structures and the structure that is ultimately mapped to the target structure. This work again can be done in several ways in terms of RIC.

As I see it, the most plausible option is to decline to directly represent any of these additional structures or mappings at all. In such a case, RIC3 remains the set of inference patterns licensed in reasoning about whatever structure is directly mapped to the target system. These inference patterns include those licensed by reasoning about further mathematical structures mapped to the first, but the way in which purely mathematical inferences about this structure are licensed is left unanalyzed. This doesn't strike me as a major drawback, since accounts of mathematical scientific representation primarily concern how such representations facilitate *mathematically-mediated* physical inferences rather than purely mathematical inferences (i.e., those with only mathematical premises and conclusions). It seems only right that explaining how these purely mathematical inferences work should require us to appeal to a further account of what licenses such inferences *within* mathematics.

That said, RIC is flexible enough that we can explicitly specify the inferences in RIC3 in terms of these structures and mappings, explicitly stating that RIC3 includes those inference patterns licensed by this particular configuration of structures and mappings. But, for the

reasons discussed in the previous paragraph, the value of doing so is dubious.⁸

Finally, these nested mappings are most crucial to mapping accounts in cases where the mathematics applied is inconsistent or otherwise unrigorous, thus requiring a complex network of structures to represent the more local and ad hoc inference strategies scientists use to apply such mathematics safely and effectively. In chapters 4 and 5, I argue that in such cases we better represent scientific practice by more directly representing these inference strategies in terms of RIC1–RIC3.

⁸Further, as I discuss in the conclusion (chapter 8), I expect that something very similar to RIC can be usefully applied to model reasoning within mathematics, especially prior to the 19th and 20th centuries. This would allow for a treatment similar to that of [Bueno & Colyvan \[2011\]](#), who can model both mathematically mediated inferences and inferences within mathematics. But again, for the reasons previously stated, I suspect this is of dubious value in giving an account specifically of mathematical scientific representations and so is well beyond the scope of this thesis.

Part II

Case Studies

CHAPTER 4

Applying Inconsistent Mathematics: The Early
Calculus

Philosophical work on applications of mathematics in the empirical sciences has largely ignored the application of inconsistent mathematical theories. But such applications are not altogether rare in the history of science. Indeed, much of the best scientific work done between the late seventeenth and early nineteenth centuries involved applying an inconsistent mathematical theory, the early calculus, to produce consistent representations of the physical world. In this chapter, I argue that these cases give us good reason to favor RIC over mapping accounts.

In §4.1, I present the central case study of the chapter, the early calculus, and argue for a particular construal of its inconsistency. In §4.2, I consider several ways in which mapping accounts might attempt to make sense of this case and argue that all of them have shortcomings. In §4.3, I show how we can better make sense of these cases in terms of RIC. In §4.4, I argue that we should take this to show not that RIC is needed as a sort of “patch” to be added to mapping accounts, but instead that we should embrace RIC over mapping accounts.

4.1 The Early Calculus

To make the following discussion more concrete, I will focus on perhaps the best known application of an inconsistent mathematical theory in the history of science: the application of the early infinitesimal calculus. The calculus was a central part of physics following its development by Leibniz and Newton in the late seventeenth century, but it had a problem. Its central techniques involved taking infinitesimals to be quantities that were zero in some cases (to eliminate terms) but non-zero in others (so that one could divide by them) within the very same proof.

For example, consider the function $f(x) = x^2$. Its derivative, according to the early calculus, is

$$f'(x) = \frac{(x + \epsilon)^2 - x^2}{\epsilon}$$

where ϵ is an infinitesimal. For this expression to be well-defined, ϵ must be non-zero, since

it involves division by ϵ . Expanding $(x + \epsilon)^2$ and simplifying yields

$$f'(x) = \frac{\epsilon(2x + \epsilon)}{\epsilon} = 2x + \epsilon.$$

Now we eliminate the remaining ϵ , yielding $f'(x) = 2x$. To carry out this step, we treat ϵ as zero (since we need $2x + \epsilon = 2x$). So we've appealed to inconsistent information (that $\epsilon \neq 0$ and that $\epsilon = 0$) in the course of a single proof.

Despite making appeal to these inconsistent pieces of information, the techniques of the early calculus were remarkably useful, producing results that allowed for the solution of previously intractable mathematical problems and for the formulation of powerful new representations of the physical world. These results could be reliably achieved without leading to absurdities like $1 = 2$ because mathematicians placed significant constraints on when infinitesimals and their inconsistent properties could be appealed to.

Infinitesimals were generally only used within the calculation of a derivative or integral, and the information that $\epsilon = 0$ and that $\epsilon \neq 0$ could only be appealed to at certain points in the process of doing so (and never at the same time). For instance, in calculating a derivative, the information that $\epsilon \neq 0$ could only be used at the beginning, when division by ϵ was necessary. It was only once one had simplified this expression, so that no remaining terms involved division by ϵ , that one could use the information that $\epsilon = 0$ to eliminate terms. Because one was never allowed to reason with both of these inconsistent pieces of information at once, and because only the final result of one's reasoning with the information that $\epsilon \neq 0$ could be used when one reasoned with the information that $\epsilon = 0$, the set of information with which one could reason at any given time was consistent.

As a result of these restrictions, one could reason classically without proving falsehoods via *ex contradictione quodlibet* inferences. So one's reasoning was locally consistent in that at a given time one was always reasoning with a consistent set of information. But it was globally inconsistent in that the information used in the course of an entire proof was inconsistent

when collected together.¹ If we ask what the world would have to be like for this reasoning to be justified, it seems we must collect up all of the information appealed to in this reasoning. Because the reasoning involved in using the early calculus appealed to globally inconsistent information, this picture is an inconsistent one.

It turned out that the calculus could be put on a consistent foundation at the cost of eliminating infinitesimals and appealing instead to the modern ϵ, δ definition of a limit, a project largely carried out by Cauchy, Bolzano, Riemann, and Weierstrass.² But this came over 150 years after the introduction of the calculus and wouldn't become widely adopted until the second half of the nineteenth century.

That the early calculus was inconsistent is not uncontroversial. In particular, [Vickers \[2013, ch. 6\]](#) argues that those who worked with the early calculus did not reason with inconsistent propositions in the way I describe above, but instead simply applied algorithms to calculate derivatives and integrals. Since neither the results of applying these algorithms nor the (inadequate) justifications for the algorithms were inconsistent, there is little sense in calling the early calculus inconsistent according to Vickers.

A thorough discussion of Vickers's arguments here would take us too far afield, as the main task of this chapter is not to show that the early calculus was inconsistent, but rather to examine how to make sense of applications of theories like the early calculus if we accept that they are inconsistent. That said, I believe that we can reconcile the inconsistency of the early calculus with Vickers's observations about the practice of those who used it. There is a sense in which the early calculus was algorithmic, but this is compatible with its being inconsistent in the sense I describe above. Consider Vickers's description of a simple algorithm for calculating derivatives:

¹This is essentially the picture of things formalized in Brown and Priest's [2004] *LN*, an implementation of their chunk-and-permeate strategy. One's information is divided into two chunks, one containing $\epsilon \neq 0$ and the other containing $\epsilon = 0$. One can reason classically with the information in each of these chunks, taken separately. One calculates $f'(x) = \frac{f(x+\epsilon) - f(x)}{\epsilon}$ and simplifies the resulting expression within the first chunk until one eliminates division by ϵ . The result of this reasoning is then allowed to "permeate" into the second chunk, where one can eliminate multiples of ϵ using the information that $\epsilon = 0$.

²For a thorough discussion of this history, see [\[Kline, 1972, ch. 40\]](#).

1. Put your equation in the form $y = f(x)$
2. Calculate $\frac{f(x+\epsilon)-f(x)}{\epsilon}$, and simplify.
3. Remove any terms which are multiples of ϵ .

The resulting term is then the derivative. [2013, p. 150]³

We can understand such algorithms perfectly well as prescribing how to reason with inconsistent propositions (about the inconsistent properties of naïve infinitesimals) in a way that produces the desired results of the calculus without leading to absurdity (beyond these inconsistent properties themselves) in just the way I describe above. In the case of calculating a derivative using this algorithm, the proposition that $\epsilon \neq 0$ may only be used in carrying out step 2, while the proposition that $\epsilon = 0$ may only be used in carrying out step 3. Since only the final result of step 2 may be used in carrying out step 3, one cannot introduce information inconsistent with $\epsilon = 0$ in step 3, and so local consistency is achieved.

At several points, Vickers seems to suggest that this cannot be the right story because those who used the calculus were simply following the algorithms and so could not have been reasoning with inconsistent propositions. For instance, he writes,

In the early calculus, scientists couldn't possibly derive a contradiction if they were making derivations by applying the relevant algorithms. Indeed, they were not reasoning with inconsistent propositions at all in such a case, but rather following a procedure. [2013, p. 182]

But certainly following an algorithm or procedure does not preclude (thereby) reasoning with propositions. For instance, in non-standard analysis, proving, say, that the derivative of x^2 is $2x$ involves following (almost⁴) exactly the steps of the algorithm described above. But in this case, it seems evident that carrying out the proof involves reasoning with propositions

³I've changed the notation for infinitesimals from o to ϵ to match the notation used in this chapter.

⁴The difference is merely that non-standard analysis treats the derivative of $f(x)$ as the standard part of $\frac{f(x+\epsilon)-f(x)}{\epsilon}$, so that step 2 is carried out on $\text{st}\left(\frac{(x+\epsilon)^2-x^2}{\epsilon}\right)$, and step 3 is the means by which one calculates the standard part of the expression arrived at in step 2.

about infinitesimals (understood as a type of hyperreal). This would seem to hold true even when one carries out the algorithm with little regard for the propositions expressed by the formulae manipulated in the process. Consider the elementary algebra one does in carrying out the algorithm to compute the derivative of x^2 . Outside of the context of the algorithm, this is also likely to be done more or less mechanically, with little regard for the propositions expressed, but it is again natural to think of performing these algebraic manipulations as reasoning with propositions about numbers.

Likewise, this understanding of the early calculus is not undermined by the fact that mathematicians generally did not take statements about naïve infinitesimals to be serious candidates for truth. Consider the period following the introduction of complex numbers. Statements including terms for complex numbers were not typically taken to be serious candidates for truth and were used to facilitate other calculations (for instance, using an algebraic solution to find the real roots of a polynomial). But again, it is natural to think of making these calculations as reasoning in terms of complex numbers and that this reasoning provides us with a picture of the mathematical world in which such numbers exist. So, for the purposes of this chapter, I will continue to suppose that the early calculus was inconsistent.

4.2 Mapping Accounts and the Early Calculus

To make sense of this case in terms of the mapping account, we must find an appropriate mathematical structure and a mapping from it to our target system. But if we take the structures involved in the mapping account to be structures of the usual sort (something like ordinary set-theoretic structures), we immediately have a problem, as no such structure satisfies an inconsistent theory. So, assuming that they had any (non-trivial) content, scientific representations appealing to inconsistent mathematics cannot be explained by this default version of the mapping account.

One way forward is to appeal to a non-standard notion of structure according to which

some structures satisfy inconsistent theories. A prominent approach that does so is the partial-structures approach, which liberalizes the notions of structure and morphism in order to represent not just inconsistency in science, but also a much broader range of phenomena, including uncertainty and the development of scientific representations over time. I discuss this approach in §4.2.1.

A similar approach is to appeal to inconsistent structures, a move suggested in passing by Colyvan [2008b, 2009] but not to my knowledge developed any further. As with the partial structures approach, the inconsistent structures approach appeals to a more liberal notion of structure, this time one in which inconsistency is explicitly represented. I discuss this approach in §4.2.2.

Finally, a more conservative approach is to explain representations appealing to the early calculus in terms of consistent structures picked out by later, more rigorous versions of the calculus, like modern calculus or even non-standard analysis. I discuss this approach in §4.2.3.

In the rest of this section, I argue that each of these approaches is unsatisfactory.

4.2.1 Partial structures

The approach

Recall from §2.4 that the notion of a partial structure is a modest generalization of the standard notion of a set-theoretic structure, in which each n -ary relation R is represented as a triple $\langle R_1, R_2, R_3 \rangle$, where R_1 is the set of n -tuples of which R holds, R_2 the set of n -tuples of which R does not hold, and R_3 the set of n -tuples for which R is undefined. (And so $R_1 \cup R_2 \cup R_3 = D^n$.) Classical structures are special cases of partial structures, in which R_3 is empty for every relation in the structure. A sentence ϕ is then said to be *partially true* in a partial structure \mathcal{S} if and only if there is a total consistent structure \mathcal{S}' extending \mathcal{S} such that ϕ is true in \mathcal{S}' .

While an inconsistent theory cannot be true *simpliciter* in any total structure (in the sense that each sentence to which the theory is committed is true in the structure), an inconsistent theory can be partially true in a partial structure (in the sense that each sentence in the theory

is partially true in the structure). Suppose we have a theory containing both Ra and $\neg Ra$. Each of these is partially true in a partial structure that puts the interpretation of a in the R_3 block of the relation interpreting R , since we can extend this structure by putting $\mathcal{I}(a)$ in either the R_1 or R_2 block of the relevant relation, thus making true Ra or $\neg Ra$, respectively. So we can understand inconsistent theories as picking out a class of partial structures—namely, those that make each sentence in the theory partially true.

To apply this approach to our case study, we'll first need to determine which class of partial structures is picked out by the early calculus. Given the account of the inconsistency of the calculus sketched in §4.1, we should count each step in calculating a derivative or integral, as well as the information needed to license each inference (and so both $\epsilon \neq 0$ and $\epsilon = 0$), as a commitment of the early calculus, which any appropriate partial structure must make partially true.⁵

In a sense, at this point, we've already arrived at the partial structures representation of the early calculus as it has been reconstructed here. As Vickers [2009, p. 244] notes, the partial structures approach represents any theory with the class of partial structures that make it partially true, with each structure representing different doubts we might have about the theory. But the resulting class of structures must be pared down if we are to understand how such structures could be used as part of the mapping account. Not just any of these structures will allow us to posit (partial) mappings that give the resulting mathematical representation the right accuracy conditions. For instance, these will include entirely uninformative structures that leave empty the R_1 and R_2 blocks of each relation. What must be true of the structures that can be used to apply the calculus?

To simplify things, consider structures only for the differential calculus. The story can naturally be extended to accommodate integrals. Appropriate structures will consist of a domain containing the denotations of expressions referring to reals and denotations for expressions

⁵If we wish to present these commitments more formally, I take them to be captured well enough by the consequences of Brown and Priest's [2004] system LN , understood as the union of the consequences of each chunk in their model.

including infinitesimals; relations corresponding to at least the usual arithmetical functions and relations, including equality; a unary, second-order function d (taking each function to its derivative); and any particular functions and relations needed for the particular application (for instance, a function to represent the trajectory of some object) together with functions to serve as the denotations of their (first- and higher-order) derivatives.⁶

A natural structure to choose is one in which the arithmetical functions and relations are defined as usual on the reals, and the function d takes each function to its derivative. Mathematicians were confident about results containing only terms referring to reals, including the derivatives they calculated, and such results were those from which one might be allowed to make inferences about a physical target system on the basis of one's representation. On the other hand, arguments including terms for infinitesimals can be left in the R_3 block of all relations in the structure,⁷ since, in practice, scientists did not make inferences about their target systems directly from intermediate steps in the calculation of a derivative, not taking mathematical claims about infinitesimals to be physically significant. Consider the calculation of the derivative of $f(x) = x^2$ in §4.1. An ordinary use of this derivation in representing a target system would interpret f as describing the displacement of some object as a function of time. A physicist would then interpret f' as describing its velocity as a function of time, but would only treat the final line of the derivation, $f'(x) = 2x$ as telling us something about the object's velocity. Physicists would commit a serious *faux pas* if instead they inferred that the object's velocity at time $x = 1$ were, say, infinitesimally greater than two (or, worse, the ratio of two infinitesimally small physical quantities). To represent this refusal to make inferences about the target system on the basis of mathematical statements including terms for infinitesimals, we must require the relevant mapping from the structure of the target system to the mathematical structure to have only reals in its image.

However, while the type of structure considered above is certainly a natural one to focus

⁶If we choose Brown and Priest's [2004] *LN* as our reconstruction of the early calculus, we will need to do a bit more, since they introduce a function symbol to provide terms for infinitesimals, as well as a λ -operator.

⁷For the sake of simplicity, I assume here that we represent functions as their graph relations.

on, we can produce representations with exactly the same accuracy conditions (and so that justify the same physical inferences) while significantly varying which elements of the domain go into the R_3 blocks of the relations. For instance, certain arguments containing the denotations of terms for infinitesimals may be put in the R_1 or R_2 blocks of some relations. Without changing the accuracy conditions of the representation, one could use a structure with, say, $\langle \epsilon, 1 \rangle$ in the R_2 block of the equality relation. Likewise, if we do not take our representation to license inferences on the basis of, say, the second derivative of $f(x) = x^2$ —perhaps we use this to represent an object’s velocity over time, but our representation only represents its velocity and acceleration—we can put $\langle 2x, g \rangle$ for any function g we choose in the R_3 block of the derivative function d without changing the accuracy conditions of our representation.

Problems

Such structures and mappings give the right accuracy conditions for typical representations appealing to the early calculus. In doing so, they succeed in the first task for an account of mathematical scientific representation discussed in section 2.2, namely explaining how inconsistent mathematical representations can latch on to their physical target systems. The central question is then how well they do at the second task: as a meta-level representational device to bring out salient philosophical issues from episodes in which the early calculus is applied.

The problem with the partial structures approach in relation to this second task is that it necessarily represents scientists’ inferential practices indirectly in terms of a class of structures. This indirect approach, I will argue shortly, results in a less perspicuous representation of the philosophically salient aspect of the relevant scientific practices than the more direct approach made possible by RIC.

In some cases, of course, such an indirect approach will be called for. As philosophers of science, we want to do more than simply describe case studies; we would like to be able to say something about the more general issues such practices raise. In many cases, an indirect ap-

proach to representing such practices brings out common features that a more direct approach would gloss over. In this case, however, there seems to be no such benefit; instead, the partial structures approach obscures philosophically critical aspects of the practice of applying the early calculus that RIC can be used to represent directly.

Ultimately, making the additional structural claim adds nothing to the original explanation in terms of the practice of applying the early calculus. The restrictions on appropriate structures and mappings above really amount to the claim that physicists took derivatives, but not infinitesimals, to require a physical interpretation, which in turn amounts to the claim that they allowed themselves to make inferences about the world from claims about derivatives in which terms for infinitesimals don't appear, but not from claims about infinitesimals. Since we are answerable only to these permissible inferences in choosing the right structure, and since the right structure is seriously underdetermined by both the practice of mathematicians who used the early calculus and that of the scientists who applied it, we have good reason to think the inferences are doing the explanatory heavy lifting, not any such structure. Note that the problem is not that the partial structures approach does not tell us which structures are appropriate to use to represent applications of the early calculus. [Vickers \[2009, p. 245\]](#) persuasively argues that this in itself is no objection to the partial structures approach. The partial structures approach treats a theory as determining a class of partial structures, each providing a different representation, but leaves open the question of which structure in this class best represents scientific practice. Rather, my complaint is that, once we answer this question by looking directly to scientific practice, it is unclear what value is added by representing this practice in terms of partial structures.

If what I say in §4.3 is right, we can provide a better account of such applications in terms of these inferences themselves without appealing to a mathematical structure picked out by the early calculus. RIC allows for a more direct representation of both the algorithms practitioners used and the ways in which they used them. Reconstructing such applications in terms of structure at best is superfluous and at worst obscures from view the very inference

strategies that allowed the early calculus to be successfully applied.

Beyond this, there are reasons to think that the partial structures approach in this case fails even at the first task, explaining how inconsistent mathematical representations latch onto their target systems. This is because the partial structures approach cannot model the full range of ways in which we might use an inconsistent mathematical theory to produce a physical representation.

One possibility is the one considered above, in which we do not treat infinitesimals as physically significant and so posit a mapping that does not take any physical quantity to an infinitesimal. The result is a consistent physical representation that will have the same accuracy conditions as one produced using the modern calculus.

Another possibility is that part of the scientific community at the time of the introduction of the calculus had an inconsistent conception of change, which mirrored the inconsistent conception of infinitesimals in the early calculus. According to such a conception, there really are infinitesimal physical quantities (both temporal and spatial), and there really is, say, instantaneous velocity, and things like the latter are to be explained in terms of the former. In this case, we would expect the relevant mapping to take infinitesimal physical quantities to infinitesimals in the mathematical structure, and so we would expect intermediate steps in the calculation of a derivative or integral to be physically significant.⁸ On the partial structures approach, such a representation would have to represent (a substructure of) the physical structure as appropriately morphic to a properly partial substructure of the relevant mathematical structure. But assuming the structure of the world is not properly partial, at least in the sense of partial structures at play here, there can be no such morphism, and so the representation is necessarily inaccurate. This seems right if we have a very coarse-grained notion of the contents of such representations. But if we move to more fine-grained accuracy conditions, we seem to get the wrong ones. Such representations represent the world as inconsistent, so

⁸Colyvan [2009, p. 167] considers this possibility. An alternative, also considered by Colyvan, is that we produce an inconsistent representation of the same kind as an idealization in order to draw out a consistent set of useful consequences. In this case, the representation and its accuracy conditions presumably stay the same.

that some physical quantities are both zero and non-zero, for instance. We take the representation to be necessarily false because we take such physical quantities to be impossible. On the other hand, the partial structures representation would seem to represent the world as incomplete (or whatever it takes to instantiate a properly partial structure, assuming that the notion of instantiating a partial structure makes sense at all), so that there are physical quantities that are in a sense neither zero nor non-zero. This too we might take to be necessarily false, but this time because we take it to be impossible for the world to instantiate a properly partial structure. The states of affairs represented—the world's having an inconsistent structure and its having a partial structure—are distinct, and the partial structures view can only represent the latter.

There is a third possibility. Suppose that instead of taking instantaneous velocity, acceleration, and so on to be explained in terms of inconsistent infinitesimals, we take them to be explained in terms of indeterminate infinitesimals. That is, we suppose that there is a physical correlate to the infinitesimals described by the early calculus, but, knowing that the latter are inconsistent, we don't take a stance on many of the properties of their physical correlates, particularly those properties that correspond to the inconsistent properties of naïve infinitesimals. This is distinct both from taking infinitesimals to be a mere artifact of the mathematics and from taking them to lead to an inconsistent representation of the world. Intuitively, such a representation should be accurate, for example, if there are infinitesimal physical quantities structurally similar to the infinitesimals described by recent reconstructions of infinitesimal calculus like those given by smooth infinitesimal analysis [Bell, 2008, Moerdijk & Reyes, 1991] or non-standard analysis [Robinson, 1966]. (That we should want to be able to account for such representations is supported by the fact that part of the interest of these reconstructions is that they have some claim to explaining the conceptions of space, continuity, change, and related notions held by scientists and mathematicians who used the early calculus that the modern calculus does not have.) On the other hand, if we had a genuinely inconsistent conception of change as described above, our representations would not be made true by

such quantities. Such a representation says *less* than a representation of inconsistent physical infinitesimals but *more* than the kind of representation we discussed earlier, which treats infinitesimals as a mere artifact of the mathematics and doesn't provide them with *any* physical interpretation.

A mapping account appealing to partial structures would have to treat indeterminate representations of this kind in the same way as it treats the earlier inconsistent representation. We are still left representing the structure of the world as partial if our mapping takes infinitesimal physical quantities to infinitesimals in the mathematical structure (and the latter are in the R_3 block of any interpreted relation). But it is hard to see how a structure that allows for a representation with a reasonably strong physical content could both keep infinitesimals out of the R_3 blocks of its interpreted relations and make each sentence in our reconstructed theory of the calculus partially true. The result is that this version of the mapping account would also assign the wrong accuracy conditions to this sort of representation. The mapping account depicts such representations as accurate only if the structure of the target system is partial—something such representations do not in fact require.

4.2.2 Inconsistent structures

Alternatively, we might try to explain applications of inconsistent mathematics in terms of inconsistent structures.⁹ We can understand an inconsistent structure as a set-theoretic structure in which each relation R is represented as a pair $\langle R^+, R^- \rangle$, where R^+ is the extension of the relation (containing the n -tuples of which the relation holds) and R^- is the antiextension of the relation (containing the n -tuples of which the relation does not hold). We require that $R^+ \cup R^- = D^n$ for each n -ary R , but don't require that $R^+ \cap R^- = \emptyset$. Consistent structures come out as a special case where $R^+ \cap R^- = \emptyset$ for each R . An atomic Ra is true in an inconsistent structure if and only if the denotation of a is in the extension of the relation assigned to R and false in an inconsistent structure if and only if the denotation of a is in the antiextension

⁹Something like this is required to make sense of some of the suggestions in [Colyvan, 2008b, 2009].

of the relation assigned to R .¹⁰ A representation appealing to an inconsistent mathematical theory represents the world as structurally similar to an inconsistent structure in which each sentence of the reconstructed mathematical theory is (at least) true. That is, such a representation posits the existence of a structure-preserving mapping between (a substructure of) the structure of the physical target system and (a substructure of) the inconsistent mathematical structure.

For the purposes of understanding applications of inconsistent mathematics, this approach has some advantages over the partial structures approach. In particular, unlike the partial structures approach, it seems to give the right fine-grained accuracy conditions for inconsistent representations making use of the early calculus. For instance, if we map an infinitesimal physical quantity to a mathematical infinitesimal ϵ in the structure, we seem to represent that physical quantity as both zero and non-zero, since $\langle \epsilon, 0 \rangle$ will have to be in both the extension and the antiextension of the equality relation.

However, all the issues that came up for the partial structures approach concerning which partial structure and mapping are relevant to a given application come up again here in determining which inconsistent structure¹¹ and mapping are relevant to a given application. Given the story in the previous subsection, we should expect the relevant mapping to be between the same bits of the empirical setup and a consistent substructure of the inconsistent structure picked out by the early calculus. Indeed, we should expect this consistent substructure to be isomorphic to the total substructures of the partial structures we considered in the previous section generated by the image of the relevant mapping. So we get what amounts to the same account of cases where the early calculus is applied to produce consistent representations of the world. And so we end up with the same problem; the substantive, interesting

¹⁰The structures and truth-definition described above are in fact those of the logic LP , first developed by [Asenjo \[1966\]](#) and popularized by [Priest \[1979\]](#). I take these structures (and corresponding truth-definition) to be the most straightforward way to make sense of various claims about ‘inconsistent structures’, but the remarks I make below should also apply to other ways of making sense of the notion.

¹¹Just as with partial structures, many inconsistent structures satisfy a given theory, differing from one another in which items in the domain are in both the extension and antiextension of each relation. For instance, every sentence is at least true in a structure where the extension and antiextension of each n -ary relation are both simply D^n .

explanation of what was going on isn't given in terms of structure, and adding the structural claim to this explanation yields nothing new. And, again, the structural claim seems to obscure the structure-independent explanation that does the real work. Moreover, like the partial structures approach, the inconsistent structure approach does not have the resources to account for cases where the inconsistency of the mathematics leads us to posit indeterminate physical correlates of the inconsistent mathematical entities. For our purposes, the two approaches differ very little.

4.2.3 Related total consistent structures

Finally, assuming that we're not interested in how the calculus might have been used to produce inconsistent or indeterminate representations of the world, we might try to explain representations appealing to the early calculus in terms of some related total consistent structure. There are two sorts of explanation we might want to provide in these terms.

The first is a retrospective explanation of the success of representations using the early calculus, given our present knowledge. It is more or less trivial to do this in terms of total consistent structures. Given that early practitioners' procedures yielded the same results as the modern calculus, their results are satisfied by structures for the modern calculus, which can in turn be mapped to the relevant target systems.

The second sort of explanation concerns the epistemic status of representations based on the early calculus from the point of view of practitioners at the time. Why was it reasonable for scientists to make physical inferences using the early calculus despite its inconsistency? This too is something we should expect a mapping account to be able to support as a meta-level representational device for representing philosophically salient features of scientific practice. Perhaps we can think of scientists at the time as using the inconsistent calculus to gesture at some consistent structure. Some early practitioners, including Newton, might seem to have gestured at something that, in retrospect, looks like an ϵ, δ limit. For example, early in the *Principia*, Newton writes:

Those ultimate ratios with which quantities vanish are not actually ratios of ultimate quantities, but limits which the ratios of quantities decreasing without limit are continually approaching, and which they can approach so closely that their difference is less than any given quantity, but which they can never exceed and can never reach before the quantities are decreased indefinitely. [...] Therefore, whenever, to make things easier to comprehend, I speak in what follows of quantities as minimally small or vanishing or ultimate, take care not to understand quantities that are determinate in magnitude, but always think of quantities that are to be decreased without limit. [1999, pp. 88f]

To modern eyes, it might appear that Newton is telling the reader that infinitesimals are just a heuristic device used to produce the same results as the consistent notion of an ϵ, δ limit.

But, while this supports the view that mathematicians were dissatisfied with infinitesimals, it does not support the idea that such representations were possible because there was such a consistent foundation in the vicinity. The fact that it took over 150 years to produce a consistent foundation for the calculus in terms of ϵ, δ limits strongly suggests that early practitioners of the calculus didn't have a suitably determinate conception of limit in mind. Moreover, as noted above, recent reconstructions of the early calculus in terms of infinitesimals—particularly smooth infinitesimal analysis and non-standard analysis—are of interest largely because they have some claim to being better reconstructions of the early calculus than the calculus built on ϵ, δ limits. That all three have a claim to being rational reconstructions of the early calculus again strongly suggests that mathematicians using the early calculus were not gesturing at a suitably determinate structure.

Now, for the purposes of the second sort of explanation, mapping accounts should require that, in using the early calculus to represent the world, scientists appealed, at least implicitly, to a consistent structure of the sort picked out by later versions of the calculus based on the modern definition of a limit. Plausible versions of the mapping account require the scientist constructing the representation to actually define the relevant mapping by pairing up the

4.3 A Robustly Inferential Account of the Early Calculus in Applications

relevant objects and relations; different mappings, even between the same two structures, can produce representations with different accuracy conditions due to differences in which objects and relations they pair up, and so merely positing the existence of a mapping is not enough. [Nguyen & Frigg \[2021\]](#) observe that this means that the scientist must also be able to single out a particular structure in the target system. But if this is so, then the scientist must also be able to single out the relevant mathematical structure in order to define the mapping in question. So, if scientists who used the early calculus weren't at least implicitly appealing to a consistent structure picked out by the modern calculus, then mapping accounts cannot explain their calculus-based representations in terms of such a structure.

4.3 A Robustly Inferential Account of the Early Calculus in Applications

I have just argued that mapping accounts have significant shortcomings in their treatment of applications of the early calculus and inconsistent mathematics more generally. In their role as meta-level representational devices, they are forced to represent the most epistemically significant features of the relevant practices—namely, algorithmic inference strategies for safely reasoning with inconsistent information—at best indirectly. Further, such accounts fail to represent the full range of ways in which we might use an inconsistent mathematical theory to produce a physical representation. I now show how we can represent this case in terms of RIC in a way that avoids both major problems. RIC-based approaches can directly represent scientists' algorithmic inference procedures and have more flexibility in how they represent the possible physical interpretations of inconsistent mathematical representations.

Recall that, according to RIC, mathematical scientific representations have three ingredients:

(RIC1) A partial physical interpretation of the language in which the relevant mathematics is expressed,

4.3 A Robustly Inferential Account of the Early Calculus in Applications

(RIC2) An initial description of the target system in this physically interpreted mathematical language, and

(RIC3) A set of mathematical inference patterns licensed by the relevant mathematics.

The representation's informational content is then given by the closure of the statements in RIC2 under the inference patterns in RIC3, under the interpretation RIC1.

Much of the story about applications of the early calculus that RIC helps us to tell parallels the account of which structure and morphism are at play in such applications on the partial structures and inconsistent structures views. Crucial there was the observation that those who applied the early calculus did not allow themselves to make inferences about their target systems on the basis of mathematical claims including terms for infinitesimals, particularly intermediate steps in calculating a derivative or integral. We can capture this in two different ways, both of which do better than the strategies examined so far.

First, we might take our set of privileged inference patterns (RIC3) to contain all inferences that mathematicians allowed themselves to make in using the early calculus. These inferences are already quite restricted, as noted before, as infinitesimals may only be used in calculating derivatives and integrals, and the propositions that $\epsilon = 0$ and $\epsilon \neq 0$ may only be used in particular parts of these calculations. To avoid inconsistency in our physical representation, we provide reals, but not infinitesimals, with a physical interpretation (RIC1). For instance, we might specify that real values of t represent time (in some unit), that real values of x represent position on some axis, that real values of $f_x(t)$ (for $t \in \mathbb{R}$) represent the position (on the same spatial axis) of a given object at the time represented by t , that for any physically interpreted function $f(x)$, real values of $f'(x)$ represent the rate of change of what is represented by f with respect to what is represented by x , and so on.

Alternatively, we might restrict the set of privileged inference patterns (RIC3). For instance, we might leave out all inferences to or from claims including terms for infinitesimals. This would leave out all inferences made within the calculation of a derivative or integral, but

4.3 A Robustly Inferential Account of the Early Calculus in Applications

would leave in inferences from what a particular function is to what its derivative is. So, for example, the inference from $f(x) = x^2$ to $f'(x) = 2x$ would be an instance of an included inference pattern, while the inference from $f(x) = x^2$ to $f'(x) = \frac{(x+\epsilon)^2 - x^2}{\epsilon}$ would not.¹² As a result, the algorithms could not yield any sentences containing terms for infinitesimals, and so we could drop the restriction to real values in the partial physical interpretation of the mathematical vocabulary (RIC1).

The constraints placed on a physical system by this kind of representation will be those that the partial structures and inconsistent structures accounts considered in §4.2 describe. But RIC can represent them in a way that more directly captures the information about the practice of applying the calculus that is ultimately used to determine which partial or inconsistent structures and mappings are appropriate to represent a given application. Which structures and mappings are appropriate are largely determined by which inferences scientists allowed themselves to make on the basis of the relevant mathematics. The mathematical part of these inferences is directly captured by the privileged set of mathematical inference patterns, while the physical part is captured by the partial physical interpretation of the mathematical vocabulary. For this reason, RIC provides a more perspicuous, explanatory account of applications of inconsistent mathematics than the partial and inconsistent structures approaches.

This framework can also explain attempts to provide infinitesimals with physical interpretations. In an inconsistent physical representation of the kind at times suggested by [Colyvan \[2009\]](#), the collection of privileged inference patterns would again contain all inferences that mathematicians allowed themselves to make in using the early calculus. But the partial physical interpretation would now assign a physical content to sentences containing terms for infinitesimals—treating terms for infinitesimals as denoting infinitesimally small physical quantities. In effect, this would involve providing an inconsistent physical justification for

¹²In this case, RIC3 would include (something like) the algorithms of the calculus, as described by [Vickers \[2013\]](#), as well as those corresponding to the algebraic inferences (restricted to the reals) one might make outside of the calculation of a derivative or integral.

4.3 A Robustly Inferential Account of the Early Calculus in Applications

algorithms taking us to or from sentences containing terms for infinitesimals and, in particular, an inconsistent physical explanation of why procedures for calculating derivatives and integrals work.

By taking such a representation and restricting the collection of privileged inference patterns and physical interpretation in the right way, we could also produce a physical interpretation that posits indeterminate physical infinitesimals of the kind discussed in §4.2.1. For instance, if we modify the physical interpretation so that statements of the form ' $a = b$ ' are not interpreted when a or b is a term for an infinitesimal, but otherwise interpreted as expressing equality up to an infinitesimal difference, we get a representation that could be accurate if there were physical quantities that behaved like the hyperreals of non-standard analysis. After some filling in, the resulting physical justification for the derivative and integral algorithms might then turn out to look something like the mathematical justifications given for the corresponding inferences in non-standard analysis. So, unlike mapping accounts, RIC can accommodate the difference between not positing infinitesimal physical quantities, positing inconsistent infinitesimal physical quantities, and positing physical quantities corresponding to infinitesimals but not saying enough about them to make the resulting theory inconsistent.

Now, at this point, one might worry that this account of representations appealing to inconsistent mathematics explains how such representations are possible at the cost of failing to explain what makes them undesirable.¹³ When the use of inconsistent mathematics results in an inconsistent physical representation, the shortcomings are clear: such a representation must misrepresent the world (assuming, *pace* dialetheists, that the world does not have contradictory properties). But what about cases in which consistent physical representations appeal to inconsistent mathematics? In such cases, the consistency of the physical representation depends on various kinds of inferential restriction. First, not all classically valid inferences can be permissible in reasoning within the inconsistent mathematical theory; otherwise, the classical law of *ex contradictione quodlibet* would trivialize the theory.

¹³I thank an anonymous referee for the *British Journal for the Philosophy of Science* for raising this objection.

4.3 A Robustly Inferential Account of the Early Calculus in Applications

Second, either the physical interpretation of the mathematical vocabulary or the collection of mathematical inference patterns (or both) must be restricted so that no inconsistent set of claims in the mathematical theory, when physically interpreted, is among the commitments of the physical representation.

Representations requiring these restrictions have several shortcomings. In the first place, producing such a representation is no easy matter. It must be clear exactly what these restrictions are for such a representation to have a determinate content. And even when such restrictions are made, it may not be clear whether they are sufficient to produce a consistent physical representation. More importantly, these restrictions rule out some ways of better understanding why calculations used in the representation work. For instance, consider the use of the derivatives of the early calculus to represent instantaneous velocity. If none of the intermediate steps in calculating a derivative are physically interpreted, the derivative is treated as a black box in the physical representation, making it mysterious why this procedure yields the function representing an object's velocity when applied to the function representing its displacement. The most straightforward way of interpreting these intermediate steps yields the inconsistent representation considered above, while ways of interpreting them that yield consistent representations give us only incomplete explanations of why the procedure works, as in the case of the representation positing indeterminate infinitesimals considered above. The modern calculus, in contrast, allows us to provide these intermediate steps with a physical interpretation, treating instantaneous velocity as the limit of an object's average velocity over a finite interval of time as that interval approaches zero—thereby also arguably shedding more light on the very concept of instantaneous velocity. That said, despite these shortcomings, such a representation can be very useful both when no suitable representation appealing to consistent mathematics is available and when it is simply more computationally convenient to continue to use the representation appealing to inconsistent mathematics.¹⁴

¹⁴For example, even once it was discovered by [Schwartz \[1945\]](#) that mathematically rigorous distributions could do the work of mathematically ill-defined delta functions in quantum mechanics, physicists largely continued to appeal to delta functions.

4.3 A Robustly Inferential Account of the Early Calculus in Applications

It is also worth pausing here to distinguish RIC from Vickers's view that the early calculus is algorithmic. It might seem that I am now advocating a very similar view, despite having rejected Vickers's view in §4.1.¹⁵ A first difference is that the two views are simply about two different things. RIC concerns how mathematics can be used to produce representations of the physical world generally, while Vickers [2013, ch. 6] does not engage explicitly with physical applications of the calculus. As a result, our claims about algorithms amount to very different things. When Vickers says that the early calculus was algorithmic, this is a statement about what mathematicians were doing when they proved results within the calculus. When I say that we can think of the inference patterns in RIC3 as algorithms, this does not entail that we should understand the mathematical inferences corresponding to these algorithms as themselves algorithmic when made in the context of carrying out a mathematical proof.

More importantly, I'm not even making the analogous claim about scientific reasoning, that scientists reason algorithmically when they apply mathematics. The role of the set of inference patterns (RIC3) in RIC is just to encode inferences about the target system that must preserve truth according to the representation: if A is a physically interpreted claim in the enhanced language of the mathematical theory, and B is a physically interpreted claim in this language at which we can arrive by applying some sequence of the algorithms to A (that is, by carrying out an algorithm with A as input, carrying out a second algorithm with the output of the first as input, and so on), then the representation is committed to B if it is committed to A . It is for this reason that I characterize the view as inferential rather than algorithmic.

On the other hand, the thought might be that I end up with a similar explanation of how the early calculus was applicable. For instance, a proponent of a standard mapping account (without partial or inconsistent structures) might claim that Vickers shows that the early calculus is applicable because it is algorithmic and thus consistent. The algorithms in RIC would similarly seem to allow us to produce a consistent representation using the early

¹⁵I would like to thank two anonymous referees for the *British Journal for the Philosophy of Science* for encouraging me to consider this point.

calculus. However, RIC can also capture uses of the early calculus to produce inconsistent representations. And, even when using it to describe cases where consistent representations are produced, the inference patterns in RIC3 need not correspond to the algorithms of the calculus as Vickers describes them.

4.4 Beyond Inconsistent Mathematics

For all I have said so far, RIC might be thought to be a mere patch, supplementing mapping accounts with a different account of applications of inconsistent mathematics. By way of conclusion, I will now briefly argue that this is not the case. The shortcomings of mapping accounts with respect to applications of inconsistent mathematics give us reason to favor RIC sketched in the previous section as a general account of mathematical scientific representations.

Proponents of mapping accounts might be tempted to treat applications of inconsistent mathematics as a marginal, pathological case, so that failure to explain such applications is not a strike against mapping accounts. But this is the wrong move for three reasons.

First, from its introduction, the early, inconsistent calculus was widely applicable and hugely successful. It was a central part of much of the most important scientific work done between its introduction and the introduction of the modern calculus. Even if we only consider the early calculus, we can hardly treat applications of inconsistent mathematics as marginal; they are among the cases we should expect a good theory of applications to explain.

Second, the case of the early calculus bears important similarities to more recent uses of mathematics in physics, in which it is not clear that the role of the mathematics is to pick out a mathematical structure due to its failure to live up to the standards of rigor of pure mathematics. As [Urquhart \[2008a, p. 410\]](#) puts it, ‘the methods that [physicists] use are frequently so far from normal mathematical practice that it is sometimes not clear that the objects [appealed to in mathematical representations] themselves are even mathematically well-defined’. Some-

times, this involves using heuristic devices, like Heaviside’s operational calculus, to solve difficult mathematical problems. These techniques allowed problems involving linear differential equations to be reduced to algebraic problems that were much easier to solve. Though it would take decades for these techniques to be given a mathematically rigorous justification, they bore a great deal of fruit in physics—particularly the study of electromagnetism.¹⁶ But sometimes such methods have been crucial to formulating our representations themselves. For instance, the Feynman path integral has yet to be put on a fully rigorous mathematical footing,¹⁷ but is central to the (extremely successful) path integral formulation of quantum mechanics among other things. In other cases, the work mostly involves formally manipulating symbols, with the help of computer simulations—for example, computer models frequently use techniques like finite differencing, which allow difficult sets of partial differential equations to be solved by brute force, together with additional *ad hoc* features to make the already simplified problem computationally tractable.¹⁸ Such techniques led Steiner [1992, p. 100] to write, “it is doubtful whether we can attribute to *today’s* physicists even a consistent mathematics.”

If it were crucial to pick out a mathematical structure and posit a mapping between it and the physical target system, then we would expect physicists to use more mathematically rigorous techniques. In appealing to mathematical objects that may not be well-defined, one would run the risk of failing to produce a representation at all. And yet mathematical objects that really were not well-defined, like the Dirac delta function (the function whose value is zero except when its argument is zero and whose integral over the real line is one), have produced representations of scientific importance, even if rigorous techniques to do the same work were subsequently developed. In particular, the delta functions have been used in representations in a number of different areas, including the study of circuits (where they played the role of the unit impulse) and quantum mechanics. Likewise, in using unjustified heurist-

¹⁶For a good history of the operational calculus and attempts to provide it with a rigorous foundation, see [Lützen, 1979].

¹⁷Though progress has been made. See [Johnson & Lapidus, 2000].

¹⁸For an accessible discussion of these and other techniques in computer modeling, see [Winsberg, 2010].

ics, one would run the risk of misrepresenting the structure essential to applications. And, in merely manipulating symbols, perhaps with the aid of computer simulations, one would not seem to have a particular mathematical structure in mind at all. It is less than clear that this activity is best reconstructed as reasoning about a mathematical structure, even when there is an appropriate structure in the vicinity.

But if mathematical representation doesn't essentially depend on relations between mathematical and physical structures, we have an explanation of this. Physicists do not need to hold themselves to pure mathematicians' standards of mathematical rigor because they ask less of the mathematics than pure mathematicians do. What physicists need is a means of representing and reasoning about physical systems. For that purpose, there is no reason to exclude computational techniques based on heuristics, inferences about mathematical objects that are not well-defined, or even formal manipulation of uninterpreted symbols. Since the aim is not to pick out and explore the properties of a particular mathematical structure or class of such structures, we should not expect that the most expedient techniques are those appropriate for pure mathematics. I consider such cases in much greater detail in the next chapter.

Finally, if proponents of mapping accounts appeal to a different account to explain applications of inconsistent mathematics, then they are forced to accept very different explanations of the applicability of the calculus pre- and post-Weierstrass, without any corresponding difference in the practice of applying the calculus. Given the continuity in the practice of applying the calculus before and after it was given a consistent foundation, we should expect continuity in the explanations of these practices. Moreover, if we take the contents of such representations to be at all fine-grained, the account of mathematical representations couched in terms of the early calculus could not attribute the same content to these representations as the default mapping account attributes to representations using the later, consistent calculus, since the latter must invoke the structure picked out by the later, consistent calculus, while the former could not. Even if we accept different explanations of the early, inconsistent

calculus and the later, consistent calculus, it seems tremendously implausible that the *content* of representations appealing to calculus would change in this way without a corresponding change in practice on the part of physicists.

All of this gives us good reason to think that RIC provides not only a good explanation of applications of inconsistent mathematics, but also a compelling, general explanation of mathematical representations of the physical world.

CHAPTER 5

Applying Unrigorous Mathematics: Heaviside's
Operational Calculus and Path Integrals in
Quantum Physics

Applications of inconsistent mathematics are an extreme case of a more general and much more common phenomenon: the use in physics of mathematics that fails in one way or another to meet the standards of rigor that apply to work in pure mathematics. From the early calculus in Newtonian physics to ill-defined path integrals in quantum field theory, physicists have leaned heavily on mathematical tools that fall well short of the standards of rigor of present-day pure mathematics.¹ These tools have made possible a number of important physical results despite the mathematics' not clearly sufficing on its own to pick out well-defined mathematical structures. The success of these techniques and their relation to more rigorous techniques are among those features of scientific practice we might want a philosophical account of applications of mathematics to shed light on in its role as a meta-level representational device for philosophers of science. But little attention has been paid to these questions in existing work.

The aim of this chapter is to show that we have reason to favor RIC over mapping accounts on the grounds that RIC is a better tool for addressing these and other questions arising from applications of unrigorous mathematics. Central to applying unrigorous mathematics is the use of inference strategies that restrict the use of incoherent, underdeveloped, or otherwise problematic concepts so that undesirable results cannot be derived. Because such techniques typically involve local inferential restrictions that do not naturally correspond to neat divisions of a mathematical structure, mapping accounts must represent them at best in a highly indirect way, which makes them less useful for reasoning about such cases. In contrast, RIC can represent such strategies much more directly and as a result opens up more possibilities for representing and reasoning about them.

In §5.1, I present a general case for RIC on the grounds that it is a more useful device for representing applications of unrigorous mathematics. In particular, RIC is more useful

¹In what follows, I will refer to such cases as “applications of unrigorous mathematics” for the sake of simplicity. Importantly, in using this phrase I do not wish to imply that the *physics* is unrigorous, but only that the mathematics itself would count as unrigorous by the standards of present-day pure mathematics. I also do not wish to imply that mathematical rigor is a simple binary. Not only is mathematical rigor a matter of degree, but the threshold and criteria for labeling a given piece of mathematics as “rigorous” depend largely on social and historical context.

for reasoning about the methodologies scientists' use to mitigate the risks associated with applying unrigorous mathematics. In the rest of the chapter, I consider in detail two case studies of applications of unrigorous mathematics: Heaviside's operational calculus (§5.2) and path integrals in quantum physics (§5.3). These cases both illustrate the central points made in the more general argument for RIC, as well as how RIC can be used to make sense of a wide variety of methodologies for working with unrigorous mathematics in physics. I conclude (§5.4) by considering two objections to my arguments: first, that these cases show at best that mapping accounts should be supplemented with an account of mathematical formalisms, inferential power, and related concepts, and, second, that these are not applications of genuine mathematics at all. I argue that each of these considerations actually highlights an advantage of RIC.

5.1 Unrigorous mathematics in general

5.1.1 Inferentially permissive and inferentially restrictive methodologies

A useful tool for thinking about applications of unrigorous mathematics can be found in the work of Davey [2003], who argues that at least some arguments that are unpersuasive when construed as mathematical arguments (due to failures of rigor resulting from the use of mathematically ill-defined concepts) are in fact persuasive *qua* physical arguments. This is possible, according to Davey, because physicists employ in these cases an “inferentially restrictive methodology” unlike the “inferentially permissive methodology” of pure mathematics.

An *inferentially permissive* methodology, like that of mathematics as ordinarily practiced, is one in which there is no restriction on what concepts one can appeal to in making an argument (or attempting to prove a theorem or solve a problem) or on the use of classically valid inferences. As mathematics is normally practiced, mathematicians may use any mathematical concept, technique or result from any area of mathematics they like in working on a

particular problem.² This is particularly fruitful in mathematics, as many theorems that have resisted proof using only the resources of the branch of mathematics to which they belong (including, perhaps most famously, Fermat’s Last Theorem) have turned out to be accessible to proof when those domains have been linked to superficially unrelated areas of mathematics, and the links between these domains are often of great mathematical interest in themselves. This certainly also has benefits for mathematical scientific representations, as it allows for the use of mathematical resources outside the mathematics explicitly used in the representation both to help explore the connections between such representations and to better understand the properties and commitments of these representations themselves.³

In contrast, an *inferentially restrictive* methodology is one in which not just any concept or classically valid inference may be used at any point in an argument. Importantly, such restrictions need not be understood in terms of the adoption of a particular subclassical logic (e.g., a paraconsistent logic⁴), but may instead involve a patchwork of more local restrictions. For instance, if what I said about the early calculus is right, the use of terms for infinitesimals was limited to the context of calculating derivatives and integrals, and the use of particular information about infinitesimals was limited to even narrower contexts within these calcu-

²There are, of course, exceptions. If what I say in the previous section is correct, the early calculus is a notable example, as certain information about infinitesimals could only be appealed to at particular points in a given proof. Early set theory, before the introduction of consistent automatizations, might also be such an example if we think its use did not commit mathematicians to the absurdities resulting from the paradoxes of naïve set theory.

But it is also worth noting that not all explorations within mathematics of mathematically unrigorous concepts require adopting such a methodology. For instance, [Mortensen \[1995\]](#) formulates inconsistent versions of several mathematical theories within a paraconsistent setting, but this work is highly rigorous and requires no inferential restrictions when (rightly) understood as reasoning *about*, rather than *within*, such theories. (That is, Mortensen’s reasoning at the meta-level is entirely classical.)

³This is not to say that these advantages are only available to physical representations allowing for an inferentially permissive methodology. Restricting the use of certain concepts and inferences to particular contexts doesn’t rule out establishing connections between different mathematical domains, though it does place limits on how this can be done.

⁴Even when the relevant restrictions are best understood in terms of the adoption of a particular logic, this need not be a paraconsistent logic. For example, smooth infinitesimal analysis requires some classically valid inferences to be unavailable, as otherwise one could derive a contradiction that would trivialize the theory. However, the appropriate background logic for smooth infinitesimal analysis is intuitionistic and so not paraconsistent. The role of the intuitionistic background logic is not to prevent contradictions from trivializing the theory, but to prevent explicit contradictions (i.e., those of the form $p \wedge \neg p$) from arising in the first place. This feature seems to be shared by the “patchwork” strategies considered in the rest of this section. Such strategies don’t necessarily aim to prevent contradictions from trivializing the representation, but rather aim to prevent explicit contradictions from arising at all.

lations. Certain classically valid inferences are unavailable provided that these restrictions are observed, but it is not at all clear that the practice of using the early calculus with these restrictions is best understood in terms of the adoption of a particular subclassical logic.

Such a methodology is required when scientists reason with concepts or information that they believe might lead us to a contradiction (or something else undesirable for their purposes) if certain inferences made available by an inferentially permissive methodology were to be carried out. Such reasoning is unrigorous by the standards of any field with an inferentially permissive methodology, but may be (rightly) persuasive in the context of an inferentially restrictive methodology. These inferential restrictions allow us to quarantine mathematically ill-defined concepts to those contexts in which they behave in the desired way, thereby making it impossible to use them to derive undesirable (or, at the limit, absurd) results.

The use of procedures for calculating integrals and derivatives in the early calculus as discussed in §4.1 is an excellent example of an inferentially permissive methodology. Because naïve infinitesimals were mathematically problematic, their use was restricted to such procedures, and even within these procedures, their use was restricted so that absurdities could not be derived. Consider the procedure for calculating the derivative of a function $f(x)$. One first takes $\frac{f(x+\epsilon)-f(x)}{\epsilon}$ and simplifies it algebraically to eliminate division by ϵ . Here one is expected to reason as if $\epsilon \neq 0$, so that division by ϵ is well-defined. Then, once all division by ϵ has been eliminated, one reasons as if $\epsilon = 0$, in order to eliminate terms in which ϵ appears as a factor. In this case, the use of the problematic concept is restricted to particular contexts (i.e., those in which these procedures are carried out), and within these contexts, not all classically valid inferences are permitted. In particular, in the first step, no classically valid inferences from $\epsilon = 0$ are permitted, and in the second, no classically valid inferences from $\epsilon \neq 0$ are permitted. As a consequence, otherwise inaccessible results not involving infinitesimals can be derived, namely the correct equations for integrals or derivatives of the relevant functions. But no inconsistent results in which infinitesimal terms do not appear can be derived. Provided the infinitesimal terms are not physically interpreted, this means the

resulting physical representation also need not be inconsistent (or otherwise epistemically deficient).

But inferentially restrictive methodologies are far more widespread than applications of inconsistent mathematics. Their use is also called for in any case in which the behavior of a mathematical concept relevant to an application is understood in some but not all contexts, either because the necessary mathematical work has been done or because the concept itself has not been fully pinned down. While such examples abound, in this chapter I focus on two: Heaviside's application of his operational calculus in electrical engineering (§5.2) and the use of path integrals in quantum field theory (§5.3). I argue that the advantages of RIC as a meta-level representational device for philosophers of science generalize from cases of inconsistent mathematics to these cases, and so the advantages of RIC aren't confined to a few fringe cases, but can be seen in connection to a number of key parts of both the history of science and current scientific practice.

5.1.2 Inferentially restrictive methodologies and accounts of applications of mathematics

Before turning to these case studies in the rest of the chapter, in the rest of this section I will present a broad argument for the claim that the use of inferentially restrictive methodologies in applications of unrigorous mathematics favors RIC over mapping accounts. In the rest of the chapter, I will substantiate these points in relation to case studies and discuss some of the more particular philosophical issues that these cases raise.

Recall from §3.1.4 that RIC recovers mapping accounts as special cases, and so adjudicating between RIC and mapping accounts comes down the question of whether RIC's greater generality has philosophical payoffs that outweigh its cost. If there were no such benefits, mapping accounts would be preferable on the grounds that they say more than RIC about applications of mathematics in general. The aim of this chapter is to show that one such benefit is that RIC is a better tool for representing episodes in which unrigorous mathematical

techniques have been applied in the history and current practice of science.

In the previous chapter I argued that, in the case of inconsistent mathematics, RIC does better at both central tasks for an account of applications of mathematics: (1) as an explanation of how inconsistent mathematics can in principle be used to represent a physical target system and (2) as a meta-level representational device to bring out philosophically salient features of episodes from scientific practice. In this chapter, however, I concede for the sake of argument that some versions of the mapping account succeed at the first task, explaining how unrigorous mathematics can in principle be used to represent a physical target system. It is enough for my purposes to show that mapping accounts have shortcomings in connection to the second task, as meta-level devices for representing episodes in which unrigorous mathematics has been applied.

Crucial to using a mapping account as a meta-level device is identifying an appropriate structure or collection of structures to represent the relevant mathematics. This is simple when applications involve well-understood mathematical theories. Even if scientists skip a few steps or rely on unarticulated assumptions, we philosophers of science can straightforwardly represent their practice in terms of a well-defined mathematical structure associated with the theory.

However, when a mathematical theory or technique involved in an application is not a well-understood piece of pure mathematics, more work must be done. In the extreme, the mathematics involved might be inconsistent, as in the case of the early calculus discussed in the previous chapter or the Dirac delta function, both of which found—and in the case of the Dirac delta function, continue to find—widespread use in physics. Since no classical, set-theoretic structure satisfies an inconsistent theory, additional work must be done to understand such cases in terms of the mapping account. Similarly, a piece of mathematics might fail to pick out a well-defined structure because it appeals to inchoate or incoherent concepts. If a mathematical theory appeals to such concepts, there will be a degree of indeterminacy in its global mathematical commitments—even if those concepts are well-behaved and

well-understood in more local contexts—due to there being multiple ways to flesh out these concepts or resolve their incoherence.

One option is to appeal to a more liberal kind of structure built to accommodate inconsistency and indeterminacy. The most promising account to do so is the partial structures approach [Bueno & Colyvan, 2011, Bueno & French, 2012, 2018, da Costa & French, 2003], as discussed in §2.4 and §4.2.1. Recall that a unary relation R in a partial structure partitions the domain into three blocks: R_1 , those items of which R holds; R_2 , those of which R does not hold; and R_3 , those for which R is undefined. (n -ary relations and functions are treated similarly.) Total structures are a special case in which the R_3 block of every relation is empty. A statement ϕ is partially true in a partial structure if ϕ is true in a total structure that extends the partial structure (by moving elements from the R_3 to the R_1 and R_2 blocks of its relations). This means ϕ and $\neg\phi$ can both be partially true in the same partial structure, provided that the structure can be extended in one way to make ϕ true and in another to make $\neg\phi$ true. The same device allows us to represent conceptual indeterminacy in addition to inconsistency. When a mathematical theory is not clearly inconsistent but also does not clearly pick out a determinate (total) structure, the theory can be represented as a collection of partial structures in which propositions to which it is not clear whether the theory is committed are partially true.

Another option is to explain applications of inconsistent and otherwise unrigorous mathematics in terms of classical structures that stand in a more complex relationship to the practice in question. For instance, one might represent scientists as reasoning about different classical structures in different local contexts, even within the same argument.⁵ Alternatively, one might appeal to a structure picked out by a later, more rigorous successor to the mathematical theory. For instance, we might understand early applications of the infinitesimal calculus in terms of modern calculus, or we might understand applications of the Dirac delta function in terms of Schwartz’s theory of distributions.

⁵This is the essence of the “chunk-and-permeate” strategy [Brown & Priest, 2004], even if Brown and Priest don’t explicitly present it in terms of the mapping account.

Why expect such strategies to produce less perspicuous explanations of the success of unrigorous mathematical techniques than RIC? Because crucial to these explanations are the inferentially restrictive methodologies scientists adopt when using these techniques, and, in typical cases, this involves a patchwork of local inferential restrictions (rather than the more general restrictions that might be involved in adopting a particular non-classical logic, say). Such restrictions allow scientists to quarantine mathematically ill-defined concepts to contexts in which they behave in the desired way, making it impossible to use them to derive undesirable (or, in the limit, absurd) results.

Consider, for example, the Dirac delta function, which, if taken to be an honest, extended real-valued function, produces an inconsistent mathematics. Nonetheless, it can be enormously useful if one adopts the inference strategy, explicitly adopted by Dirac, of using it only as a factor within an integrand.

Dirac [1967, §15, p. 58] defines the delta function as a quantity $\delta(x)$ satisfying

$$\int_{-\infty}^{+\infty} \delta(x) dx = 1 \tag{5.1}$$

and

$$\delta(x) = 0 \text{ for } x \neq 0. \tag{5.2}$$

Of course, there can be no such δ if it is construed as an extended real-valued function (and the integral is an ordinary Lebesgue integral). Consider the function $f(x) = 0$. From (5.2), it follows that $f(x) = \delta(x)$ almost everywhere. But from this and an elementary property of Lebesgue integrals, it follows that $\int_{-\infty}^{+\infty} f(x) dx = \int_{-\infty}^{+\infty} \delta(x) dx$ and so, from this and (5.1), we have $0 = 1$. As a result, any mathematical theory that posits the existence of such a function must be inconsistent.

Dirac was, of course, well aware of this, writing, “ $\delta(x)$ is not a function of x according to the usual mathematical definition of a function” [1967, §15, p. 58]. His response was to restrict its use to particular contexts in which it behaved in the desired way and could not

lead to inconsistency:

Thus $\delta(x)$ is not a quantity which can be generally used in mathematical analysis like an ordinary function, but its use must be confined to certain types of expression for which it is obvious that no inconsistency can arise. [Dirac, 1967, §15, p. 58]

The important properties of the delta function for Dirac’s purposes concerned its behavior within integrals. The most important of these for Dirac’s [1967, §15, p. 59] purposes were

$$\int_{-\infty}^{+\infty} f(x)\delta(x)dx = f(0) \tag{5.3}$$

and

$$\int_{-\infty}^{+\infty} f(x)\delta(x - a)dx = f(a) \tag{5.4}$$

for any continuous function $f(x)$ and $a \in \mathbb{R}$.⁶ And even when Dirac introduces elementary properties of his delta function outside of integrals, he makes it clear that these are meant only to inform the manipulation of expressions containing expressions for his delta function within the scope of an integral. They are “essentially rules of manipulation for algebraic work involving δ functions. The meaning of any of these equations is that its two sides give equivalent results as factors in an integrand” [Dirac, 1967, §15, p. 60]. And Dirac justifies these rules by appealing to the properties of the delta function within integrals. So Dirac avoids inconsistency by constraining his use of his delta function to within integrals—or, more accurately, uses it in such a way that it can ultimately end up only as a factor of an

⁶Indeed, these properties are so central to the value of the Dirac delta function that presentations of it sometimes dispense with (5.2) and even (5.1) in favor of conditions more closely related to (5.3) and (5.4). For instance, Messiah [1961, ch. V, §8, pp. 181f] defines the Dirac delta function as the singular function $\delta(x)$ satisfying

$$\int_a^b f(x)\delta(x - x_0)dx = \begin{cases} f(x_0) & \text{for } x_0 \text{ in the interval } (a, b) \\ 0 & \text{for } x_0 \text{ outside the interval } (a, b) \end{cases}$$

for any function $f(x)$ continuous at $x = x_0$. He then mentions (5.1) and (5.2) only as a way to “visualize” the Dirac delta function. Likewise, Swanson [1992, p. 2] defines it via (5.4), and Davey [2003] defines it via (5.1) and (5.3).

integrand, so that no inconsistency can arise:⁷

although an improper function [e.g., the delta function] does not itself have a well-defined value, when it occurs as a factor in an integrand the integral has a well-defined value. In quantum theory, whenever an improper function appears, it will be something which is to be used ultimately within an integrand. Therefore, it should be possible to rewrite the theory in a form in which the improper functions appear all through only in integrands. One could then eliminate the improper functions altogether. [Dirac, 1967, §15, p. 59]

Regardless of how local or global a set of inferential restrictions is, RIC can directly represent it via the specification of the RIC3 component of the representation, the set of privileged inference patterns; disallowed inferences may simply be excluded from RIC3. With these inferences excluded, the result of applying the inference patterns in RIC3 to the initial specification of the target system (RIC2), given the interpretation of the mathematical vocabulary (RIC1), needn't be inconsistent or contain otherwise undesirable propositions meant to be avoided via the inferential restrictions. The inferential restrictions may be specified in an entirely piecemeal fashion or at a very high level of generality, but in applications of unrigorous mathematics an intermediate level of grain will almost always be called for. Consider Dirac's inference strategy. We can very naturally specify RIC3 in this case by including all inference patterns licensed in real analysis given the assumption that the Dirac delta is an extended real-valued function, apart from those in which it doesn't appear as a factor in an integrand. Provided the problematic concepts really do behave in the desired way in these restricted contexts, we then have a quite straightforward explanation of the success (and epistemic legitimacy) of their application: thus restricted, the behavior of those concepts was sufficiently well understood for scientists to judge whether the *physical* inferences they licensed under a given physical interpretation (RIC1) were in accord with their understanding

⁷It is straightforward to see how this allows us to block the reasoning by which we earlier derived $0 = 1$ from the definition of the Dirac delta function. This involved making essential appeal to properties of the function outside an integral (namely, the property that $\delta(x) = 0$ almost everywhere).

of the target system. In Dirac’s case, the restrictions allowed him to do so by showing that the delta function could ultimately be dispensed with altogether.

In contrast, representing such inferential restrictions in terms of a mapping account is considerably less straightforward. Consider, for example, how the partial structures approach would treat the Dirac delta function. Because the properties $\int_{-\infty}^{+\infty} \delta(x)dx = 1$ and $\delta(x) = 0$ for $x \neq 0$ are inconsistent, each must be merely partially true in the partial structure representing the mathematical theory as a whole. This structure is directly mapped to the target structure.⁸ But the important properties of this structure are those it shares with structures in which one or the other is strictly true. When one reasons as if $\int_{-\infty}^{+\infty} \delta(x)dx = 1$, one reasons about a partial structure in which that sentence is strictly true. This reasoning is connected back to the structure in which it is merely partially true by means of a partial morphism. The restriction of the delta function to where it appears as a factor in an integrand can be reflected only indirectly in which structures are mapped back to the one that is interpreted physically and by which morphisms.

Compared to the RIC-based explanation sketched above, the partial structures explanation is positively baroque. It requires a proliferation of new resources—at least four structures and three mappings, even in this relatively simple case—to indirectly represent surface features of Dirac’s practice that RIC can capture by appealing directly to syntactic restrictions that Dirac makes explicitly. And this divergence in complexity becomes more pronounced when

⁸This is a simplification. As Dirac saw it, this inferential restriction meant that the delta function could be dispensed with entirely, albeit at the cost of expressing quantum mechanics in a more cumbersome way. Whenever we find the delta function as a factor within an integral of the form given on the left-hand side of (5.4), we can produce an equivalent expression in which the delta function does not appear by substituting the appropriate instance of the right-hand side of (5.4) for the integral in which the delta function appears. Dirac suggests that this means that the unrigorous delta function introduces no lack of rigor into the physical theory, as the Dirac delta function “is merely a convenient notation, enabling us to express in a concise form certain relations which we could, if necessary, rewrite in a form not involving improper functions, but only in a cumbersome way which would tend to obscure the argument” [Dirac, 1967, §15, p. 59].

If we wished to represent this in terms of the partial structures approach, we would do so by adding a delta-free structure mapped to both the target structure and the partial structure just described. In terms of RIC, we would note that the restriction of the delta function to contexts in which it appears as a factor in an integrand, together with properties connecting such expressions to expressions in which the delta function does not appear, means that we could achieve an equivalent representation by replacing all inference patterns in which the delta function appears in RIC3 with inference patterns in which it doesn’t appear.

scientists' inference strategies are themselves more complex, resulting in a spider's web of structures and morphisms.

This in itself needn't be a problem. A reconstruction of complicated scientific reasoning should itself be expected to be complicated. The question is whether the additional layer of complexity introduced by mapping accounts adds any value, and it's not at all clear to me that it does in this case. I think it's telling that [Bueno & French \[2018, pp. 131ff\]](#), who treat this case at length, never explicitly appeal to partial structures, apart from writing, "we would speculate—although we shall not go into it here—that our framework of partial homomorphisms could quite naturally capture both the open-ended nature of Dirac's theory and the manner in which it can be related to Schwartz's" (pp. 136f). Instead, they write directly about the restriction of the delta function to contexts in which it is a factor in an integrand and how that together with Dirac's algebraic rules for manipulating the delta function in such contexts ensures that the delta function is dispensable. These features of Dirac's practice are at the center of his successful use of the delta function, but they can only be represented very indirectly in terms of partial structures. Explicitly representing the case in terms of partial structures would have distracted from these more important features. In contrast, they could be much more directly represented in terms of the set of inference patterns RIC3. And so this complexity arguably hinders the partial structures approach in its role as a meta-level device for representing philosophically significant features of scientific practice.

Note that the problem here isn't the move from classical to partial structures. The first type of classical structure-based approach, appealing to multiple classical structures to capture different aspects of the mathematical reasoning (cf. [Benham, Mortensen & Priest, 2014](#)), yields a similar proliferation of structures and morphisms without introducing any further resources for representing inferential restrictions. And an approach based on later, rigorous structures has even fewer resources for representing such restrictions. The right diagnosis, I think, is that the inferential restrictions used to successfully apply unrigorous mathematics in practice rarely correspond to neat divisions of some mathematical structure—or even divisions

between several such structures. Even when a structure and mapping that correctly capture the informational content of the representation can be found, these don't lend themselves to a neat account of how they relate to the relevant inferential restrictions. It's in this sense that I claim that RIC facilitates more perspicuous representations of applications of unrigorous mathematics. Both RIC- and mapping-based explanations of the success of such applications, particularly of why it was reasonable for scientists to reason as they did, must make essential appeal to strategies of inference restriction used by those scientists. RIC directly represents such information, while mapping accounts can represent it at best very indirectly, placing more emphasis on the structural gymnastics required to accommodate such episodes.

Still, one might think that the mapping-based explanations are deeper or more substantial than their RIC-based counterparts in that they provide an explanation in terms of structures and mappings of the restricted inferences that RIC directly appeals to. Here, it is worth recalling an important similarity between RIC and mapping accounts. A mapping account's structures and mappings do the same work as the three components of RIC; they represent the informational content of the relevant mathematical scientific representation, which in turn is used to explain mathematically mediated inferences licensed by the representation. In either case, scientists' mathematically mediated inferences are ultimately justified by the supposition that the representation is in fact accurate. And so explanations of why particular mathematically mediated inferences are justified ground out in either case in the informational content of the relevant mathematical representation.

But what about the explanation of Dirac's *success* in using the delta function? Dirac's inferential restrictions are at the foundation of the RIC-based explanation of this success: by choosing the appropriate inferential restrictions, Dirac ensured the delta function behaved as desired, so that the resulting representations had the desired content. A proponent of a mapping account should flesh this out as follows: By making the appropriate inferential restrictions, Dirac ensured that he was reasoning about mathematical structures related in the right way to the relevant quantum mechanical structures. But it might seem that a mapping-based

explanation achieves greater depth than the RIC-based explanation by going one step further, explaining those inferential restrictions in terms of structures and mappings. While the mere fact that mapping accounts represent such restrictions in terms of structures and mappings doesn't entail that they *explain* them in those terms, there is a feature of these restrictions that can be profitably explained in terms of structures and mappings: their appropriateness. Why are the restricted inferences appropriate? Because they are licensed in reasoning about a collection of structures and morphisms that collectively support using the delta function to reason about quantum mechanical structures. But this is just to say, in terms of the mapping account, that they are appropriate because they ensure that the representation has the desired informational content. And as before, this informational content can be specified (very often with greater ease) in terms of RIC to yield a parallel explanation. And so I see no reason to think explanations given in terms of structures and mappings are in virtue of that deeper than other explanations formulated in terms of RIC.

The rest of the chapter is devoted to substantiating the points made here through an examination of two more detailed case studies: Heaviside's operational calculus (§5.2) and path integrals in quantum physics (§5.3).

5.2 Heaviside's operational calculus

In this section, I substantiate the points I made in the previous section by considering Oliver Heaviside's application of his operational calculus, which was notoriously unrigorous. Perhaps most infamous were his operational treatment of fractional differentiation and his treatment of the divergent series expansions that arose in this work. His cavalier approach to these series, coupled with his apparent ignorance of the then-burgeoning theory of divergent series produced by Poincaré and others, led the Royal Society to take the unprecedented step of subjecting his work to peer review and ultimately refusing to publish his work on the subject.⁹ Rather than leading Heaviside to pursue greater mathematical rigor, this seems to have

⁹For useful discussions of this episode, see [Yavetz \[1995, pp. 318–20\]](#) and [Nahin \[2002, pp. 222f\]](#)

had the opposite effect, leading him to more enthusiastically embrace a lack of rigor:

Shall I refuse my dinner because I do not fully understand the process of digestion? No, not if I am satisfied with the result. Now a physicist may in like manner employ unrigorous processes with satisfaction and usefulness if he, by the application of tests, satisfies himself of the accuracy of his results. At the same time he may be fully aware of his want of infallibility, and that his investigations are largely of an experimental character, and may be repellent to unsympathetically constituted mathematicians accustomed to a different kind of work. [Heaviside, 1899, §224, pp. 9f]

For all this, Heaviside's operational work was remarkably successful. There have been no credible challenges to his physical results, and his unrigorous techniques underpinned real conceptual (not just computational) advancement. In particular, his "resistance operators," which were central to his operational techniques, generalized the concept of resistance so that he could extend Ohm's law to time-varying circuits with reactive elements, anticipating the later concept of generalized s -plane impedance, which did the same work more rigorously.

How is it that Heaviside could be so successful in applying his operational calculus despite the highly unrigorous nature of his work? As in the cases considered so far, Heaviside employed a number of strategies to restrict the inferences one could make with incoherent, underdeveloped, or otherwise questionable mathematical concepts. That is, he adopted an inferentially restrictive methodology.

What is distinctive about Heaviside's approach in relation to the cases considered so far is that Heaviside often made inferential restrictions in a strikingly local and ad hoc way, and he frequently used the physical interpretation of the mathematics in a given case to inform the mathematical inferences he ultimately took to be licensed. (The restrictions used in the early calculus and those made by Dirac look relatively simple, general, and well-motivated in comparison.) Explaining how and why these strategies worked requires us to represent Heaviside's inferential practices at a finer level of grain than mapping accounts naturally

allow for. I will argue that RIC can do much better in this regard.

In section 5.2.1, I introduce the operational calculus. In section 5.2.2, I discuss several ways in which the operational calculus fails to meet the standards of rigor of pure mathematics, with a focus on the inferential restrictions Heaviside introduces to manage this lack of rigor. In contrast to Dirac's, Heaviside's inference restriction strategies were strikingly local, piecemeal, and ad hoc. I argue that these features put further stress on mapping accounts but are easily accommodated by RIC. I then consider two ways in which mapping accounts may be thought to have an explanatory benefit outweighing these shortcomings: explaining the role of the physical reasoning in informing Heaviside's mathematical reasoning (§5.2.3) and explaining Heaviside's success retrospectively in terms of later rigorous techniques (§5.2.4). I argue that neither of these considerations favors mapping accounts.

5.2.1 Heaviside's operational calculus and resistance operators

The operational calculus

As a bare mathematical device, Heaviside's operational calculus was a method for solving differential equations algebraically by treating differentiation as an operator. As a toy example, consider the equation $dx/dt = f(t)$. The first step was to reformulate the equations in terms of differential operators. In our example, this would yield $px(t) = f(t)$, with p the operator corresponding to d/dt . Next, Heaviside treated these operators as ordinary algebraic quantities, allowing him to solve the reformulated equations algebraically in terms of functions of these operators, construed as algebraic quantities. Call this the "operational solution." In our example, solving for $x(t)$ in the operationalized equation yields the operational solution $x(t) = p^{-1}f(t)$. Heaviside would then "algebrize" this solution to eliminate all reference to functions of differential operators, yielding the desired solution to the original differential equations. He most often achieved this by expanding functions of p in his operational solution in ascending or descending powers of p and applying rules for replacing particular expressions containing p with expressions for functions of t , though he occasionally used substitutions

that did not require such power series expansions.¹⁰ In our simple example, no power series expansion is required, but the expression p^{-1} requires an interpretation. Heaviside usually interpreted p^{-1} as the definite integral $\int_0^t du$, which in this case gives us $x(t) = \int_0^t f(u)du$.

Heaviside is far from the originator of methods treating differential operators as algebraic quantities independent of the functions they operate on. This was made possible by Leibniz's d/dx notation for differentiation and subsequently developed by many of the biggest names in late-eighteenth- and nineteenth-century mathematics, including Lagrange, Laplace, Fourier, Cauchy, and Boole, among others.¹¹ Heaviside is known to have studied the relevant work of Fourier and Boole in particular [Cooper, 1952, p. 12], on the basis of which he likely developed his own version of the operational calculus. While the novelty of Heaviside's approach, construed purely as a piece of mathematics, is up for debate,¹² the *primary* contribution of Heaviside's operational calculus was not that it taught mathematicians anything they did not already know about differential equations. Indeed, these mathematicians were able to prove more general results—and with greater rigor—than Heaviside. Rather, Heaviside's primary contribution in this work was, as Nahin [2002, p. 218, emphasis in original] puts it, to show “*how to apply to real, physical problems of technological importance analytical techniques that had up till then been symbolic abstracts.*”

Resistance operators

Central to this task were what Heaviside called “resistance operators.” These he used to present an early generalization of Ohm's Law.¹³

¹⁰Cf. the schema presented by Lützen [1979, §I.5, pp. 170–2].

¹¹For detailed treatment of these historical antecedents, albeit without reference to Heaviside, see Koppelman [1971]. For a discussion that more explicitly ties this history to Heaviside, see Cooper [1952].

¹²For instance, Cooper [1952] argues that Heaviside's results are not very novel at all. Lützen [1979] responds by citing a number of distinguishing features of Heaviside's techniques. Petrova [1987] argues in turn that these distinguishing features were anticipated by Cauchy, Gregory, and Boole, but there is no consensus on this [see, e.g., Yavetz, 1995, p. 310, note 4].

¹³Prior to introducing this concept in 1887 in “On Resistance and Conductance Operators” [Heaviside, 1894, pp. 255–74], Heaviside used these techniques extensively in “On the Self-Induction of Wires” [Heaviside, 1894, pp. 168–323], first published in 1886–7. Aspects of them appear as early as “The Induction of Currents in Cores” [Heaviside, 1892, pp. 353–416], first published in 1884–5.

If we regard for a moment Ohm's law merely from a mathematical standpoint, we see that the quantity R , which expresses the resistance, in the equation $V = RC$,¹⁴ when the current is steady, is the operator that turns the current C into the voltage V . It seems, therefore, appropriate that the operator which takes the place of R when the current varies should be termed the resistance-operator. To formally define it, let any self-contained electrostatic and magnetic combination be imagined to be cut anywhere, producing two electrodes or terminals. Let the current entering at one and leaving at the other terminal be C , and let the voltage be V , this being the fall of potential from where the current enters to where it leaves. Then, if $V = ZC$ be the differential equation (ordinary, linear) connecting V and C , the resistance-operator is Z . [Heaviside, 1894, p. 355]

While Ohm's Law could be used only for the analysis of DC circuits in steady state or time-varying circuits with only resistive elements (i.e., with no reactance), the concept of resistance operator could be used to generalize it to circuits with reactive elements and arbitrary time-varying voltages or currents: $V = ZC$ with Z the resistance operator for the circuit rather than its total resistance. While some work had been done to generalize Ohm's law to time-invariant AC circuits in the years immediately preceding Heaviside's work,¹⁵ Heaviside's resistance operators were a significant step forward in that they did not require the assumption that the relevant voltages and currents were sinusoidal. After all, a resistance operator is *whatever* turns the current between two points in a circuit into the voltage between those points. Due to this high degree of generality, resistance operators could be used to analyze the behavior of a much wider range of systems.

Heaviside's operational calculus was central not just to how resistance operators came to be able to play these roles, but also to the very possibility of expressing them. The resistance

¹⁴Heaviside deviates from modern usage here in using C rather than I for current. I follow him in this throughout this paper for the sake of consistency.

¹⁵The first to do so seems to have been Wietlisbach [1879], and these techniques were subsequently refined by Oberbeck [1882] and popularized in Britain by Lord Rayleigh [1886a, 1886b, 1891]. For a useful discussion of this history, see Kline [1992, pp. 77ff].

operators for circuit elements were derived from the usual equations describing their relation to voltage and current by solving for V/C . In the case of a resistor with resistance R , the relevant equation is Ohm's Law, and solving for V/C yields $Z_R = V_R/C_R = R$. On the other hand, in the case of reactive elements, the operational calculus becomes crucial, as the relevant equations are differential equations relating inductance or capacitance (respectively) to voltage and current. In the case of an inductor with inductance L , the relevant equation is $V_L(t) = L \frac{dC_L(t)}{dt}$, derived from Faraday's law. Here, to derive the resistance operator $Z_L = V_L/C_L$, we must use the operational calculus. Substituting Heaviside's p operator for d/dt yields $V_L = LpC_L$, and solving for V_L/C_L yields $Z_L = V_L/C_L = Lp$.

The resistance operators for circuit elements can then be combined to yield the resistance operator for the whole circuit, just as resistances are combined to yield the resistance of circuit as a whole in the case of steady-state DC circuits or circuits with no reactive elements. In the latter case, the total resistance of a collection of resistors R_1, \dots, R_n in series is $R_{\text{total}} = R_1 + \dots + R_n$, and for resistors in parallel,

$$\frac{1}{R_{\text{total}}} = \frac{1}{R_1} + \dots + \frac{1}{R_n},$$

Resistance operators behave exactly the same in this respect.¹⁶ For circuit elements in series, $Z_{\text{total}} = Z_1 + \dots + Z_n$, and for circuit elements in parallel,

$$\frac{1}{Z_{\text{total}}} = \frac{1}{Z_1} + \dots + \frac{1}{Z_n}.$$

A simple example: Step response of an RL circuit

For a simple example of how this worked in practice, consider a circuit consisting of an ideal resistor with resistance R and an ideal inductor with inductance L in series—or, equivalently, a coil with resistance R and inductance L —with a constant external voltage e applied at $t = 0$.

¹⁶Not only that, but they can be derived from Kirchhoff's laws and Ohm's law for resistance operators in exactly the same way as the rules for combining resistances are derived from Kirchhoff's laws and Ohm's law for resistors.

One thing we might want to know about this circuit is the resulting current C as a function of time. Heaviside discusses this problem in §283 of the second volume of *Electromagnetic Theory* [Heaviside, 1899, pp. 129f].




Figure 5.1: A DC circuit with a resistor and inductor in series, as treated in [Heaviside, 1899, §283, pp. 129f].

Since the resistor and inductor are in series, the resistance operator for the whole circuit is just the sum of the resistance operators of these elements—i.e., $Z = Z_R + Z_L$. And so, making the appropriate substitutions, we have $Z(p) = R + Lp$. To represent the external voltage's being applied at time $t = 0$, we represent the voltage by that external voltage e multiplied by the unit step function¹⁷:

$$\mathbf{1}(t) = \begin{cases} 0 & \text{if } t \leq 0 \\ 1 & \text{if } t > 0 \end{cases}.$$

Making the relevant substitutions in $V = ZC$ and solving for C then yields $C = e\mathbf{1}/(R + Lp)$.

This is the operational solution of the problem.

Heaviside's next step is to “algebrize” this solution. This is required to yield an expression for C as a function of t , rather than p . To do this, Heaviside expands the right-hand side of

¹⁷In more recent texts, it is more common to use $H(t)$ for the unit step function in Heaviside's honor. Here I use a boldface $\mathbf{1}$ to remain close to Heaviside's own notation, while marking the difference between the unit step function and the integer 1. Heaviside typically left multiplication by the unit step function $\mathbf{1}$ implicit in such cases, much as one normally leaves multiplication by the integer 1 implicit. In the rest of this discussion, I've added in instances of $\mathbf{1}$ where Heaviside leaves them implicit in his treatment of this example in [Heaviside, 1899, pp. 129f] for the sake of clarity.

the operational solution in descending powers of p , yielding

$$C = \frac{e}{R + Lp} \mathbf{1} = \frac{e}{Lp(1 + R/Lp)} \mathbf{1} = \frac{e}{R} \left(\frac{R}{L} \cdot \frac{1}{p} - \left(\frac{R}{L} \right)^2 \cdot \frac{1}{p^2} + \left(\frac{R}{L} \right)^3 \cdot \frac{1}{p^3} - \dots \right) \mathbf{1}.$$

To make sense of this, Heaviside had to give meaning to the expression $1/p$ (or p^{-1}). In cases like this, Heaviside interpreted this as the inverse operator of $p = d/dt$ —since algebraically we should have $p \cdot p^{-1} = 1$, and multiplication by 1 should correspond to the identity operator—and took this inverse operator to be $1/p = \int_0^t du$. $1/p^n$ then comes to represent n -fold integration.

So we have

$$\frac{1}{p} \cdot \mathbf{1} = \int_0^t \mathbf{1} \, du = \begin{cases} 0 & \text{if } t \leq 0 \\ t & \text{if } t > 0 \end{cases}$$

and so

$$\frac{1}{p^n} \cdot \mathbf{1} = \begin{cases} 0 & \text{if } t \leq 0 \\ t^n/n! & \text{if } t > 0. \end{cases}$$

So the power series can be rewritten as

$$C = \frac{e}{R} \left(\frac{R}{L} \cdot \frac{1}{p} \mathbf{1} - \left(\frac{R}{L} \right)^2 \cdot \frac{1}{p^2} \mathbf{1} + \left(\frac{R}{L} \right)^3 \cdot \frac{1}{p^3} \mathbf{1} - \dots \right) = \frac{e}{R} \left(\frac{R}{L} t - \left(\frac{R}{L} \right)^2 \frac{t^2}{2!} + \left(\frac{R}{L} \right)^3 \frac{t^3}{3!} - \dots \right),$$

which Heaviside recognized as the power-series expansion of

$$C = \frac{e}{R} \left(1 - \epsilon^{-(R/L)t} \right)$$

for $t \geq 0$, the correct result.¹⁸

¹⁸Here again, I follow Heaviside in using ϵ rather than e for Euler's number.

5.2.2 Failures of rigor and inferential restrictions

Layered local inferential restrictions

We are already in a position to observe several ways in which these techniques failed to live up to the standards of mathematical rigor—in the hands of Heaviside at any rate. In each case, the failures of rigor do not involve the derivation of any straightforwardly incorrect result, but rather, as [Yavetz \[1995, p. 317\]](#) puts it, “the use of terms and procedures that are seldom fully defined”. As in other examples of unrigorous mathematics, the result is a degree of indeterminacy in the global mathematical commitments of the operational calculus, and Heaviside avoided deriving incorrect results by adopting an inferentially restrictive methodology. However, as we will see, these restrictions were more local, piecemeal, and ad hoc than in other cases, and his ultimate justification for them was atypical. As a result, these inference strategies are even more difficult to treat straightforwardly in structural terms than those considered in section 2, while RIC can again represent them with relative ease.

An illustrative problem present in the example in §3.3 was that Heaviside's techniques often involved interpreting the inverse of his time-differentiation operator $p = d/dt(\cdot)$ as $p^{-1} = \int_0^t \cdot dx$, but, thus interpreted, these are not generally inverse operators. Following [Nahin \[2002, p. 232\]](#), consider again the simple differential equation $dx/dt = f(t)$. As we saw in the previous section, applying Heaviside's operational techniques yields the solution $x(t) = \int_0^t f(u)du$. But this entails that $x(0) = 0$, since $\int_0^0 f(u)du = 0$ for any f we like. Since not every function x has this property, p and p^{-1} are not generally inverse operators. This problem is blatant and easily avoided in this simple case by restricting the use of this reasoning to cases in which we know that $x(0) = 0$ (for all functions x that might be the operand of p or p^{-1}). Heaviside observed this in practice by restricting his attention almost exclusively to circuits in which the voltage and current are enveloped by the unit step function 1.

However, further restrictions turn out to be required. If p and p^{-1} are inverse operators, they must also be commutative, and Heaviside ultimately rests his justification of this

property on an interpretation of $p\mathbf{1}$ as the (inconsistent!) Dirac delta function:

Thus $pp^{-1}\mathbf{1} = pt = 1$ but $p^{-1}p\mathbf{1} = [p^{-1}]0 = 0$, unless we say $p^{-1}p\mathbf{1} = p^{-1}\frac{t^{-1}}{(-1)!} = \frac{t^0}{0!} = 1$. This property has to be remembered sometimes. [Heaviside, 1899, §358, p. 298]

$p\mathbf{1}(t)$ must be zero except when $t = 0$, since it jumps from 0 to 1 at $t = 0$ and is otherwise constant. If $p^{-1}p\mathbf{1} = \int_0^t p\mathbf{1}(u)du = 1$, then $p\mathbf{1}$ can only be the Dirac delta function.¹⁹

We should interpret Heaviside's move here not as an attempt to rigorously justify the commutativity of the p and p^{-1} operators but instead as a way of managing inconsistent demands on the behavior of these operators. In essence, Heaviside is adopting an inferential strategy that makes the line of reasoning that leads to the contradiction in this case—namely, $p^{-1}p\mathbf{1} = p^{-1}0 = 0$ —off-limits. Rather, $p\mathbf{1}$ must be reasoned with as if it were the Dirac delta function, so that only the second line of reasoning— $p^{-1}p\mathbf{1} = p^{-1}\frac{t^{-1}}{(-1)!} = \frac{t^0}{0!} = 1$ —is allowed. This second line of reasoning is still unrigorous, since it appeals to the Dirac delta function under the guise of $p\mathbf{1} = \frac{t^{-1}}{(-1)!}$. But by limiting oneself to one or the other of these lines of reasoning, one rules out an obvious way of deriving an explicit contradiction from the inconsistent properties of the Dirac delta function.

While this on its own is not enough to guarantee that one won't derive explicit contradictions via the inconsistent properties of the delta function, Heaviside does seem to have implicitly adopted a strategy similar to Dirac's strategy of using the delta function only as a factor in an integrand. In Heaviside's case, it was the step of "algebrization" that allowed instances of $p\mathbf{1}$ to be dispensed with via identification of $p^{-n}\mathbf{1}$ with $t^n/n!$ only after the algebraic manipulation of p -expressions required to derive the operational solution had been completed.

¹⁹Recall that the Dirac delta function is defined by the properties $\delta(x) = 0$ for $x \neq 0$ and $\int_{-\infty}^{\infty} \delta(x)dx = 1$. In the presence of the first property, the second is equivalent to $\int_0^t \delta(x)dx = p^{-1}\delta = 1$. This, incidentally, is why the many attempts to provide the operational calculus with a purely algebraic, as opposed to analytic, foundation failed. Making p and p^{-1} commute means living with the Dirac delta function or something very much like it. For a useful historical treatment of such approaches, see Lützen [1979, pp. 188ff].

So to ensure that p and p^{-1} were inverse operators while avoiding nasty side effects, Heaviside had to observe several inferential restrictions at different levels of grain: (1) That property could only be appealed to when those operators were applied to functions enveloped by the unit step function, (2) p must be taken to be the ordinary time derivative *except when applied to step functions*, in which case its integral is non-zero, and (3) reasoning as if $p1$ were the Dirac delta was restricted to a certain part of Heaviside's overall operational procedure, during which those operators are treated merely algebraically prior to the (counterintuitively named) step of "algebrization," during which they were again interpreted analytically.

While Heaviside's practice was far from ideal, RIC allows us to express a natural account of why it was nonetheless reasonable. Each of these local inferential restrictions served to rule out certain ways of working with p and p^{-1} that produced undesirable algebraic behavior by appealing to properties of those operators that were otherwise useful. When these concepts were constrained to contexts in which their algebraic behavior was understood—represented by RIC in terms of excluding the relevant patterns of inference from RIC3—they yielded determinate numerical results, which could be compared with known results or with Heaviside's physical understanding of his target systems via the interpretation given by RIC1. In essence, by restricting those operators to contexts in which they are well-behaved, Heaviside ensured that his representations (at least as far as those operators are concerned) have determinate accuracy conditions and so could be evaluated in terms of their agreement with both theoretical understanding of the target phenomena and experimental results.

Mapping accounts may be able to tell a similar story, but only in a cumbersome way. Consider how the partial structures approach might do so, which can again be naturally adapted to an approach based on classical structures. Directly mapped to the target structure will be a structure in which, for each problematic property of p and p^{-1} , it is merely partially true that they have that property: that p and p^{-1} are inverse operators, that p and p^{-1} *aren't* inverse operators, that p and p^{-1} commute, that p and p^{-1} *don't* commute, that $p^{-1}p1 = 0$, that $p^{-1}p1 = 1$, and so on. In particular contexts, one reasons with partial structures extending

this one in which some of these statements are strictly true. Each of these structures is then mapped back to the structure in which all these statements are merely partially true either directly or through mappings to intermediary structures. Heaviside's inferential restrictions are then represented indirectly in terms of which structures can be reasoned about in which contexts and how they can be mapped to other structures in this family. As in the case of the Dirac delta, I suspect that the inclusion of partial structures gives only a false appearance of explanatory depth here. Each explanation ultimately bottoms out in properties of the inferential restrictions Heaviside observes. Representing those restrictions indirectly in terms of structures adds unnecessary complexity and obscures the features of Heaviside's practice that do the explanatory heavy lifting.

Ad hoc inferential restrictions

A further striking feature of Heaviside's inference strategies is their often ad hoc nature. Heaviside generally didn't lay them out in advance and expressed comfort with the possibility that such techniques might, if used injudiciously, lead to inconsistent results.

One striking example is that Heaviside frequently treated his operators in general (not just p and p^{-1}) as if they were commutative, though they don't generally have this property. Heaviside explicitly noted this but made no attempt at a general explanation of when such moves were permissible. In a representative passage, he writes,

The reader may have noticed in the above, and perhaps previously, that we change the order of operations at convenience, as in $f(p)\phi(p)\mathbf{1} = \phi(p)f(p)\mathbf{1}$, and that it goes. But I do not assert the universal validity of this obviously suggested transformation. It has, however, a very wide application, and transforms functions in a remarkable manner. *Reservations should be learnt by experience.* [Heaviside, 1899, §251, pp. 59f, my italics]

The idea seems to be to freely appeal to the commutativity of his operators in contexts where this *works*, determining which contexts these are through experimentation.

Heaviside expressed the same attitude toward one of his most central results, the expansion theorem, which Lützen [1979] calls “Heaviside’s most important tool in algebrizing procedures” (p. 170): where $e = ZC$ is an operational solution, e is a constant multiplied by the unit step function, and “the form of Z [is] such as to indicate the existence of normal solutions for C ,”

$$C = \frac{e}{Z_0} + e \sum \frac{e^{pt}}{p \frac{dZ}{dp}}$$

[Heaviside, 1899, §282, p. 127]. This was a powerful tool because it worked in such a wide range of circumstances, including for algebrizing operational solutions of partial differential equations describing continuous telegraph circuits. But again it was not true generally. Heaviside insisted that it was actually *undesirable* to state precisely the conditions in which it can be used:

Now it would be useless to attempt to state a formal enunciation to meet all circumstances. [...] *It is better to learn the nature and application of the expansion theorem by actual experience and practice.* [Heaviside, 1899, §282, p. 128, my italics]

So what do we make of this? What Heaviside seems to be proposing is what we might call an *ad hoc* inference restriction strategy: inferences are restricted not in advance, but only as we discover that certain patterns of inference produce incorrect or undesired results. Until we make such a discovery we must “keep [our] eyes and [our] mind open, and be guided by circumstances” [Heaviside, 1899, §223, p. 3]. How do we make sense of the success of such a strategy, and, in particular, to what extent was it reasonable for Heaviside to adopt it?

Ultimately, such a strategy balances the benefits of a mathematical opportunism with a recognition that results thus obtained are more fallible than those obtained by rigorous means. Recognizing this fallibility means continually scrutinizing the results achieved with suspect mathematics on the basis of their agreement with known results and one’s independent understanding of the target system. RIC offers considerable flexibility in how we represent this

scrutiny.

On one hand, we might think of such a strategy as one in which inferences are added to RIC3 in piecemeal fashion, so that the commitments of the representation never outstrip those results that have actually been derived. When an inference would lead to an inconsistent result or to a result that can otherwise be ruled out—for instance, because it has physical consequences that are known to be incorrect—that inference is simply not added to RIC3. As a result of this extreme conservatism, the representation cannot be committed to anything undesirable. This might be relaxed in contexts in which the problematic mathematics is better understood. Heaviside writes of “numerical groping” only as a technique of last resort “when [physical] intuition breaks down” [Heaviside, 1899, §437, pp. 461f]. In less desperate cases, these piecemeal additions might be higher-level inference *types* in which one has gained confidence.

On the other hand, we might think of such a strategy as one in which inferences are *excluded* from RIC3 in piecemeal fashion, so that inferences are only disallowed when they are shown to lead to results that are inconsistent or can otherwise be ruled out. In this case, the representation is very likely to be inconsistent, but its use does not require its users to commit themselves to its accuracy. Instead, a user of such a representation might only commit themselves to the accuracy of those results actually derived and scrutinized (in parallel fashion to the previous case), reserving judgment about other commitments of the representation or at least keeping their fallibility well and truly in mind.

Ultimately, a combination of the two approaches is likely to be most useful, with the former accounting for Heaviside's own commitments and the latter accounting for his opportunistic heuristic use of less well-understood techniques. Inference patterns from the latter are incorporated into the former only after they have been suitably scrutinized. This allows epistemically suspect results of unrigorous techniques to be quarantined even when one does not yet have a good understanding of the contexts in which they can safely be applied.

Now consider how the partial structures approach would treat this inference strategy.²⁰ In this case, the move corresponding to the addition of an inference to RIC3 is extending the existing morphisms between the partial structures at work in the representation in a way that licenses the inference and, when the representation cannot be made to license the inference otherwise, adding a further partial structure and morphism. In this way, the partial structures approach represents ad hoc reasoning in terms of ad hoc choices of structure and mapping. Now, as the addition of inferences becomes more and more piecemeal, it becomes less clear to me that the partial structures approach has the resources to license exactly those inferences without licensing further inferences not licensed by the corresponding RIC representation. But more importantly, this again introduces more complexity to do the same explanatory work. Because the relevant structures will often in practice be several morphisms away from the one that is directly mapped to the target structure, particularly when other strategies of inference restriction are also in use, the requisite additions will often be less than straightforward. And it is again unclear that the formal apparatus used to represent these inferential restrictions adds depth or substance to the explanation. What seems to do the work in explaining why such techniques were epistemically justifiable is simply the practice of withholding judgment about certain mathematical inferences until further support for their conclusions is found. And this practice is again more straightforwardly represented in terms of RIC.

5.2.3 Failures of rigor and “physical mathematics”: The physical demand for fractional differentiation

Rather than the failures of rigor considered so far, it was Heaviside's treatment of fractional differentiation and divergent series that most upset his contemporaries. The latter ultimately led the Royal Society to begin subjecting his submissions to peer review and justified the rejection of his final submission [Cooper, 1952, p. 14].

The second volume of Heaviside's *Electromagnetic Theory* [Heaviside, 1899] begins with a

²⁰It can again be adapted to an approach positing multiple classical structures in a straightforward way.

spirited, but sometimes bitter and defensive, justification of his unrigorous techniques in the wake of this rejection. In addition to presenting a number of practical virtues of his unrigorous techniques, he presented an approach to mathematics that is deeply grounded in physical reasoning. In physics, the physical interpretation of the mathematics is to be kept in mind at all times, so that physical knowledge and intuition can guide one's mathematical work. [Heaviside, 1899, §224, pp. 4f; see also §437, pp. 460f] We've already seen one way in which it can do so: serving as means of checking results derived mathematically. But mathematics, according to Heaviside, is also answerable to physics in that, if a physical representation requires us to use a mathematical expression that appears to be mathematically meaningless, we can conclude that that piece of mathematics is indeed meaningful and use the physics to elucidate its behavior [Heaviside, 1899, §224, pp. 6f]. Even if we reject the idiosyncratic empiricist philosophy of mathematics Heaviside used to bolster these claims, we can make sense of them in a context where the needs of physics outstrip existing mathematical resources. In such cases, physics can legitimately guide the development of new mathematical theories and techniques tailored to particular kinds of physical problem. This sort of inferential move from physics to mathematics is common in Heaviside's practice.

Heaviside was driven to the topics of fractional differentiation and divergent series by his representation of semi-infinite, continuous systems, particularly the representation of a semi-infinite transmission line via the telegraph equations

$$-\frac{\partial V(x, t)}{\partial x} = RC(x, t) + L \frac{\partial C(x, t)}{\partial t}$$

$$-\frac{\partial C(x, t)}{\partial x} = KV(x, t) + S \frac{\partial V(x, t)}{\partial t},$$

where R , L , K , S are resistance, inductance, "leakance" (conductance between the signal and return wires), and capacitance per unit length of the telegraph wire, respectively.²¹ Replacing

²¹Heaviside's treatment of this case can be found in [Heaviside, 1899, chapter 7].

$\partial/\partial t$ with Heaviside's p and doing some algebra yields

$$V = \sqrt{\frac{R + Lp}{K + Sp}} C.$$

Producing numerical solutions or indeed any solutions expressed in terms of functions of t rather than p , even for special cases in which we ignore one or more of R , L , K , and S , requires making mathematical sense of expressions like ' $p^{\frac{1}{2}}\mathbf{1}$ '. Heaviside's confidence in the representation of such systems given by the telegraph equations grounded his conviction that this operational equation must have numerical solutions (if it is to adequately represent these systems) and thus that there was indeed sense to be made of such expressions.²²

Heaviside did so via the equation $p^{\frac{1}{2}}\mathbf{1}(t) = (\pi t)^{-\frac{1}{2}}$, a result known at least as early as 1819 by Sylvestre Lacroix [1819, pp. 409f], but derived independently by Heaviside by more "experimental" means.²³ One such derivation is the following. If we ignore leakage and inductance, we can derive the equation

$$C = (Sp/R)^{\frac{1}{2}} e\mathbf{1}$$

for the current at $x = 0$ where $e = V(0, 0)$. One way to algebrize this operational equation is by considering the case in which the wire has finite length l , producing a Fourier series expansion for the finite case (via Heaviside's expansion theorem), and taking the limit as $l \rightarrow \infty$. Using the result to calculate the current at $x = 0$ yields a new expression for the

²²Heaviside made a similar move in interpreting the divergent series expansions of various operational solutions of the telegraph equation. He made no reference to the then-burgeoning theory of divergent series, but he rightly recognized certain divergent series expansions as what we would now call "asymptotic expansions" of the relevant functions on the basis of their physical meaning (i.e., the meaning of the physically interpreted mathematical expression). He even correctly conjectured that the divergent parts of these asymptotic expansions are meaningful, carrying information about the exact value of the function that they approximate, again based on a physical interpretation of the components of the series. Nonetheless, as John R. Carson observed, "the precise sense in which the expansion asymptotically represents the solution cannot be stated in general, but requires an independent investigation in the case of each individual problem" [Carson, 1926, p. 78]. This relied on a notion of "equivalence" between convergent and divergent series that Heaviside left undefined [e.g., Heaviside, 1899, §340, p. 250].

²³For an extended treatment of the history of fractional differentiation, see Ross [1977].

current at $x = 0$:

$$C = \frac{2e}{R\pi} \int_0^\infty e^{-s^2 t/RS} ds = (S/R\pi t)^{\frac{1}{2}} e.$$

Comparing our two formulae for current at $x = 0$ and doing a little more algebra yields $p^{\frac{1}{2}} \mathbf{1} = (\pi t)^{-\frac{1}{2}}$ (Heaviside, 1899, §350; cf. §240). After deriving this, he writes, “The above is only one way in a thousand. I do not give any formal proof that all ways properly followed must necessarily lead to the same result” [Heaviside, 1899, §350, p. 288]. Despite the lack of assurance that this is the unique possible result and despite its reliance on a particular special case (the telegraph equation without inductance or leakage), in the rest of the same chapter Heaviside uses this equation to provide an account of more general fractional differentiation, including half-integer differentiation and cases in which polynomials in p occur under the radical, as well as the operational solution of the telegraph equations in their full generality.

Here Heaviside moves from a physical system represented in operational terms, a semi-infinite telegraph cable with no inductance or leakage, to a conclusion about the mathematics apparently used to represent this very system. The thought seems to be that, *whatever the underlying mathematics*, if it is to play the right role in representing this particular physical case, then it must interpret $p^{\frac{1}{2}} \mathbf{1}$ as $(\pi t)^{-\frac{1}{2}}$ (at least in this instance). And if it does that, it must interpret $p^{\frac{n}{2}} \mathbf{1}$ for odd n as $p^{\frac{n-1}{2}} p^{\frac{1}{2}} \mathbf{1}$. Consideration of other cases determines a value for $p^{\frac{1}{4}} \mathbf{1}$. Using these values that are necessitated by the physics (in its operational representation), Heaviside then can test possible extensions of p^x for arbitrary $x \in \mathbb{R}$. It is agreement with these results, rather than reasoning from first principles, that justifies his choice of extension of p^x to arbitrary real x .²⁴ [Heaviside, 1899, ch. 7]

It might seem at first that this is a case in which mapping accounts can provide a significant explanatory benefit even in the absence of a well-understood mathematical theory, contrary to my claims in section 2 and in contrast to the cases considered in the previous

²⁴This reasoning from first principles would likely involve the gamma function, which extends the factorial function to the complex numbers. In the end, Heaviside does much of the work of defining the gamma function for arguments with non-positive real parts via analytic continuation (albeit restricting his attention to the reals), but this again is used because it is the most expedient way to define p^x for negative x to suit the physical cases at hand. [Heaviside, 1899, §425, p. 435]

section. Such a benefit might in turn justify the greater complexity with which mapping accounts must represent Heaviside's strategies of inference restriction. According to this line of thought, Heaviside learned about the relevant mathematical structure by making inferences about that structure on the basis of its structural relation to its target system. So even if Heaviside didn't start with a well-understood mathematical theory that neatly picked out a particular structure, mapping accounts can nonetheless explain how he came to understand some of the properties of the mathematics needed to represent his chosen target systems.

But closer analysis of the case doesn't bear this out. In this case, the only conclusion that Heaviside used considerations about the physical system to directly support is that fractional powers of p must be able to be used meaningfully. The operational version of the telegraph equations, which Heaviside took to accurately represent the relevant physical systems, can only be algebrized if such expressions can be manipulated, and only the algebrized equations, expressed in terms of functions of time rather than p , can be used to derive numerical solutions, which, when interpreted, yield determinate predictions about the behavior of the target system. That is, if the operational representation can be used at all, fractional powers of p must be able to be manipulated. So far, this means only that inferences in which such expressions appear can't be ruled out wholesale by any inference restriction strategy Heaviside might adopt for reasons like those discussed in the previous section.

When Heaviside determined how $p^{\frac{1}{2}}1$ should be reasoned with, it was by comparing an operational expression for a given physical quantity— $C = (Sp/R)^{\frac{1}{2}}e1$ —with an expression for the same quantity that can be derived by extending existing algebrization procedures to this new context— $C = (S/R\pi t)^{\frac{1}{2}}e$. The result can be interpreted as telling us how $p^{\frac{1}{2}}1$ must be reasoned with provided that the expansion theorem can be extended to this context in this way. But, as we saw earlier, Heaviside made a point of not committing to the uniqueness of his interpretation of $p^{\frac{1}{2}}1$. A simple extension of his algebrization strategies to this new case might yield inconsistent results, in which case some of the inferences involved in deriving these inconsistent results might need to be restricted, perhaps even those that allowed him to

derive $p^{\frac{1}{2}}\mathbf{1}(t) = (\pi t)^{-\frac{1}{2}}$. So the inference of the properties of fractional powers of p shouldn't be represented as a simple inference of the properties of one (mathematical) structure on the basis of another (physical) structure and a morphism between them. Instead, the same sorts of strategies of inference restriction discussed in the previous section must still be at the heart of an explanation of why Heaviside's treatment of fractional differentiation was epistemically respectable—at least from a physical, rather than mathematical, perspective. And so the arguments from that section apply here as well.

5.2.4 The Laplace transform in heavy disguise?

So far, I have largely ignored the possibility of making sense of Heaviside's success in terms of a structure picked out by later, more rigorous alternatives to his techniques. This is because I have focused on why his techniques were epistemically justifiable at the time he used them, and I see no reason to think that a mapping account appealing to later rigorous structures would be any better situated to provide that sort of explanation than the versions of the mapping account I've considered so far. But there is another sort of explanation of Heaviside's success that explains why the *results* of those techniques were correct, regardless of their epistemic status at the time, in terms of later mathematics.

We can appreciate the importance of such explanations regardless of whether we feel the pull of mapping accounts. For instance, [Wilson \[2006, ch. 8\]](#) uses the case of Heaviside's operational calculus to argue against a classical view of concepts according to which their meaning must be grasped once and for all at the outset. This presents a difficulty for mapping accounts, as it means that many of scientists' mathematically mediated inferences aren't explicitly grounded in the existence of structural relations, as scientists' concepts don't suffice to pick out the needed structures. But Wilson makes sense of Heaviside's inferences in a local way through what he calls "correlational pictures", "generic stories that speakers tell themselves with respect to how their predicate's usage matches to worldly support within normal circumstances of application" (p. 516). Heaviside's contemporaries gave him grief because "he

was unable to supply orthodox mathematical underpinnings for his procedures in terms of an adequate associated picture" (p. 521). But Heaviside was vindicated by subsequent rigorous work that provided such a picture. And mapping accounts have a neat story to tell about how this later work did so. The question is then whether this gives us good reason to favor mapping accounts over RIC, and in this section I argue that it doesn't.

This later mathematics was largely built on the Laplace transform and its inverse.²⁵ Heaviside biographer Paul J. Nahin goes as far as to write, "in fact, Heaviside's operational calculus is just the Laplace transform in heavy disguise" [2002, p. 218]. The Laplace transform maps functions in the time domain to functions in the s -domain, where s is a complex number whose imaginary component represents the function's periodic behavior (frequency response) and whose real component represents its non-periodic behavior (e.g., its decay). The Laplace transform $F(s)$ of a function $f(t)$ is given by the integral $F(s) = \int_{0^-}^{\infty} f(t)e^{-st} dt$. The result typically looks a lot like Heaviside's operational solution of the same problem, but with s taking the place of p . Consider again the example in §3.3. Just as Heaviside started by calculating the resistance operators for the components of the circuit, a contemporary electrical engineering student could start by calculating the s -plane impedance of each circuit component by applying the Laplace transform to the differential equation characterizing it, yielding $Z_R(s) = R$ where Heaviside has $Z_R(p) = R$, $Z_L(s) = Ls$ where Heaviside has $Z_L(p) = Lp$, and so on. Corresponding to Heaviside's operational solution $C = e1/(R + Lp)$, they would arrive at the Laplace transform for the circuit $I(s) = 1/(R + Ls)$ (using the modern notation, I instead of C , for current) and multiply it by the Laplace transform of the input signal, in this case V_0/s (the Laplace transform of the unit step function multiplied by the value of the voltage source).²⁶ Corresponding to Heaviside's "algebrization", they would translate this back to the time do-

²⁵ As Wilson [2006, p. 531] points out, the Laplace transform on its own doesn't suffice to vindicate all of Heaviside's techniques. In particular, those involving his use of $p1$, interpreted as the Dirac delta, should be understood in terms of Schwartz's theory of distributions. Here I focus on the Laplace transform, but the points I make can be naturally extended to other pieces of mathematics used to retrospectively vindicate Heaviside.

²⁶ Strictly speaking, one should take the Laplace transform of each side of the differential equation characterizing the circuit as a whole, and most treatments work directly with this equation. But this can be done in terms of the Laplace transforms of the circuit elements thanks to the linearity of the Laplace transform, which allows for an approach closer to Heaviside's.

main via the the inverse Laplace transform, given by $f(t) = \frac{1}{2\pi i} \lim_{T \rightarrow \infty} \int_{\gamma-iT}^{\gamma+iT} F(s)e^{st} ds$. In this case, applying the inverse Laplace transform to both sides yields $i(t) = \frac{V_0}{R} (1 - \exp(-\frac{R}{L}t))$, the same result Heaviside achieved.²⁷

Now, the relationship between Heaviside's operational calculus and the Laplace transform cannot be one of simple identification, as Nahin suggests. Heaviside himself certainly didn't think so.²⁸ More importantly, significant differences arise in practice between the two techniques. For one thing, it is certainly not the case that we can simply substitute s for each instance of p . Note that even in this simple case, we cannot do this for Heaviside's operational solution ($C = \frac{e1}{(R+Lp)}$) and its equivalent in the s domain (in modern notation, $I(s) = \frac{V_0}{s(R+Ls)}$) as a result of the additional factor of $1/s$ in the latter.²⁹ Calculations involving the Laplace transform were also often more cumbersome than the corresponding calculations in the operational calculus. Harold Jeffreys, who harshly criticized Heaviside's lack of rigor, nonetheless wrote

[A]s a matter of practical convenience there can be no doubt that the operational method is far the best for dealing with the class of problems concerned. [...] [I]t is certain that in a very large class of cases the operational method will give the answer in a page when ordinary methods take five pages, and also that it gives

²⁷For a representative recent treatment of this example, see [Salivahanan et al. \[2000, pp. 157f\]](#).

²⁸In fact, he wrote to Bromwich, the first to rigorize the operational calculus via the inverse Laplace transform, "I never could stomach your complex integral method" (letter to Bromwich on 7 April, 1919, quoted in [Nahin \[2002, p. 230\]](#)). For a useful summary of Bromwich's [1928, 1916] approach, see [Lützen \[1979, pp. 176–180, 184–7\]](#).

Similarly, Heaviside's operational treatment of AC looks similar to Steinmetz's phasor method, the basis of the current treatment of impedance in terms of the complex plane, but Heaviside again resisted any such identification. In such cases, Heaviside interpreted p as ni , where n was angular velocity and i a differential operator that *behaved like* the imaginary unit. Heaviside explicitly contrasted this approach with approaches like Steinmetz's, which take the use of complex numbers seriously. In the latter case, one "[assumes] a complex form of solution at the beginning. It comes out complex at the end. [...] The algebra is that of the real imaginary." [[Heaviside, 1899, §284, p. 132](#)] In contrast, the i in Heaviside's $p = ni$ "must be finally interpreted correctly, as a differentiator, of course" [[Heaviside, 1899, §284, p. 132](#)]. Putting it more strongly later, he wrote, "if i be used at all, it is only a spurious imaginary" [[Heaviside, 1899, §437, p. 459](#)]. There is interesting historical and philosophical work to be done to make sense of Heaviside's claim not to be working with complex numbers in the same way as the likes of Steinmetz. Regrettably, space constraints mean I cannot discuss this further here. For a useful survey of early approaches to AC in terms of complex numbers and their relation to Steinmetz's phasor method, see [Kline \[1992, pp. 77ff\]](#).

²⁹What is more, the means of reaching these two equations will generally differ. See footnote 26.

the correct answer when ordinary methods, through human fallibility, are liable to give a wrong one. [Jeffreys, 1927, p. v]

But if we're only interested in a retrospective explanation of Heaviside's success, regardless of the epistemic status of his techniques at the time, no such identification is required. For example, in response to Wilson [2006], Pincock [2012, ch. 13] concedes that a scientist might not grasp concepts sufficient to pick out a determinate mathematical structure, but he suggests that we adopt a kind of semantic externalism according to which such scientists can be understood to articulate claims that go beyond the features of their concepts (both mathematical and physical) that they explicitly grasp. If so, we can understand Heaviside as unknowingly appealing to the sort of structure picked out with Laplace transform techniques. We then have a new explanation of Heaviside's success: Heaviside's mathematically mediated physical inferences succeeded because, unbeknownst to him, their mathematical part correctly characterizes structures picked out by the theory of Laplace transforms, and those structures stand in the right relationship to his target systems.

I am happy to concede that this is a perfectly good retrospective explanation of Heaviside's success, but I don't think it gives us reason to favor a mapping account. For one thing, retrospective explanations are also available in terms of RIC. Because RIC recovers mapping accounts as a special case, one option is to simply coopt the explanation just considered. But this is unnecessary. Why are Laplace transform techniques appropriate to do the work of the operational calculus? Because the inferences licensed by applications of the operational calculus, under the inference restriction strategies discussed so far, are a subset of those licensed by applications of Laplace transform techniques. And so the informational content of the former, treated in terms of RIC, is a subset of the informational content of the latter. Due to the rigor of Laplace transform techniques, we can be about as sure of their consistency and coherence as we can of any mathematical theory. As a result, experimental and theoretical agreement with representations involving Laplace transform methods confers a higher degree of epistemic support than similar agreement with representations appealing only to

the operational calculus. When we can't be so sure of a mathematical theory's consistency and coherence and must therefore adopt flexible inference restriction strategies, we have less assurance that any such agreement won't be undermined by the derivation of problematic results that necessitate further inferential restrictions. In this way, the relationship between Heaviside's techniques and subsequent rigorous ones retrospectively bolsters the epistemic standing of applications of the former.

Is there reason to favor the mapping-based explanation on the grounds that it further explains the informational content of representations using Laplace transforms in terms of a mathematical structure? I don't think so. RIC provides an alternative account of representations' informational content to explain how it licenses mathematically mediated inferences. We therefore don't need to appeal to later rigorous theories to explain how Heaviside's representations came to have informational content that justified his inferences. We only need to appeal to such theories to explain why these representations enjoyed a stronger epistemic status than he or his contemporaries could have appreciated.

Moreover, RIC has more to say about the relationship between applications of unrigorous techniques and their unrigorous counterparts than this sort of mapping account allows for. For example, different uses of the operational calculus bear remarkably different relationships to more rigorous mathematics, which must be understood in terms of different inferentially restrictive methodologies. Like Jeffreys, Bromwich recommended working with Heaviside's operational calculus rather than more rigorous mathematics as a matter of practical convenience. But he suggested an alternative inferentially restrictive methodology to Heaviside's "experimental" one: the operational calculus should not be used to derive any result that could not be derived independently via the method of Laplace transforms.³⁰ While Heaviside's infer-

³⁰Bromwich suggested such an approach in a letter to Heaviside:

After coming back to these questions after 2½ years of war-work, I found myself able to work more readily with operators than with complex-integration. [...] I at once saw that I must make the operator-method take the leading place: and complex-integrals have accordingly been pushed into footnotes. I still regard the complex-integral as a useful method for convincing the purest of pure mathematicians that the p -method rests on sound foundations: but I am sure that the p -method is the working-way of doing these things. [Letter to Heaviside on 5 April, 1919, quoted in

entially restrictive methodology may have been appropriate in the absence of more rigorous alternatives, Bromwich's approach seems entirely more appropriate once Laplace transform methods have been shown to do rigorously what Heaviside's methods could only do unrigorously. Once one can calculate inverse Laplace transforms via the Bromwich integral, one has a reliable, general means of checking results derived via the operational calculus, so that more ad hoc inferential restrictions serve little purpose.

Finally, even limiting our attention to "correlational pictures" in Wilson's sense, the cost of allowing ourselves only retrospective explanations is high. *Pace* Wilson, the explanation of Heaviside's success is not simply "because he was lucky" [Wilson, 2006, 528] to have picked out algebraic rules that both were useful and could be vindicated by later rigorous work, but rather that he took great care to calibrate his techniques to the physical problems in which he used them. As I argued in section 4, Heaviside's inference restriction strategies served to ensure that his representations had determinate accuracy conditions in restricted contexts and so could be evaluated in terms of their agreement with theoretical and experimental results. His ad hoc restrictions ensured that he could continually submit his results to theoretical and experimental scrutiny when he ventured out onto shakier ground. This certainly seems to allow him to tell a convincing story about how the usage of his operational techniques "matches to worldly support within normal circumstances of application" [Wilson, 2006, p. 516], albeit not one with all of the epistemic benefits of a more rigorous approach. An approach to supplying correlational pictures to justify Heaviside's inferences that is limited to retrospective explanations in terms of rigorous theories therefore misses an important part of the justification of Heaviside's practice.

Nahin [2002, p. 229]]

Bromwich was not rewarded for his kindness. Heaviside responded,

I rejoice to know that you have seen the simplicity and advantages of my way [...]. Now let the wooden headed rigorists go hang, and stick to differential operators and leave out the rigorous footnotes. It is easy enough if you don't stop to worry. [...] I never could stomach your complex integral method. [Letter to Bromwich, 7 April, 1919, quoted in Nahin [2002, pp. 229f]]

5.2.5 Conclusion

Central to how unrigorous mathematics can be successfully applied are the inferentially restrictive methodologies scientists use to manage the risks of working with an unrigorous theory. This is particularly clear in the case of Heaviside's operational calculus, which required him to make largely piecemeal inferential restrictions and to appeal, among other things, to the physical interpretation of the mathematics to determine how his mathematical tools ought to behave. In this section, I have argued that these practices are naturally represented in terms of RIC, but at best in a cumbersome and indirect way in terms of mapping accounts. As a result, RIC can be used to formulate better explanations of the success of physical inferences based on unrigorous mathematics—both why it was reasonable for Heaviside to adopt such techniques and why, in light of later developments, the results of these techniques were correct.

5.3 Path integrals in quantum physics

In all of the cases considered so far, the success of the unrigorous techniques of interest could ultimately be explained retrospectively in terms of more rigorous mathematics. The early calculus was succeeded by the modern calculus, which dispensed with infinitesimals, as well as rigorously defined infinitesimals, like those of non-standard analysis. Dirac's delta function could be understood in terms of Schwartz's theory of distributions. And Heaviside's operational techniques could be understood in terms of their relation to techniques based on the Laplace transform and its inverse, among others. I have already argued that this does not favor the mapping account. Such explanations don't help with the (at least) equally important question of why scientists should have seen themselves as justified in using those techniques before the existence of rigorous alternatives. And, as I argue in §5.2.4, RIC is a better tool even for reasoning about the relationships between unrigorous and more rigorous mathematical techniques that are crucial to retrospective explanations of the success of applications

of unrigorous mathematics.

But there is more to be said in response to the thought that the success of applications of unrigorous mathematics should ultimately be explained in terms of structures picked out by more rigorous alternatives. In particular, there is an important class of cases from present-day science in which such explanations are simply unavailable. Perhaps most striking are the mathematical tools underlying quantum field theory, which has made some of the most precisely confirmed predictions in the history of science, such as the value of the fine-structure constant α and the magnetic moment of the electron. Nonetheless, the various mathematical techniques used in quantum field theory are notoriously unrigorous, so much so that in most physically relevant cases there is no proof that a structure satisfying the relevant mathematics even exists. For instance, the existence of a non-trivial model for Yang-Mills theory on \mathbb{R}^4 with a mass gap is a (still unsolved) Millennium Prize problem. Since Yang-Mills is the basis of a significant part of the Standard Model of particle physics, this is striking indeed. If we are limited to explanations of the success unrigorous mathematics in terms of more rigorous mathematics, then we seem to be at a loss to explain some of the most important results in 20th- and 21st-century physics.

In this section, I consider one group of such techniques, those based on path integrals. Path integrals are a tool used in several areas of physics, most notably in quantum mechanics and quantum field theory, to make sense of the notion of a sum over a continuous space of functions. The introduction of the path integral allows for representations with a number of advantages. It allows quantum mechanics and quantum field theory to be presented directly in terms of the classical Lagrangian, yielding intuitive representations that are more manifestly compatible with special relativity, as well as allowing for an intuitive study of the semiclassical limit of quantum mechanics. In quantum field theory, perhaps the most important contribution of path integrals is to the derivation of the Feynman rules for perturbation series (the mathematics behind Feynman diagrams) in cases where this more difficult in terms of the canonical formalism (especially non-abelian gauge theories), but they have a number

of non-perturbative uses as well.

However, like the other cases in this chapter, these path integrals largely fall short of modern standards of mathematical rigor. In almost all cases, the relevant path integral cannot be understood straightforwardly as an honest integral with respect to a well-defined measure. A number of diverse strategies have been proposed to make these path integrals more rigorous, but these generally involve significant deviation from the naïve path integral formalism, losing much of the intuitiveness of this formalism in the process. Nonetheless, in many cases, the path integral formalism is used with little if any reference to these rigorization strategies. As one author puts it, “insofar as this fiction [expressed by the naïve path integral formalism] can be maintained it has many virtues [...]”. Nonetheless, it is not surprising that in some circumstances the formalism breaks down [...]” [Rivers, 1987, p. 123].

As in the cases of the Dirac delta function and Heaviside’s operational calculus, I argue that understanding representations appealing to the naïve path integral presents a challenge for mapping accounts. As in those cases, we cannot identify path integrals with the entities proposed to rigorize them, and while we can posit appropriate structures to make sense of such representations in terms of the mapping account, we don’t have a grip on what such structures are like independently of the inferential affordances of the naïve path integral formalism, which seem to do most of the explanatory heavy-lifting.

However, this example has a number of distinctive features, which present further challenges to the mapping account. Unlike in the cases of the Dirac delta function and Heaviside’s operational calculus, different inferential restrictions appear to be called for in different contexts. An apparent consequence of this is that no rigorization strategy for path integrals seems to cover all cases in which the path integral is used. Moreover, it is not always clear which inferential restrictions are called for in a given case, particularly when rigorous foundations are lacking. Nonetheless, these cases are held together by what seems to be a single, albeit mathematically incoherent, path integral concept expressed by the naïve path integral formalism. So, in addition to the challenge of making sense of particular applications of path integrals,

there is the challenge of making sense of how *different* applications of this formalism are related in this way, though they seem to call for quite different structures (reflected in different rigorization strategies) if represented in terms of the mapping account.

5.3.1 Path integrals in quantum mechanics

As it was originally used in quantum mechanics, the Feynman path integral serves to specify the probability amplitude for a particle to go from one spacetime point to another. The core idea of the path integral formulation of quantum mechanics is that each continuous path between these points makes an equal contribution—though with different phases corresponding to the classical action of the path—to this overall probability amplitude. Feynman & Hibbs [1965, p. 29] give the rough, qualitative idea in the following way.³¹ Let $K(b, a)$ be the probability amplitude of the particle moving from point $x_a \in \mathbb{R}^d$ at t_a to point $x_b \in \mathbb{R}^d$ at t_b , and ϕ a functional taking each path $x : [t_a, t_b] \rightarrow \mathbb{R}^d$ from a to b to its complex contribution to the probability amplitude. We then have something like

$$K(b, a) = \sum_{\text{paths } x(t) \text{ from } a \text{ to } b} \phi[x(t)]. \quad (5.5)$$

The phase of each path's contribution to the overall probability amplitude is given by the classical action $S[x(t)] = \int_{t_a}^{t_b} L(\dot{x}, x, t) dt$ (where L is the Lagrangian of the system) for the path in units of \hbar , so this becomes

$$K(b, a) = \sum_{\text{paths } x(t) \text{ from } a \text{ to } b} \text{const } e^{i\hbar^{-1}S[x(t)]} \quad (5.6)$$

with const a normalizing constant. Of course, more needs to be done to make sense of sums over continuum-many paths. The natural way to do so—and the route Feynman followed in developing this approach to quantum mechanics—is to treat the sum as (if it were) an

³¹In the following discussion I largely follow the presentation in chapter 2 of [Feynman & Hibbs, 1965], but for the sake of generality I consider the case of a system with d spatial dimensions, rather than the one-dimensional system considered there.

integral supported by a suitable measure on a suitable space of paths. This is the Feynman path integral.

Understood in this way, the rough, qualitative idea expressed in (5.5) and (5.6) becomes

$$K(b, a) = \int_{C_{x_a, x_b}^{t_a, t_b}} e^{i\hbar^{-1}S[x(t)]} \mathcal{D}x(t) \quad (5.7)$$

with $\mathcal{D}x(t)$ understood as a measure on the space of paths that weights all paths equally and $C_{x_a, x_b}^{t_a, t_b}$ the space of \mathbb{R}^d -valued continuous functions $x(t)$ such that $x(t_a) = x_a$ and $x(t_b) = x_b$. Intuitively, this captures quantum behavior in the following way. Due to the factor of i in the exponential, the integrand in (5.7) will oscillate in the complex plane, resulting in constructive and destructive interference between the contributions of different paths—thus explaining, e.g., the interference patterns observed in the double-slit experiment. This oscillatory behavior also helps provide an intuitive picture of the semiclassical limit of quantum mechanics. As \hbar becomes very small relative to S , this behavior becomes more extreme, so that small changes to a path tend to yield large changes in phase. Paths for which this is the case will tend not to contribute much to the overall “sum” due to destructive interference with nearby paths. However, small changes to the path that minimizes S will produce only a small change in S . Since the contributions of paths very close to this path are nearly in phase, the main contribution to the integral as a whole should come from paths near the path that minimizes the classical action—i.e., the classical trajectory. This approach also has the advantage of being more manifestly compatible with special relativity than standard, Hamiltonian-based presentations of quantum mechanics, given the Lorentz covariance of the Lagrangian (though doing this right generally requires a move from quantum mechanics to quantum field theory).

Feynman & Hibbs [1965] give a preliminary “definition” of this sum by analogy to the Riemann integral. According to this definition, we consider an approximation of the set of paths in terms of piecewise-linear paths. We get such an approximation by partitioning the time interval from t_a to t_b into n steps of width $\epsilon = (t_b - t_a)/n$. Let $t_0 = t_a$, $t_{i+1} = t_i + \epsilon$, and so

$t_n = t_b$. Our paths $x(t)$ are those that can be produced by choosing some value x_i for each t_i and connecting these points with straight lines. Then we can treat the sum over such paths as $n - 1$ d -dimensional integrals over all values for x_i from 0 to n so that

$$K(b, a) \sim \int_{\mathbb{R}^d} \dots \int_{\mathbb{R}^d} e^{i\hbar^{-1}S[x(t)]} dx_1 \dots dx_{n-1} \quad (5.8)$$

up to a constant factor. We then get the integral over all continuous paths by introducing a normalization factor C and taking the limit as $n \rightarrow \infty$:

$$K(b, a) = \lim_{n \rightarrow \infty} C \int_{\mathbb{R}^d} \dots \int_{\mathbb{R}^d} e^{i\hbar^{-1}S[x(t)]} dx_1 \dots dx_{n-1}. \quad (5.9)$$

The value of C will depend on the system represented and will be a function of ϵ (or, equivalently, n). C intuitively plays the role of the partition width Δx_i of an ordinary, finite-dimensional Riemann sum $\sum_{i=0}^{n-1} f(x_i)\Delta x_i$, except that, unlike in the case of Δx_i in the ordinary Riemann sum, C tends to be infinite in the limit $n \rightarrow \infty$. For instance, for a system whose Lagrangian is $L = \frac{m}{2}\dot{x}^2 - V(x, t)$, $C = \left(\frac{m}{2\pi i\hbar\epsilon}\right)^{nd/2}$ is the correct normalization factor.

We already have a problem. Feynman & Hibbs [1965] note that this sort of definition won't work in all cases. For instance, they point out that if the Lagrangian (and so the action) depends on \ddot{x} , then problems may arise because \dot{x} will almost always be discontinuous at (x_i, t_i) for a given piecewise-linear path $x(t)$ and time t_i , and so \ddot{x} will be infinite at those points.³² And so they allow that other definitions of the path integral may be required in such cases. They also provide no general way to determine the normalization factor C , determining what C should be only for particular Lagrangians. As they put it, "to define such a normalizing factor seems to be a very difficult problem and we do not know how to do it in general" (p. 33). However, they downplay the importance of the lack of a fully general definition, pointing to the wide variety of integrals—even when we restrict our attention to integrals over the reals—

³²For related reasons, more recent presentations making use of similar definitions typically explicitly introduce a time-slice approximation of the action as well [see, e.g., Rivers, 1987, p. 122]. Of course, this still doesn't address the closely related problem, considered below, that typical paths in the infinite limit are *nowhere* differentiable.

required in pure mathematics. In cases where we cannot use the Riemann integral, we may need to appeal to the Lebesgue integral, and in cases where we cannot use the Lebesgue integral, we may need to appeal to yet other definitions. This need for a variety of definitions of integral that do not hold in full generality does not, according to Feynman and Hibbs, mean that there is no core concept of an integral common to each case, and this should hold as much for the path integral as for the finite-dimensional integral. So, while other definitions may be required to make sense of the path integral in certain cases, Feynman and Hibbs use the more general notation $\int \phi[x(t)]\mathcal{D}x(t)$ to pick out a general path integral concept common to these different definitions.

However, the situation regarding the mathematical rigor of the path integral is worse than Feynman and Hibbs make it out to be. Except in special cases, there is *no* rigorous definition of the path integral that allows us both to understand it as an honest integral supported by a well-defined measure on the relevant space of paths and to use it for its intended purposes.

One natural interpretation of the naïve path integral formalism as reflected in (5.7) is that $\mathcal{D}x(t)$ is a reference measure with respect to which the integral is defined, analogous to the Jordan measure for finite-dimensional Riemann integrals or the Lebesgue measure for finite-dimensional Lebesgue integrals. For instance, since the right-hand side of (5.9) is the limit $n \rightarrow \infty$ of $n - 1$ d -dimensional integrals and so equivalent to the infinite limit of an $d(n - 1)$ -dimensional integral, we might expect $\mathcal{D}x(t)$ to denote an extension of, say, the Lebesgue measure supporting the finite dimensional integrals—i.e., the measure given by $dx_1 \dots dx_{n-1}$ —to an infinite-dimensional Lebesgue-type measure on the space $C_{x_a, x_b}^{t_a, t_b}$ of paths $x(t)$. But, given reasonable restrictions on what such a measure should be like (e.g., σ -finite, translation-invariant, countably additive), there is no such measure.³³ On the other hand, we might

³³See, e.g., Hall [2013, p. 446] and Mazzucchi [2009, p. 5]. The prospects seem to be better if we relax these restrictions somewhat, so that, in particular, the “measure” is allowed to be a generalized measure (i.e., a distribution). For instance, Montaldi & Smolyanov [2017] show that one can use translation invariant generalized measures on a locally convex vector space to define the path integral in (5.7) in a fairly natural way, with $\mathcal{D}x(t)$ interpreted as one of these generalized measures. They claim—rightly, I think—that this is closer to what Feynman and others originally did with the path integral than other strategies for providing it with rigorous foundations. However, just as Schwartz’s work capturing the behavior of the Dirac delta function in terms of distributions doesn’t show that the Dirac delta *really* is a distribution (as used by Dirac), Montaldi and Smolyanov’s work

think that we could understand not $\mathcal{D}x(t)$ but the whole expression $Ce^{i\hbar^{-1}S[x(t)]}\mathcal{D}x(t)$ as a complex-valued measure on the space of paths. But again this cannot generally be true. Cameron [1960] showed that such a measure must have infinite total variation and so is not suitable to support the path integral.³⁴

We have yet more trouble taking the path integral formula (5.7) literally when we consider what the space $C_{x_a, x_b}^{t_a, t_b}$ must be like. Such paths will be rough in a way incompatible with a natural understanding of $S[x(t)]$. In fact, the paths that are differentiable at at least one point in $[t_a, t_b]$ can be shown to form a set of Wiener measure zero; almost all (in the sense of the Wiener measure) paths in $C_{x_a, x_b}^{t_a, t_b}$ are nowhere differentiable (see Johnson & Lapidus, 2000, pp. 2, 102; Rivers, 1987, pp. 114–9). But $S[x(t)]$ depends on $\dot{x}(t)$, since it is just the time integral of the Lagrangian, and so S is singular for many paths in $C_{x_a, x_b}^{t_a, t_b}$. A side effect of this is that in the special case where the Feynman path integral can be defined in terms of the Wiener measure (a move that requires deviating from the original path integral formula by moving from real to imaginary time), only nowhere differentiable paths (and so paths with infinite action) can contribute to the integral!³⁵

The situation is better if we take the path integral formula (5.7) to be a shorthand notation for a limit of finite-dimensional integrals as in (5.9), but serious problems remain. In particular, if S is real-valued, as it is in most problems, even the finite-dimensional integrals $\int_{\mathbb{R}^d} \dots \int_{\mathbb{R}^d} e^{i\hbar^{-1}S[x(t)]} dx_1 \dots dx_{n-1}$ do not converge, as the integrand then has a constant absolute value of one. But even if this is resolved, say, by shifting to imaginary time (thereby making S complex-valued)³⁶, it is not clear that the limit itself exists. [Johnson & Lapidus,

doesn't show that the Feynman path integral really is an integral with respect to one of their Lebesgue-Feynman generalized measures. (It also remains to be shown that such an integral has the right properties to serve as the Feynman path integral more generally, though Montaldi and Smolyanov do show that it suffices to derive analogues of some of the results in Feynman's original paper on the path integral approach to quantum mechanics [Feynman, 1948].)

³⁴Cf. Albeverio *et al.* [2008, p. 4], Mazzucchi [2009, pp. 8–10], and Hall [2013, p. 447].

³⁵Rivers [1987, pp. 109–19] shows this for the one-dimensional case, and there is no reason to think that things get better as we move to more dimensions. What is worse, an analogous result holds even if we follow Feynman and Hibbs in defining the path integral as a limit of finite-dimensional integrals [Rivers, 1987, pp. 121–7].

³⁶A similar trick frequently used by Feynman was to add an infinitesimal imaginary component to the time to make it more plausible that the limit exists, where this infinitesimal is closer to those considered in chapter 4 than to the infinitesimals of non-standard analysis!

2000, p. 110]

In light of this, it looks like any attempt to put the path integral on a rigorous footing must deviate significantly from what we find in expressions like (5.7). Nonetheless, typically it is unrigorous expressions like (5.7) and (5.9), rather than these rigorous replacements, that have been used in path-integral-based, non-relativistic quantum mechanics. In a sense, this is to be expected. After all, as we saw, Dirac and others continued to use the delta function even after the work of Schwartz, and Heaviside continued to prefer his operational calculus to the rigorous techniques introduced by Bromwich and others, largely because the unrigorous techniques were more computationally efficient.

What is striking, however, is that the pull of the unrigorous path integral approach to quantum mechanics, in contrast to these other unrigorous devices, was (and is) largely *conceptual*.³⁷ For the most part, it doesn't greatly simplify calculations in the quantum mechanics of simple non-relativistic systems—in contrast to the path integral approach to quantum field theory, as we'll see shortly. Instead, it provides a formulation of quantum mechanics that produces the same results as the canonical formulation, but in a way that can (perhaps) be more intuitively motivated, that connects the concepts of classical and quantum mechanics (including reinstating concepts, like that of trajectory, that had been banished from quantum mechanics previously), that is more compatible conceptually with relativity, and so on.³⁸

This fact suggests a useful way to understand the use of the unrigorous, quantum mechanical Feynman path integral in terms of an inferentially restrictive methodology quite different from those at work in the use of the Dirac delta function and Heaviside's operational calculus. Since the path integral formulation of non-relativistic quantum mechanics is meant to be equivalent to the canonical, operator-based formulation, we can use the latter to place con-

³⁷In a sense, we might think this is true of Heaviside's attitudes toward his operational calculus as well, in that Heaviside saw his unrigorous techniques and proofs as promoting a better understanding than more rigorous techniques, as we saw earlier. But this kind of conceptual advantage seems to be grounded in computational efficiency and is not what I have in mind here.

³⁸Such advantages were first described by Dirac [1933] and motivated Feynman's work on the path integral approach in his PhD thesis [Feynman, 2005]. More recently, the physicist Fay Dowker has been a strong proponent of the approach [Dowker, 2012, Dowker *et al.*, 2010].

straints on the inferences we can make on the basis of the unrigorous path integral. There are a couple of ways in which we might try to do this.

First, following Davey [2003, pp. 450f], we might look to how Feynman (Feynman, 1948, §§5–6, pp. 372–7; Feynman & Hibbs, 1965, §4.1, pp. 76–84) derives the Schrödinger equation from the path integral formulation. As Davey notes, for this derivation to work, it is necessary only that the path integral share certain properties with ordinary, finite-dimensional integrals, as well as the property $K(b, a) = \int_{\mathbb{R}^d} K(b, c)K(c, a)dx_c$. These properties on their own would certainly seem to be consistent, entailing only that the path integral *behaves like* a finite-dimensional integral in certain ways, not that the path integral *is* in fact an integral supported by a well-defined, well-behaved measure. And since they suffice to derive the Schrödinger equation, they suffice to do non-relativistic quantum mechanics.

However, this approach is too restrictive, as it does not allow for some of the conceptual advantages of the path integral formulation to be exploited. Despite the equivalence between the path integral and canonical formulations, the path integral formulation does go beyond the canonical formulation particularly where classical concepts are concerned, something Davey’s proposed strategy does not address. For example, Davey’s inferential restrictions would seem to rule out the argument I sketched earlier concerning the path integral approach to the semiclassical limit of quantum mechanics. A less hand-wavy version of that argument—and taking advantage of the path integral as a tool to study the semiclassical limit of quantum mechanics more generally—requires that the path integral share important properties with more specific finite-dimensional integrals, viz. those properties required to make the stationary phase approximation work out. This requires us to allow more inference patterns involving the path integral than those at work in deriving the Schrödinger equation. And as we begin ascribing more and more properties to the path integral to enable us to take advantage of the conceptual resources of the path integral formulation of quantum mechanics in this way, it becomes less clear that these properties are in fact consistent—and less clear still that the equivalence with the canonical formulation is preserved. Perhaps this is partly

what Davey has in mind when he writes, “it is unclear *exactly* how literally one can take the idea of a path integral without getting into trouble. [...] If the physicist wants to go beyond the safest ways of treating the path integral, he must make conjectures about consistency which he may well be forced to withdraw at some later point. [...] Thus, in the case of the path integral, the precise bounds of the physicist’s inferential restrictiveness are not entirely clear.” [Davey, 2003, p. 451, emphasis in original]

However, we have another way to make sense of the inferential restrictions involved in using the Feynman path integral in non-relativistic quantum mechanics by appealing to the equivalence between the path integral and canonical formulations. In brief, the idea is to get as much mileage out of the analogy between the Feynman path integral and the relevant finite-dimensional integrals as one can without losing the equivalence between the path integral and canonical formulations of quantum mechanics. One is allowed to reason as if the Feynman path integral were a genuine integral, except when that reasoning leads to a result expressible in the canonical formalism that disagrees with the canonical formulation. This, I think, much better accounts for the more opportunistic uses of the path integral in quantum mechanics.

Now we come to the question of how to make sense of all of this in terms of the mapping account. As in the case of the Dirac delta function and Heaviside’s operational calculus, it is clear what *not* to do. We cannot simply identify the Feynman path integral with one of its rigorous replacements and understand applications of it in quantum mechanics in terms of the structures picked out in this way. In fact, this is clearer in the case of the path integral than in these earlier cases, as none of these rigorization strategies seem to succeed in making sense of the whole range of techniques used in the path integral formulation of quantum mechanics. For instance, Johnson & Lapidus [2000] show that several strategies to provide the path integral with a rigorous foundation have equivalent results and suffice to justify several core uses of the path integral, but these aren’t enough to underwrite the use of the path integral to study the semiclassical limit of quantum mechanics. While similar arguments can be given a rigorous treatment in finite dimensions, Johnson and Lapidus write of the

infinite-dimensional case only that it “has been extremely difficult to prove and will not be discussed rigorously in this book” (p. 106).

On the other hand, given the equivalence between the path integral and canonical formulations of quantum mechanics, we might think that the appropriate structure is not picked out by a rigorization of the path integral, but instead picked out by (the mathematical apparatus of) the canonical formulation. But this too would be a mistake for two reasons. First, as we saw earlier, the equivalence between these two formulations of quantum mechanics puts constraints on the behavior of the path integral, but the path integral formulation nonetheless goes beyond the canonical formulation in significant ways by making use of different concepts (particularly those of classical mechanics) and related methods of argument (for instance, the study of the semiclassical limit of quantum mechanics via the method of stationary phase). Indeed, this seems to be the whole point of using the path integral approach to quantum mechanics in the first place. Representing such applications of the path integral in terms of a structure picked out by the canonical formulation will fail to account for these deviations from the canonical formulation, particularly the use of classical concepts not present in the canonical formulation. Second, even if the path integral were just a convenient computational tool for deriving the same results as the canonical formulation and taken to have no conceptual import whatsoever, representing applications of the path integral in terms of a structure picked out by the canonical formulation would still fail to account for the distinctive patterns of reasoning involved in these computations.

So again we find ourselves in a situation where the right structures seem to be just *whatever* structures support the inferences allowed by physicists’ inferentially restrictive methodology. As in the earlier cases, we don’t have a good grip on what such a structure is like independently of these permissible inferences. So yet again, RIC seems to represent much more straightforwardly the philosophically relevant features of the relevant practices.

5.3.2 The path integral beyond quantum mechanics

Most problems of quantum field theory can be thought of as problems of finding a correct definition and a computation method for some Feynman path integral. From a mathematician's viewpoint almost every such computation is in fact a half-baked and *ad hoc* definition, but a readiness to work heuristically with such *a priori* undefined expressions [...] is necessary in this domain. [Manin, 1989, p. 234]

Things get more interesting as we move from the case of the Feynman path integral in quantum mechanics to the use of related path integrals in other areas, where the foundations of such approaches are often even less clear. In such cases, often the naïve fiction of the path integral as analogous to a finite-dimensional integral is particularly fruitful, enabling insights apparently at odds with efforts to rigorize the path integral. Here one needs to venture out onto shakier ground, as one often cannot rely on an equivalent but more rigorous theory to provide inferential guardrails. And so we need a different story about the required inferential restrictions than we had in the case of non-relativistic quantum mechanics.

This often happens when the path integral is connected to other mathematical considerations. For instance, there is a fair amount of very interesting work connecting the Feynman path integral to topology, which is then applied in physics. A classic example, discussed by Schulman [1988, p. 13; 2005, chs. 23–4], is the path integral formula for the propagator for a particle moving on the group manifold for $SO(3)$, the group of rotations of three-dimensional Euclidean space about the origin under composition. Paths on this space naturally form two homotopy classes (classes of paths that can be continuously deformed into one another). If we take the analogy between the path integral and ordinary finite-dimensional integrals seriously, it is natural to think that we can break up the path integral for this propagator into two integrals, one over each of these homotopy classes. This has the advantage of allowing us to think about the relative phases of these two parts of the integrand, which in turn makes

available an elegant model of particles of integer or half-integer spin. But the moves required to make sense of this elegant representation are at odds with more rigorous ways of understanding the path integral, especially in terms of Feynman’s definition of the path integral as the infinite limit of a finite-dimensional integral. Schulman puts it nicely when he writes, “I do not know how to justify this breakup rigorously, since the ‘sum’ is in the end only the limit of a multiple integral, but the ease with which the path sum language leads to a profound property of nature suggests that some underlying truth has been found” [Schulman, 1988, p. 13]. The physical fecundity of the naïve path integral formalism is very much at odds with the attempts to provide the formalism with a rigorous mathematical interpretation.

This is even more dramatic when we consider how the path integral for non-relativistic quantum mechanics was generalized to be used in quantum field theory. At the level of the formalism, it is easy to understand how the path integral is thus extended; instead of integrating over paths construed as functions of time alone, we can integrate over paths construed as classical field configurations and define the action S for these field configurations via the Lagrangian for the classical field theory to be quantized. At the level of rigor, however, this jump from integration over a space of particle trajectories to integration over a space of field configurations is a very significant one. When we make the move to the path integrals for quantum field theory, even some of the simplest cases are significantly more difficult to put on a rigorous foundation.

For example, consider the simplest case, the path integral for real scalar fields. This just involves integrating over a space of functions of more than one variable, so that we have a very close analogue of the path integral formula (5.7) for the propagator of a non-relativistic particle:

$$K(b, a) = \int_{\mathcal{F}_d} e^{i\hbar^{-1}S[\phi]} \mathcal{D}\phi \quad (5.10)$$

with \mathcal{F}_d the space of fields ϕ in d dimensions satisfying the appropriate boundary conditions. For the sake of concreteness, consider the special case of a ϕ^4 field theory—i.e., with self-

interactions requiring a Lagrangian with a term proportional to ϕ^4 . This is one of the better cases in quantum field theory with respect to mathematical rigor, but understanding it in a rigorous way still involves significant deviation from what is suggested by the formalism. This first requires performing a Wick rotation from real time to imaginary time (corresponding to a move from Minkowski space to Euclidean space), yielding

$$K(b, a) = \int_{\mathcal{F}_d} e^{-\hbar^{-1}S_E[\phi]} \mathcal{D}\phi. \quad (5.11)$$

Then, in a sense, it is possible for all $d \geq 1$ to define a Gaussian measure on \mathcal{F}_d with respect to which the integral in (5.11) is defined. However, for $d \geq 2$, this measure is not supported on any space of functions, but only on spaces of distributions. (Glimm & Jaffe, 1981, §8.5, pp. 152ff; Hall, 2013, pp. 451f) This result is analogous to the result for the path integral in non-relativistic quantum mechanics that almost all (in the sense of Wiener measure) paths in $C_{x_a, x_b}^{t_a, t_b}$ (for one spatial dimension) are nowhere differentiable, but worse: the relevant paths are so rough that they cannot even be functions! And this issue arises *after* we've already made the simplifying move from Minkowski to Euclidean space-time.

Despite being generally more difficult to put on a rigorous footing, path integrals have had a far greater impact in quantum field theory than in non-relativistic quantum mechanics. Perhaps most importantly, unrigorous path integrals were used to derive the Feynman rules for perturbation series in a number of quantum field theories before this could be done via the canonical formalism. Feynman [1949] used his path integrals in part to derive the rules for Feynman diagrams for quantum electrodynamics, though other physicists tended to derive these rules using the canonical operator formalism [Weinberg, 1995, p. 376]. Later, path integrals were shown to provide a method for deriving the Feynman rules for the much more difficult cases of non-Abelian gauge theories (DeWitt, 1964; Faddeev & Popov, 1967) and spontaneously broken gauge theories [’t Hooft, 1971] well before this could be done in terms of the canonical formalism.

In such cases, we need to look beyond the strategies for inferential restriction considered in the previous section. The equivalence between the path integral and canonical formalisms could not be used to provide inferential restrictions at the time DeWitt, Faddeev, Popov, and 't Hooft did this work because such work had not yet been done with the canonical formalism. This leaves us with a couple of options.

First, as Davey [2003] suggests, we might treat the path integral methods in, say, Faddeev & Popov [1967] as providing an efficient algorithm for generating perturbation series for non-Abelian gauge theories. In that case, path-integral-based inferences made outside the context of applying this algorithm would be disallowed. Such restrictions do seem to be frequently observed within the context of perturbative quantum field theory. For instance, Rivers [1987] often draws attention to how much work can be done on the assumption that the path integral is merely an algorithm for generating perturbation series (e.g., p. 75). Nguyen [2016] goes as far as to present this as a sort of rigorization strategy, justifying unrigorous manipulations of path integrals in terms of the properties of these perturbation series (and clarifying some of the analogies and disanalogies between perturbative path integration and ordinary, finite-dimensional integration in the process).

On the other hand, we might want to take these path integrals to be more than algorithms for producing perturbation series for a few reasons. First, the point of a perturbation series is to approximate the results of a fuller theory. This suggests an understanding of the path integrals themselves as the objects the relevant perturbation series approximate. Second, as Rivers [1987, p. 81] observes, the analogy between path integrals and ordinary finite-dimensional integrals suggests other ways of producing useful (and tractable) approximations that cannot be understood without relaxing the inferential restrictions described above—for instance, series expansions in \hbar rather than the coupling strength, which cannot always be understood as rearrangements of the perturbation series produced by the Feynman rules. Third, there are non-perturbative effects (i.e., that cannot be captured by perturbation theory at any finite order) that are usefully represented via the path integral. For instance, color confinement can

usefully be represented in terms of the path integral approach to quantum chromodynamics, with approximate calculations made by moving to a discrete representation (lattice QCD) in which the path integrals become finite-dimensional and can be approximated via Monte Carlo simulations.³⁹

In these cases, it is necessary to venture out onto shakier ground. Davey puts it very nicely when he writes

We know that the path integral may be used unproblematically to generate perturbation theory, and we know that the path integral cannot be taken completely literally as an integral supported by a well-defined underlying measure. There are, however, many “degrees of seriousness” between these extremes for which it is generally *not* clear whether one is on safe ground or not. If the physicist wants to go beyond the safest ways of treating the path integral, he must make conjectures about consistency which he may well be forced to withdraw at some later point. [Davey, 2003, p. 451]

For the cases considered in the previous paragraph, what is needed is a more general path integral concept to serve as the object that these less objectionable mathematical techniques—perturbation series, other series-based approximations including saddle-point/stationary-phase approximations, and fully discretized path integrals—serve to approximate. It therefore needs to allow for a more extensive set of inferential possibilities that allow it, among other things, to provide inferential connections between these various approximate techniques with their more limited inferential possibilities. Such an understanding of the path integral seems to be required to make sense of the contributions of these various techniques to our more unified quantum-field-theory-based representations, including the Standard Model. We cannot identify the path integrals in, say, the formulation of the Standard Model with any of these more restricted, approximate path integral concepts, as each is best suited to capture only

³⁹This was first discovered by Wilson [1974] and has since become an important part of the toolkit of quantum field theory (See, e.g., the introductory textbook [Gattringer & Lang, 2010]).

some of the phenomena that the Standard Model is supposed to explain. But the precise behavior required of such a path integral concept for this purpose is far from clear, and thus so too are the necessary inferential restrictions. These must be tentative and subject to revision.

Turning to how to understand all this in terms of the mapping account, it is clear that the problem that arises for all of the previous cases arises again here. The right structures for representing these cases in terms of the mapping account seem to be just those that make the inferences allowed by physicists' inferentially restrictive methodology, whatever that may be in the given case, come out truth-preserving. And we don't really have any understanding of what such a structure might be independently of these inferences. But this case adds some important wrinkles.

First, these inferential restrictions are tentative in a way that the other inferentially restrictive methodologies considered in this chapter—for the early calculus, the Dirac delta function, Heaviside's operational calculus, and even the path integral in the path integral formulation of non-relativistic quantum mechanics—are not. This presents us with a new problem: how do we make sense of the use of the path integral when such a tentative inferentially restrictive methodology is adopted? The most natural way seems to be in terms of a range of different representations involving different inferential restrictions. And the most natural way to explain the relationships among these different representations is in terms of their restricting in various ways a shared collection of inference patterns suggested by the path integral notation and the formal analogies between it and finite-dimensional integrals. We might then represent these inferential restrictions in terms of collections of partial structures, say, but it is not clear what is to be gained by doing so in cases like this, since, as before, our grasp of what such structures are like is entirely dependent on our understanding of these inferential restrictions.

Second, it would be nice to be able to explain the relationships between different path integral concepts requiring different inferential restrictions, especially when they are used in the context of a single representation, like the Standard Model. And the most natural way

to do this is also at the level of the inferential affordances of the path integral formalism rather than at the level of mathematical structure. Here is one possible way to go. What the various approximation strategies appealing to path integration considered above have in common is a basis in some formal analogy between the naïve path integral formalism and ordinary, finite-dimensional integrals. Perturbative uses of path integrals take path integrals to be analogous to ordinary integrals in that they bear a certain relationship to certain infinite series—viz., that a certain way of constructing series produces series that are asymptotic to the integral. Other uses of path integrals exploit analogies with ordinary oscillatory integrals, allowing for an analogue of the stationary phase method. Yet other uses exploit the analogy between such integrals and finite dimensional integrals in a discrete space-time. Each of these analogies puts constraints on whatever the “full” path integral might be. As [Manin \[1981, p. 93\]](#) writes, albeit in a different context, “From the viewpoint of a mathematician, each such calculation of a Feynman integral simultaneously defines what is calculated, i.e. it constructs a text in a formal language whose grammar has not previously been described.” But these are just formal analogies, analogies in inferential behavior, and explanations of them in terms of underlying structure have not generally been given. And so they are most straightforwardly understood in terms of the inferential affordances of this formalism. Again, while we can in a sense come up with structures that do the job, these will again be *whatever* structures support the target inferences. And again, we don’t have a grasp on such structures independently of these inferences. (After all, even showing the existence of structures for Yang-Mills in four dimensions is an open problem.)

5.3.3 Conclusion

In this section, I’ve argued that the case of path integrals in quantum mechanics and quantum field theory gives us further reason to favor RIC over mapping accounts. Unlike other cases considered so far, it is not (yet) possible to explain the success of path integral techniques in terms of more rigorous mathematics, which picks out structures that can be used by mapping

5.4 Conclusions: The scope of an account of mathematical scientific representation

accounts. Further, these cases illustrate the diversity of the inferentially restrictive methodologies that may be called for in applications of unrigorous mathematics, even concerning what appears to be a single mathematical concept (viz., the path integral), as well as how tentative such methodologies must be when the limits of the relevant unrigorous concepts are not well understood. This means, among other things, that Heaviside was in good company (or at least was not as much of an outlier as it might seem) when he proposed similarly tentative methodologies in relation to his operational calculus. Because RIC again provides a much more straightforward representation of these philosophically salient aspects of scientific practice, we have further reason to favor it over mapping accounts.

5.4 Conclusions: The scope of an account of mathematical scientific representation

In this chapter, I have argued that applications of unrigorous mathematics give us good reason to favor RIC over mapping accounts on the grounds that RIC more perspicuously represents the philosophically salient aspects of the relevant practices. By way of conclusion, I now consider two objections a proponent of a mapping account might have to the arguments I have presented here, each of which concerns the appropriate scope of an account of mathematical scientific representation. I briefly argue that these illustrate advantages rather than liabilities of RIC.

5.4.1 Formalisms and inferentially restrictive methodologies

One possible response is that it is no objection to the mapping account that the functional aspects of this or that scientific practice cannot be directly captured in terms of structures and mappings. Structures and mappings are meant to capture the informational content of scientific representations. It is open to the mapping account to incorporate further information to explain some of the functional aspects of these practices in terms of that informational con-

5.4 Conclusions: The scope of an account of mathematical scientific representation

tent. Indeed, this seems to be exactly what is called for and exactly how mapping accounts are intended to be used in their role as meta-level devices for representing philosophically salient aspects of scientific practice.

In particular, we might think that the cases discussed in this and the previous chapter show not that we should reject the mapping account in favor of RIC, but that we should supplement the mapping account with a further account of mathematical formalisms as the means by which human beings pick out and reason about mathematical structures. After all, in each of the cases considered before, we see scientists reasoning using mathematical formalisms without sufficient means to pick out associated structures. But that doesn't rule out the possibility of philosophers' using associated structures to explain why those formalisms were appropriate for reasoning about the relevant target systems.

In a recent paper, [Vincent *et al.* \[2018\]](#) argue that mapping accounts need to be supplemented in just such a way. In their view, a major shortcoming of mapping accounts—especially those, like the inferential conception of [Bueno & Colyvan \[2011\]](#), that emphasize the inferential role of mathematics in science—is that notions like “mathematical resources” and “inferential power” remain unanalyzed. Cashing out such notions requires not just looking to mathematical structures, but also the formalisms through which they are presented. And these cannot simply be “two sides of [the] same inferential coin” (p. 4) because the same structure may be presented in terms of quite different formalisms, which in turn open up different inferential possibilities. For mathematical structures to be useful in scientific applications, it is necessary to have a formalism that allows one to gain information about the structure that is useful relative to one's epistemic goals. Thus, mapping accounts, including those that emphasize the inferential role of mathematics, must be supplemented with an account of the role of things like formalisms in order to explain what provides mathematical models with their inferential power.

I don't dispute that mapping accounts can (and should) be supplemented in precisely this way. Certainly, we shouldn't expect an account of mathematical scientific representation to

5.4 Conclusions: The scope of an account of mathematical scientific representation

directly represent functional aspects of such practices in general. Imagine an account that formally represented scientists' intentions and other psychological states relevant to their mathematically mediated reasoning. Such an account might capture something interesting about scientific practice, but proponents of it couldn't rightly criticize other accounts on the grounds that they didn't represent scientists' psychological states. Scientists' psychological states are simply outside the scope of those accounts.

However, in their role as meta-level devices for representing philosophically salient aspects of mathematical practice, mapping accounts are typically intended to explain a number of functional aspects of scientific practice, especially mathematically mediated surrogative inference, in terms of the informational content of mathematical scientific representations. RIC is intended to do the same. I take the cases in this and the previous chapter to favor RIC primarily because RIC represents their informational content in a way that makes it easier to reason about other philosophically important aspects of scientific practice, especially those relevant to understanding the epistemic credentials of applications of unrigorous mathematics both from the perspective of scientists at the time and retrospectively. In other words, the problem with mapping accounts is not that they don't directly represent functional aspects of scientific practice in their account of the informational content of mathematical scientific representations, but rather that their account of the informational content of these representations is less useful than an alternative account (RIC) as a meta-level tool for reasoning about these functional aspects.

For concreteness, consider how [Bueno & Colyvan \[2011\]](#) would explain the epistemic credentials of a mathematical formalism. Mathematically mediated inferences about a physical target system can be reasonably made provided that mappings establish structural relations appropriate to support those inferences. The use of a particular mathematical formalism to make such inferences is then justified when it is used to explore the properties of a structure that is connected to the target structure via appropriate mappings.

The problem with this in the cases considered here is that, in each of these cases, the relev-

5.4 Conclusions: The scope of an account of mathematical scientific representation

ant formalism seems to float free of any particular mathematical structure. This is quite clear in the path integral case, but it is also present in the other cases I've discussed so far. The Dirac delta function facilitates formal rules for manipulating certain integrals, with an incoherent (but quite intuitive) underlying justification, and Heaviside's operational calculus largely involves formal rules for manipulating expressions that again cannot be justified by underlying structure unless we appeal to heavy-duty mathematics that wasn't in the picture at the time, e.g. Bromwich's Laplace transforms. Indeed, it seems that wherever we find an inferentially restricted methodology, we will find a formalism that floats free from structure in this way. Nonetheless, these free-floating formalisms sufficed for extremely successful applications of mathematics. Indeed, calculations made using (in part) path integral techniques have yielded predictions confirmed by some of the most precise measurements ever made. This suggests that the structures and mappings appealed to even by inferential versions of the mapping account are besides the point, at least as far as the epistemic credentials of these applications are concerned. RIC does better because it allows us to bypass the structural epicycles necessary to represent these cases in terms of a mapping account. As a result, we have a much simpler way of making sense of how the informational content of these representations supports the mathematically mediated inferences that scientists made, which in particular allows us to appeal more directly to the most epistemically relevant features of their practice, the inferentially restrictive methodologies that they employed to keep their unrigorous mathematical tools under control.

5.4.2 Are these applications of *mathematics*?

Finally, one might object that the cases considered in this chapter do not tell against mapping accounts because these are not applications of genuine mathematics at all. These are cases in which scientists have developed computational and representational devices on the fly to suit their needs in particular cases. To the extent that these devices fall short of the standards of mathematical rigor, we should treat them as a sort of pseudo-mathematics, not genuine

5.4 Conclusions: The scope of an account of mathematical scientific representation

mathematics. Thus, the thought goes, it is not a shortcoming of the mapping account if it does not provide the best account of such cases; it was never intended to do so.

However, I think we should resist this thought for two reasons. First, the claim that these cases do not constitute applications of mathematics is not very well motivated. And, second, even if we do not count these cases as applications of genuine mathematics, we should recognize that they are very much *like* mathematics—enough so that we should tell a similar philosophical story about both kinds of case.

Regarding the first point, not to treat them as genuine mathematics would seem to rule out a great deal of historical mathematics, as, historically, mathematical work failed to meet the modern standards of mathematical rigor (and often quite badly). In particular, this would certainly have to rule out the calculus of Leibniz and Newton and the work of subsequent mathematicians who built on it prior to the introduction of ϵ, δ limits. It would also seem to rule out much of the work of Fourier and Boole, who each had operational calculi much like that of Heaviside. Likewise, the Fourier integral theorem as presented in [Fourier \[1822, p. 525\]](#) is essentially equivalent to the introduction of the Dirac delta function as $\delta(x - \alpha) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \cos(px - p\alpha) dp$. I see no principled reason to count, say, the work of Fourier and Boole as genuine mathematics but the work of Heaviside and Dirac as “pseudo-mathematics”—except perhaps that Fourier and Boole were *considered* to be mathematicians by their contemporaries, while Heaviside and Dirac were not.

But even if we do not count these cases as applications of “genuine mathematics,” there is reason to think a good account of applications of mathematics should be able to accommodate them—or, in any case, that it is advantageous for such an account to do so. First, for the reasons spelled out above, the resources needed to make sense of such cases will be also be needed to make sense of historical applications of mathematics (including most notably the early calculus) that predate our current standards of mathematical rigor. Second, these cases are similar enough at the level of practice to applications of rigorous mathematics that we should expect our philosophical story about them to be substantially similar to our story

5.4 Conclusions: The scope of an account of mathematical scientific representation

about applications of rigorous mathematics.

Part III

Explanation and Understanding

In this part of the thesis, I shift from considering the role mathematics plays in scientific representations to considering its role in scientific explanations and understanding. Unlike in the previous parts of the thesis, my primary concern here is not to compare RIC to mapping accounts, but rather to show RIC in action.

In chapter 6, I use RIC to help describe a metarepresentational role for mathematics in science that goes beyond its simple representational role. This role, I argue, helps to explain how mathematics contributes to the high degree of generality possessed by certain scientific explanations.

In chapter 7, I turn to the closely related question of the contributions of mathematics to scientific understanding. I show how RIC can be combined with an influential inferential account of explanation to produce a highly nuanced account of the relationship between mathematical rigor and scientific understanding.

CHAPTER 6

The Metarepresentational Role of Mathematics:
Mathematical Scientific Explanations in the RIC
Framework

6.1 Introduction

Mathematics clearly plays some role in many scientific explanations. Whenever we represent a physical domain mathematically, mathematics can be used to represent explanatory physical facts, the explanantia of explanations in that domain.¹ For instance, to explain why Jamie and Jon took six hours to drive from Detroit to Milwaukee, we might appeal to the facts that the route they traveled was approximately 360 miles long and that they traveled at approximately 60 miles per hour. While these facts are represented with numbers, they are not themselves facts about numbers. They are facts about a particular physical system in which Jamie and Jon drive a particular distance at a particular speed. In these cases, mathematics is said to play a *merely representational* role in the explanation. According to some—call them *representationalists*—such cases exhaust the role of mathematics in scientific explanations [e.g., [Daly & Langford, 2009](#), [Melia, 2000, 2002](#), [Saatsi, 2011, 2016](#)].

In contrast, *explanationists* have argued that mathematics sometimes plays a further, *distinctively explanatory* role. In these cases, mathematical facts would be among the explanantia of non-mathematical explananda. For instance, suppose Jamie and Jon brought five sandwiches to eat during their drive. We might take the purely mathematical fact that five isn't divisible by two to explain the fact that they didn't manage to split their sandwiches evenly (without cutting). This is what [Saatsi \[2016\]](#) calls a “thick” explanatory role: mathematical facts stand in an ontic relation of explanatory relevance to their explananda.² In such cases, facts about the target system would have to stand in an objective dependence relation to mathematical facts, construed as genuine constituents of the world.

This role for mathematics has been defended in part on the grounds that some explanations have features that can only be explained if we take the mathematics involved to play a

¹Here and throughout this chapter, I use ‘explanans’ and ‘explanandum’ to refer to the relata of the worldly dependence relations involved in successful explanations, rather than the communicative devices used to pick those relata out. While not all accounts of explanation entail that all successful communicative acts of explanation must pick out such a dependence relation, for the purposes of this chapter I assume that such dependence relations must be involved in one way or another in the explanations I discuss.

²Note that this means that the influential modal account of mathematical explanations in science proposed by [Lange \[2013\]](#) is not explanationist in the sense I have in mind here.

distinctively explanatory, rather than merely representational role in the explanation.³ These explanations are thought to be more general than explanations in which mathematics plays a merely representational role, in that they carry more counterfactual information about their target systems (scope-generality) and in that they share an “explanatory core” with explanations of phenomena in other domains (topic-generality). For instance, parallel explanations in terms of divisibility explain why Jamie and Jon couldn’t evenly split seven—or nine, or any odd number of—sandwiches (scope-generality) and why a philosopher can’t evenly divide their five good ideas between two papers (topic-generality). In contrast, we couldn’t similarly generalize the explanation of why it took six hours to travel from Detroit to Milwaukee, since the explanantia of that explanation are concrete features of the target system, rather than more general features shared by, say, systems in which different distances are traveled (scope-generality) or systems in entirely different domains (topic-generality). For much the same reason, we couldn’t similarly generalize an explanation of why Jamie and Jon couldn’t evenly split their sandwiches that appealed only to concrete properties of the physical system consisting of Jamie, Jon, and their sandwiches.

The debate between representationalists and explanationists has primarily been carried out in the context of evaluating indispensability arguments for mathematical platonism. While the classic Quine-Putnam indispensability argument supports platonism on the grounds that mathematics is indispensable to our best scientific theories, more recent versions of the argument appeal to its *explanatory* indispensability. In short, the thought is that we should believe in mathematical objects for the same reason we believe in non-mathematical theoretical posits—due to their indispensable role as explanantia in our best scientific explanations.⁴ While such arguments might fail for other reasons, their success depends in large part on whether mathematics plays the right sort of role in our best scientific explanations—

³For example, [Colyvan \[2002\]](#), [Baker & Colyvan \[2011\]](#), [Lyon \[2012\]](#), [Plebani \[2016\]](#), [Baker \[2017\]](#), and [Baron \[2020\]](#) all argue for a distinctively explanatory role for mathematics on the grounds of explanatory generality or cognate notions like robustness or unification.

⁴Perhaps the most influential formulation in terms of explanatory indispensability is in [Baker \[2005\]](#), which has spawned an enormous literature. For a survey, see [Mancosu \[2018, §3.2\]](#).

in particular, the distinctively explanatory role described above, according to which physical explananda stand in ontic explanatory dependence relations to mathematical explanantia [Saatsi, 2016]. And it is on precisely this point that representationalists and explanationists disagree, with the former denying and the latter affirming that mathematics plays such a role.

But the highly abstract scientific explanations at the center of this debate are independently important. Explanations of this kind are crucial to the burgeoning literature on non-causal explanations in science (see, for example, the papers in Reutlinger & Saatsi [2018]). Understanding how they work—and in particular how mathematics contributes to their explanatory generality—is crucial to understanding scientific explanation more generally. In this regard, existing explanationist and representationalist approaches *all* leave something to be desired.

If explanationists are right, we have to explain the very possibility of a mathematical fact's standing in an ontic relation of explanatory relevance to a physical fact. This is far from straightforward. Most realists about mathematics take mathematical and physical entities to belong to different ontological categories. But even nominalist interpretations and paraphrases of mathematical language—for instance, in terms of modal information about possible structures (modal structuralism) or facts about mathematical practice (some forms of fictionalism)—don't obviously pick out the right sort of entities to stand in such relations. We also would have to explain why increasing the *degree* of generality of an explanation eventually yields a different *kind* of explanation altogether—viz., one in which the explanans must be a mathematical, rather than physical, fact. But it is not obvious that such a difference in degree should yield a difference in kind.⁵

On the other hand, representationalist accounts, when they explicitly address the nature of mathematical representation, typically rely on an austere conception according to which mathematical facts simply “index” physical ones [Daly & Langford, 2009, Melia, 2000], which fails to do justice to the full range of contributions mathematics makes to science. Such

⁵For an argument to that effect, see Jansson & Saatsi [2019].

accounts have little to say about how mathematics contributes to explanatory generality. Arguably the best move for the representationalist to make in the context of debates over the explanatory indispensability of mathematics is to deny that degrees of scope- or topic-generality beyond those straightforwardly available to the nominalist are explanatory *virtues* that would support the inference to the best explanation central to explanatory indispensability arguments [Knowles & Saatsi, 2021]. But regardless of whether these degrees of generality are virtuous—a question on which we can remain agnostic for present purposes—they are properties of certain explanations in which mathematics plays a prominent role, and a good philosophical account of such explanations should be able to account for them.

In this chapter, I argue that by looking more closely at the role of mathematics in scientific *representations* we can understand how scientific explanations can achieve such high degrees of scope- or topic-generality without including mathematical facts among their explanantia. In addition to representing particular target systems, mathematics allows us to represent properties of these representations that remain stable as the mathematics involved and its physical interpretation are allowed to vary. This is what I call the *metarepresentational role* of mathematics. In using the word ‘metarepresentational’ here, I do not mean to imply that in this role mathematics is just a “meta-level device” for reasoning about science from the outside, but rather to describe the use of mathematics within science to reason about features shared by collections of individual representations.⁶ This metarepresentational contribution of mathematics allows us to reason about the very abstract features of physical target systems in virtue of which they are accurately represented by certain *kinds* of mathematical representation. These abstract features of target systems, rather than the mathematical facts relevant to representing them, are the explanantia of highly general explanations. While mathematics is at least practically necessary to pick out these features, this use of mathematics does not commit us to the truth of any purely mathematical (as opposed to physically interpreted)

⁶That said, if the reader balks at the use of the word ‘metarepresentational’ for any reason, what I call the “metarepresentational role” can without great loss be taken to be part of a significantly enriched account of the representational role of mathematics, rather than a distinct role.

claim, and so does not support explanatory indispensability arguments for mathematical platonism. The result is a richer conception of the role of mathematics in scientific explanations that significantly improves on existing representationalist and explanationist accounts.

In §6.2, I develop a more nuanced account of the role of mathematics in scientific representations, showing how RIC as an account of mathematical representation can be extended to explain the metarepresentational contributions of mathematics. I then distinguish two kinds of explanatory generality (§6.3) and present a scientific explanation, the famous number-theoretic explanation of the cycles of periodical cicadas, that has been taken to exhibit both (§6.4). In reference to this example, I show how we can explain these features in terms of the metarepresentational role of mathematics without appealing to pure mathematical facts (§§6.5–6.6). Finally, I respond to the objection that explanations in which mathematics plays only a metarepresentational role have less explanatory depth than those in which mathematics plays a distinctively explanatory role (§6.7).

6.2 Representation and Metarepresentation

In its representational role, mathematics serves to construct individual epistemic representations of particular target systems. In its metarepresentational role, it serves to elucidate the properties shared by collections of individual representations, as well as very general properties of these representations' target systems that can be picked out only by reasoning about the features of collections of representations. The purpose of this section is to clarify and illustrate this central distinction before applying it to the case of scientific explanations appealing to mathematics in the rest of the chapter.

6.2.1 Representation

In its representational role, mathematics helps scientists to construct epistemic representations of particular non-mathematical target systems. These representations support making

inferences about their target systems on the basis of the relevant mathematics.

To focus the discussion in this chapter, I will frame both the representational and metarepresentational roles of mathematics in terms of RIC. I do this not just because I develop RIC earlier in the thesis, but because I think it more perspicuously represents the central features of mathematical representations relevant to the discussion here. Nonetheless, with minor adjustments, the rest of this chapter could be cast in terms of a version of the mapping account instead.⁷

Recall that, according to RIC, mathematics places constraints on what the target system of a mathematical scientific representation must be like by specifying inferences about the target system that must preserve truth if the representation is accurate. Such representations have three ingredients:

- (RIC1) a *physical interpretation* of the language of the mathematical theory sufficient to provide at least some expressions in this language with physical truth conditions,
- (RIC2) an *initial description of the target system* in the language of the mathematical theory, given this interpretation, and
- (RIC3) a *collection of privileged inference patterns* from those licensed by the original mathematical theory.

The commitments of the representation are the physically interpreted versions of the claims in the language of the mathematical theory that are in (RIC2) or can be derived from these via the inference patterns in (RIC3).

Consider a model of a weight suspended from a spring as a damped harmonic oscillator. Its behavior is represented by the differential equation $m \frac{d^2x}{dt^2} + c \frac{dx}{dt} + kx = 0$ and perhaps by equations determining values for the constants m , c , and k . These equations provide component (RIC2), the initial description of the target system. In this case, x is interpreted as the

⁷Recall that mapping accounts are best understood as special cases of RIC, with (RIC1) provided by the relevant structure and mapping, (RIC2) by a “structure-generating description” [Nguyen & Frigg, 2021] or something similar, and (RIC3) by those inferences that preserve truth when interpreted in terms of the relevant mathematical structure, so this recasting could be done with relative ease.

vertical displacement of the object suspended from the spring (and its derivatives with respect to t the vertical components of velocity and acceleration); t is interpreted as time elapsed; m is interpreted as the mass of the object suspended from the spring; c represents the effect of damping (related to the force F_f due to friction by the equation $F_f = -c\frac{dx}{dt}$); and k represents the effect of the tension of the spring (related to the force F required to extend the spring by length x by the equation $F = -kx$). This suffices to provide numerical solutions of this and related equations with physical truth conditions, and so it suffices for (RIC1). Finally, (RIC3) is simply the set of inference patterns licensed by real analysis.⁸

Now, consider the closure of the (mathematically interpreted) set of equations in RIC2 under the inferences licensed by real analysis. A subset of the expressions in this set are assigned physical truth conditions by the interpretation RIC1. The expressions in this subset, under that physical interpretation, constitute the informational content of the representation. This then allows us to see how such a representation can support surrogate reasoning about the target system via mathematical reasoning about the real numbers. Whatever results we can derive purely mathematically in real analysis by means of the equations in RIC2 can be brought to bear on the physical target system by interpreting them according to the interpretation rules in RIC1, provided that RIC1 supplies those results with physical truth conditions. Since such physically interpreted claims are by definition part of the informational content of the representation, such inferences preserve truth on the condition that all of the representation's informational content is true, and are therefore licensed by the representation.⁹

⁸Restriction of the inference patterns in (RIC3) is only required where the relevant mathematical theory is inconsistent or otherwise unrigorous, so that allowing all inference patterns made available by the mathematical theory might have undesired results. See chapters 4 and 5.

⁹Of course, most—if not all—scientific representations have some informational content that is untrue. What matters to scientists is instead whether their representations are accurate in the appropriate respects. In light of this, I might seem to be presupposing an implausibly close relationship between truth-preservation and accuracy. But I need not make any such presupposition. If a representation is accurate in some respects but has some false informational content, there is no guarantee that the inferences licensed by the representation are actually truth- or accuracy- preserving. Whether such an inference should be made depends also on whether the representation is accurate in the right respects.

6.2.2 Metarepresentation

A *representation schema* is a collection of mathematical representations that share common features but vary with respect to some aspect of the representation. In particular, this means that one may vary the physical interpretation assigned to the mathematical vocabulary (RIC1), the specification of the target system (RIC2), or the mathematical apparatus itself (RIC3).

To see how this works, consider how we might extend the representation of a weight suspended from a spring to a representation schema. Since a number of very different physical systems can be usefully modeled as harmonic oscillators, we might be interested in various schemata that allow either the physical interpretation of the mathematical language (RIC1) or the initial description of the target system (RIC2)—or both—to vary in certain ways.

For instance, we might start with a representation of a weight-and-spring system that specifies the values of the constants m , c , and k , but we might be interested in features that representation shares with representations that specify different values for those constants. In that case, we could reason about the schema consisting of representations that share the same physical interpretation (RIC1) and underlying mathematical framework (RIC3) but differ with respect to the initial description of the target system (RIC2) in that (at most) different values are given to the constants m , c , and k .

We might also be interested in features that this representation shares with those of other kinds of systems that can be represented as damped harmonic oscillators, such as pendulums or certain electrical circuits. In that case, we might start with the representation schema considered above, which allows the values of the relevant constants to vary, and further allow the physical interpretation (RIC1) to vary in a limited way (so that the relevant mathematics can be interpreted in terms of pendulums and electrical oscillators but not other sorts of target system). For instance, to allow representations of pendulums to be instances of our schema, we must allow x to be angular (rather than vertical) displacement, m to be rotational inertia (rather than mass), and so on. If we also wish to include representations that do not specify these constants directly, but derive them from other features of the target system, we may

need to allow more substantial changes to the initial specification of the target system (RIC2), allowing it to incorporate the equations used to derive the values of these constants. It is also possible that deriving these constants would require appealing to mathematical resources beyond those supplied by real analysis. In that case, the set of available inference patterns (RIC3) would also have to be allowed to vary accordingly.

For instance, we might be interested in the features the representation from the previous section shares with representations that specify other values for the constants m , c , and k . In that case, we could reason about the schema consisting of representations that share the same physical interpretation (RIC1) and underlying mathematical framework (RIC3) but differ with respect to the initial description of the target system (RIC2) in that (at most) different values are given to the constants m , c , and k .

We might also be interested in features that this representation shares with those of other kinds of systems that can be represented as damped harmonic oscillators, such as pendulums or certain electrical circuits. In that case, we might start with the representation schema considered above, which allows the values of the relevant constants to vary, and further allow the physical interpretation (RIC1) to vary in a limited way (so that the relevant mathematics can be interpreted in terms of pendulums and electrical oscillators but not other sorts of target system). For instance, to allow representations of pendulums to be instances of our schema, we must allow x to be angular (rather than vertical) displacement, m to be rotational inertia (rather than mass), and so on. If we also wish to include representations that do not specify these constants directly, but derive them from other features of the target system, we may need to allow more substantial changes to the initial specification of the target system (RIC2)—namely by allowing it to incorporate the (physically interpreted) equations used to derive the values of these constants.

Reasoning with representation schemata allows us to do two things we cannot do with individual representations alone. First, the mathematics common to the instances of the schema (or a mathematical framework into which these instances are embedded) can be used to shed

light on features shared by the individual representations in the schema. Second and closely related, reasoning with the schema allows for reasoning about general features shared by the target systems represented by individual representations in the schema, which are not captured (in their full generality) by these individual representations. As might be expected, capturing these general features, both of representations and their target systems, is central to achieving the kind of generality possessed by scientific representations in which mathematics seems to play an explanatory role, to which I turn in the next section.

6.3 Explanatory Generality

Many explanationists have claimed that the degree and kind of generality possessed by certain scientific explanations requires them to include mathematical facts among their explanantia. [Baker \[2017\]](#) gives perhaps the clearest statement of this idea by distinguishing two kinds of generality that mathematics might help us achieve: scope-generality and topic-generality.

6.3.1 Scope-generality

An explanation is more *scope-general* the wider the range of counterfactual situations to which it applies in which the explanans is varied (but involving the same sort of target system). Consider the explanation of why Jamie and Jon couldn't evenly split five sandwiches that appeals to the mathematical fact that five is not divisible by two. This explanation is scope-general in that a parallel explanation can be used to determine what happens whatever number of sandwiches or passengers there are: one can divide n sandwiches evenly among m passengers if and only if n is divisible by m . And so we can explain by exactly the same means why it's not possible to evenly split nine sandwiches between two people, five sandwiches among three people, and so on. A less scope-general explanation would explain why Jamie and Jon cannot evenly split five sandwiches without telling us anything about what happens when the number of sandwiches or passengers is different. And an explanation with an intermediate

degree of scope-generality might tell us what would happen if there were n sandwiches and m passengers, provided n and m don't exceed some finite upper bound.

Some explanationists have argued that maximal scope-generality can only be achieved by explanations that appeal to mathematical facts, like the explanation in terms of the fact that five is not divisible by two. Any explanation that holds no matter how many sandwiches and passengers there are, the thought goes, must appeal to a mathematical relation like divisibility. For instance, Baker & Colyvan [2011, p. 331] argue that explanations in which mathematics plays a merely representational role must be “less general and less robust” (i.e., less scope-general) than those in which the mathematics is explanatory. Lyon [2012] similarly claims that explanations in which mathematics plays an explanatory role are more “robust”¹⁰ with respect to causal-historical details (p. 567). Plebani [2016] criticizes Liggins [2016] on the grounds that his nominalist-friendly explanations operate “at the wrong level of generality” (p. 553) for the similar reason that a nominalistic explanation's explanans would have to include too many concrete details.

After presenting a more sophisticated example of a scope-general explanation in §6.4, I will argue in §6.5 that these arguments miss the mark. Maximal scope-generality can be achieved by an explanation that does not appeal to mathematical facts as explanantia, but only if mathematics is allowed to play a metarepresentational role in the explanation. In that case, an explanation can be formulated in terms of a representation schema that carries more counterfactual information than individual representations of the relevant target system.

6.3.2 Topic-generality

An explanation is *topic-general* to the extent that it does not depend on the concrete features of a particular target system, but has an explanatory core that can be used to formulate parallel explanations about target systems of radically different types. Consider again the mathemat-

¹⁰The sense of ‘robust’ in both Baker & Colyvan [2011] and Lyon [2012] should be distinguished from its meaning in more general discussions of modeling. Baker, Colyvan, and Lyon have something narrower in mind—viz., a kind of explanatory scope-generality.

ical explanation of why Jamie and Jon cannot evenly split five sandwiches. The core of this explanation has nothing to do with Jon, Jamie, or sandwiches. A parallel explanation can be formulated for any system of discrete objects that one might want to partition into a finite number of equinumerous groups. The explanatory core is something like

1. It is possible to divide n F s into m even groups iff n is divisible by m .
2. Five is not divisible by two.
3. So, it is not possible to divide five F s into two even groups.

We might think of this as an explanation schema à la [Kitcher \[1989\]](#) that can be filled in by replacing ‘ F ’ with a term referring to some kind of discrete object to produce a concrete explanation.¹¹ In that case, we get our toy explanation by substituting ‘sandwiches’ for ‘ F ’. But we get just as good an explanation by replacing ‘ F ’ with any other expression picking out a group of discrete objects. For instance, we can produce in this way an explanation of why it is not possible to divide five students into two even discussion groups or of why it is not possible to distribute one’s five good ideas evenly across two papers (assuming, of course, that ideas are discrete objects).

Alternatively, we might think of this as a set of propositions that is itself topic-general in that it doesn’t appeal to concrete features of any type of non-mathematical target system in particular. Rather, if we take (1) and (3) to implicitly quantify over all types F of discrete objects and (1) to implicitly quantify over all $n, m \in \mathbb{N}$, it appeals only to a mathematical fact about natural numbers (2) and a proposition connecting mathematical facts with very general facts about dividing discrete objects into groups (1). Particular explanations are produced

¹¹Central to Kitcher’s unificationist account of explanation is the notion of an argument pattern, which consists of a deduction with the non-logical vocabulary replaced with schematic letters, a set of “filling instructions” specifying how those letters can be filled in to produce a legal instance of the schema, and a classification of the schematic argument. According to the unificationist account, an ideal explanation is an instance of an argument pattern belonging to the set of such patterns that best unifies the set of beliefs accepted by scientists at a particular time. But we need not accept the entirety of this unificationist account to represent explanations in terms of argument patterns in this way. A similar strategy is pursued by [Baron \[2020\]](#), who proposes a hybrid unificationist-counterfactual account of mathematical explanations in science.

by supplementing this explanatory core with propositions about the particular domain of the explanation. We get our toy explanation by adding the proposition that sandwiches are discrete objects and concluding from this and (3) that five sandwiches cannot be divided into two even groups. But we might add other propositions to produce parallel explanations—for instance, that students are discrete objects or that ideas are.

Explanationists have argued that topic-generality is a distinguishing feature of explanations in which mathematics plays such a role. In several places, Colyvan [2002, p. 72; 2013, p. 1042] claims that mathematics is genuinely explanatory due to its unifying power—i.e., its ability to produce topic-general explanations. The central claim of the paper in which Baker [2017] articulates the distinction between scope- and topic-generality is that topic-generality requires mathematics to play a genuinely explanatory role even if scope-generality does not. Along similar lines, Baron [2020] develops a view according to which topic-generality (alongside a few other conditions) distinguishes genuinely mathematical explanations from explanations in which the mathematics is merely representational.

After considering a more sophisticated example (§6.4) and the case of scope-generality (§6.5), I will argue in §6.6 that we can explain the topic-generality of such scientific explanations without treating the mathematics as playing a distinctively explanatory role. Again the mathematics should be understood as playing a metarepresentational role. In this case, the core of an explanation can be formulated in terms of a representation schema that allows the interpretation of the mathematics to vary. The very abstract features of physical target systems in virtue of which they are correctly represented by some instance of this schema then serve as explanantia in the topic-general explanation.

6.4 Example: The Prime Cycles of Periodical Cicadas

A classic example of an explanation that is both scope- and topic-general is the number-theoretic explanation of the prime periods of periodical cicadas, first discussed in the philo-

sophical literature by Baker [2005]. Periodical cicadas spend most of their lives in a larval stage underground, emerging as adults for a single season, during which they reproduce and die. In extant species, this life-cycle lasts either thirteen or seventeen years. Why are these period lengths adaptive?¹²

According to Goles *et al.* [2001], it is because prime periods optimize evasion of periodical predators with shorter life-cycles. Under certain assumptions, where x is the length of the cicadas' life-cycle and y the length of the predators' life-cycle, the cicadas' average fitness over a period of xy years is a decreasing function of $\gcd(x, y)$, while the predators' average fitness over the same period is an increasing function of $\gcd(x, y)$. By definition, all and only prime numbers minimize \gcd with all smaller positive integers. Assuming the predators must have a shorter period than the cicadas, this means that it is only when cicadas have a prime period that neither the cicadas nor their predators could increase their fitness by changing their period. To reflect biological constraints, Goles *et al.* constrain cicada periods to be between 12 and 18 years, resulting in 13- and 17-year periods.¹³

This explanation is both highly scope general and highly topic general, and these features will be the focus of the next two sections. I will argue that these two kinds of generality do not give us reason to think the mathematics is playing a distinctively explanatory role. Instead, the move from representation to metarepresentation allows us to pick out features of the phenomenon that are sufficiently general to ground explanations with the same degree

¹²Most presentations of this case, including Baker's original paper, treat the *existence* of 13- and 17-year cycles as the explanandum, and what I say in the rest of this chapter can naturally be adapted to such an explanation. However, there is good reason to doubt that the optimality of prime cycles actually played a role in their selection [Wakil & Justus, 2017]. As Wakil and Justus conclude, we can avoid this problem by instead treating the *adaptiveness* of 13- and 17-year cycles as the explanandum of the number-theoretic explanation.

¹³My presentation of this example is simplified in several ways in order to streamline the rest of the chapter. For one thing, the explanation I present is not the only number-theoretic explanation of this phenomenon. Another discussed by Baker [2005] is the suggestion that prime periods are advantageous because they minimize hybridization with other periodical cicada species. What I say in the rest of the chapter can straightforwardly be adapted to this explanation. For another, I ignore several important biological details, such as the pressures of nymphal crowding and deviations from 13- and 17-year periods (most often by four years) sometimes observed in existing cicada species. (See, for example, [Wakil & Justus, 2017].) While these details are essential to explaining why cicadas actually developed these periods, one can explain why those periods are *adaptive* without them. Since I only discuss the latter sort of explanation, and since my aim is not to give a complete account of this particular case, I take these omissions to be harmless.

of scope-generality (§6.5) and topic-generality (§6.6) as explanations in which mathematics plays a distinctively explanatory role without appealing to pure mathematical facts.

6.5 Metarepresentation and Scope-Generality

The number-theoretic explanation is highly scope-general. It does not just explain why the actual periods of actual cicadas are advantageous, but also why prime cycles would be advantageous if cicadas had different biological constraints on their life-cycles. Supposing cicada periods were biologically constrained to another range (say, 12 to 25 years), we have a parallel explanation of why prime periods within that range would be advantageous. The number-theoretic explanation indeed seems to be more scope-general than explanations that appeal only to nominalistically acceptable properties, which would seem to be limited to cases in which the cicadas' periods cannot exceed some finite bound.

The best existing responses to this line of thought concede for the sake of argument that explanations in which mathematics plays a distinctively explanatory role enjoy a higher degree of scope-generality than other explanations, but downplay the importance of such high degrees of generality. For instance, [Knowles & Saatsi \[2021\]](#) argue persuasively that scientists have no reason to prefer an explanation that holds for all ranges of possible cicada periods to one that holds only for periods up to some suitably high finite bound. If, say, the nominalistic explanation works for cicada period lengths up to the age of the universe, the additional counterfactual information provided by the explanation in which mathematics plays a distinctively explanatory role will be of dubious scientific value indeed!

But regardless of whether the high degree of scope generality enjoyed by the mathematical explanation is a virtue, it does seem to be a genuine feature of this and other scientific explanations in which mathematics plays a prominent role. A complete philosophical understanding of the role of mathematics in scientific explanations therefore still requires an account of what makes this greater degree of generality possible (or why, despite appear-

ances, it is not possible). In the rest of this section, I provide such an account in terms of the metarepresentational role of mathematics. If this account is correct, nominalists need not settle for explanations with limited scope-generality like those defended by Knowles and Saatsi.

The scope-generality of the explanation can naturally be understood in terms of a representation schema in which only the range of biologically feasible life-cycles for the cicadas and their predators is allowed to vary. Instances of this schema are those that can be produced from Goles et al.'s initial representation by varying the initial description of the target system (RIC2) with respect to the range of possible cicada periods. For the explanation to go through for all instances of the schema, we must also require that the maximum predator period is shorter than the minimum cicada period, that the range of predator periods is sufficiently broad to make all non-prime cicada periods unstable, and that the range of cicada periods contains at least one prime.

Instances of the resulting schema represent particular counterfactual situations in which the cicadas and predators are biologically constrained to have life cycles within a particular range. For each of these instances, we can run a parallel explanation of why the prime cycles within that range are optimal for predator avoidance. Where p_1, \dots, p_n are the prime numbers in the range of numbers representing biologically feasible cicada periods and x and y represent the cicadas' and predators' periods in years, respectively, a parallel series of applications of the inference patterns in (RIC3) and purely physical inferences yields the conclusion that periods of p_1 or ... or p_n years are optimal for predator avoidance:

1. Via the inference patterns licensed by number theory, infer that p_1, \dots, p_n are prime.
2. From 1, infer $\gcd(p_i, y) = 1$ for $1 \leq i \leq n$ and $y_{\min} \leq y \leq y_{\max}$.
3. From 2, infer that there is no y' such that $\gcd(p_i, y') < \gcd(p_i, y)$ when $1 \leq i \leq n$, $y_{\min} \leq y \leq y_{\max}$, and $y_{\min} \leq y' \leq y_{\max}$.
4. From 2, infer that there is no x' such that $\gcd(x', y) < \gcd(p_i, y)$ when $1 \leq i \leq n$,

$$x_{\min} \leq x' \leq x_{\max}, \text{ and } y_{\min} \leq y \leq y_{\max}.$$

5. From the physical interpretation of 3, infer that there is no way for the predators to increase their average fitness¹⁴ by changing their period, provided the cicadas' period is p_1 or ... or p_n years.
6. From the physical interpretation of 4, infer that there is no way for the cicadas to increase their average fitness by changing their period, provided that period is p_1 or ... or p_n years.
7. From 5 and 6, infer that periods of $p_1, \dots,$ and p_n years are optimal with respect to periodical predator avoidance.
8. By a similar chain of reasoning, infer that, when the cicadas' period is not p_1 or ... or p_n years, either the predators or the cicadas can increase their fitness by changing their period, and so those periods are not optimal with respect to periodical predator avoidance.
9. Conclude from 7 and 8 that the cicadas' optimal periods with respect to periodical predator avoidance are exactly $p_1, \dots,$ and p_n years.

Recognizing that this is the case for each instance of the schema allows us to formulate a more general explanation of why periodical cicadas have prime cycles: no matter which instance of the representation schema picks out the right ranges of biologically feasible periods for the cicadas and predators, the chain of inferences above goes through. This explanation goes beyond the explanations made available by the individual instances of the representation schema by incorporating all of the counterfactual information expressed by the schema; it tells us not just what happens given that cicada lifecycles are constrained to be (say) between 12 and 18 years, but also what would happen if that range were different. This is made possible

¹⁴This is how Goles et al.'s [2001] model works, anyway. Really, we should take their predator and cicada fitness functions to represent only the ways interactions between the predator and cicada periods contribute to overall fitness.

by reasoning not about the particular physical facts represented by a given instance of the schema, but the features shared by the systems represented by all these instances, which are the explanantia of the scope-general explanation. And this is achieved by reasoning about the schema as a whole. The result is scope-generality: the explanation generalizes over all relevant ranges of biologically feasible predator and cicada periods because it generalizes over all instances of the representation schema.

At this point, one might worry about the use of mathematical vocabulary in setting out the explanation above. In particular, (1) appears to be nothing but a pure mathematical fact, since ‘prime’ is not assigned a physical interpretation by the RIC1 component of the representations in the schema. And (2) and (3) also appear to state pure mathematical facts, albeit ones that can also be given a physical interpretation. And if this is the case, it seems wrong to say that mathematical facts are not among the explanantia of this explanation.

But this rests on a misunderstanding the role of (1)–(8) above. (1)–(4) result solely from the application of the inference patterns in the RIC3 component—the collection of privileged mathematical inference patterns from those licensed by the original mathematical theory—of the instances of the representation schema. The role of the inference patterns in RIC3 is solely to determine which inferences from physically interpreted claims in the language of the mathematical theory to physically interpreted claims in that language preserve truth according to the representation. For this purpose, whether (1)–(4) are true under their standard mathematical interpretation is irrelevant. Even supposing a form of mathematical error theory were true, so that the mathematical interpretations of (1)–(4) were false, (5)–(8) would still be consequences of each representation in the schema, since the inference patterns yielding (1)–(4) would still be in RIC3, and each of these representations has as part of its informational content the physically interpreted versions of any claims derivable from the initial specification of the target system RIC2 via the inference patterns in RIC3. So, while (1)–(4) have the surface appearance of statements of mathematical facts, in this context they merely express the inferences licensed by the instances of the representation schema. As a result, the

mathematics involved in their expression does not play a distinctively explanatory role.

6.6 Metarepresentation and Topic-Generality

The cicada explanation is also topic-general. According to Baker [2017] (cf. [Baron, 2020]), it has an explanatory core that is not specific to periodical cicadas but that applies generally to phenomena involving unit cycles with certain features.

Following Baker [2017, pp. 201f], but with adjustments to fit the description of the example here, this explanatory core consists of the propositions:

(M₁) The gcd of numbers m, n is minimal if and only if m and n are coprime. (pure mathematical fact)

(UC₁) The number of co-occurrences of the same pair of cycle elements of two unit cycles of periods m and n in an interval of length mn is equal to $\text{gcd}(m, n)$. (fact about unit cycles)

(UC₂) So, any pair of of unit cycles with periods m and n minimizes the number of co-occurrences of the same pair of cycle elements over an interval of length mn if and only if m and n are coprime. (fact about unit cycles, from M₁, UC₁)

(M₂) All and only prime numbers are coprime with all smaller numbers. (pure mathematical fact)

(UC₃) So, given a unit cycle p_m of length m and a range of unit cycles q_i with lengths $i < m$, p_m minimizes the number of co-occurrences of the same pair of cycle elements over an interval of length mi for all q_i if and only if m is prime (fact about unit cycles, from M₂, UC₂)

From UC₃, thus derived, and facts about periodical cicadas and their predators in particular, we can then derive the conclusion that 13- and 17-year periods are optimal with respect to predator avoidance, reasoning in much the same way as above.

This explanatory core is then (at least in part) shared by explanations of phenomena in other domains. For instance, we might give a parallel explanation of why the number of teeth on the front and rear gears of brakeless fixed-gear bicycles are optimal when they are coprime [Baker, 2017, pp. 203f]. In this case, we need only M_1 , UC_1 and UC_2 , since appeal to primeness is unnecessary. We then can add particular facts about fixed-gear bikes to derive the explanandum. Stopping such a bike involves locking the pedals at a particular point (so that the front gear is always in the same position), thus locking the rear tire, causing the bike to skid to a stop. This causes wear on the tire where it skids. To maximize the life of the bike's tires, one should choose a gear ratio that minimizes the frequency with which the same part of the tire skids when the pedals are in braking position. Since we can treat the positions of the front and rear gears of the bike as unit cycles, UC_2 tells us that this happens just when the numbers of teeth on the two gears are coprime.

The thought is then that topic-generality of this kind can only be achieved when mathematical facts are included as part of the explanatory core. The explanatory core of an explanation in which mathematics doesn't play an explanatory role must appeal to particular facts about the domain of the explanation that prevent it from generalizing to other domains in this way. If we explain why 13- or 17-year cicada periods are advantageous in terms of properties of temporal intervals, as Saatsi [2011] does, this rules out using the same explanantia to explain the optimal gear configurations of fixed-gear bicycles.

In the rest of this section, I argue that topic-generality is possible because mathematics plays a metarepresentational role in the explanation. The explanatory core doesn't include pure mathematical facts like M_1 and M_2 above, but only very general, mathematically represented physical facts. And representation of these very general facts is made possible by the move from individual mathematical representations to a representation schema. Recall the pendulum example from section 2. Moving to a representation schema in which the physical interpretation RIC_1 was allowed to vary made it possible to reason about the features such a pendulum shares with radically different systems, like electrical oscillators. A similar move to

a schema in which RIC1 is allowed to vary allows us to capture features shared by periodical cicada populations and fixed-gear bicycles.

All that was needed for the scope-general explanation above to go through was for there to be two entities with cycles of length x and y in some unit such that $x_{\min} \leq x \leq x_{\max}$ and $y_{\min} \leq y \leq y_{\max}$ for $x_{\min}, x_{\max}, y_{\min}, y_{\max} \in \mathbb{N}$, with a guarantee that $x > y$, and that it is optimal to minimize the intersection between these cycles over an xy -year period. These features constitute the “explanatory core” of the scope-general explanation in the previous subsection. The basic predator-prey dynamics represented in Goles et al.’s model explain how these features are realized, but the explanatory core is independent of those details. In this respect, the explanatory core I present here is analogous to the explanatory core proposed by Baker [2017]. The crucial difference is that Baker’s explanatory core includes pure mathematical facts (M_1 and M_2) relating prime numbers, gcd, and coprimeness, while mine includes only very general, mathematically represented physical facts. But the explanatory core I present here allows for the formulation of explanations just as topic-general as those built on the explanatory core presented by Baker.

We can achieve this degree of topic-generality by moving to a yet more general representation schema that allows the mathematical representation to vary, as long as the three features above (or potentially a subset thereof) are preserved. Clearly, it is necessary to allow the physical interpretation of the mathematical language (RIC1) to vary, so that x and y can come to represent, say, spatial magnitudes, or numbers of teeth on interlocking gears, rather than temporal magnitudes. In addition, it will be necessary to allow the initial description of the target system (RIC2) to vary with respect to its mathematical formulation in order to capture the different ways in which the three conditions above might be realized.

Parallel explanations can then be given for phenomena represented by instances of this schema that represent phenomena in radically different domains. Consider again the explanation of why it is optimal to have coprime numbers of teeth on the gears of brakeless fixed-gear

bicycles, as discussed by Baker [2017].¹⁵ In this case, x and y come to represent the period of the front and rear gears, respectively, where the appropriate unit is teeth. Since the numbers of teeth on the gears do not tend to be prime (but only coprime), it is not necessary to restrict the representation schema to cases in which $x > y$ (though this happens to hold in realistic cases). There is some finite possible range of values for x and y determined by various constraints, such as that the gear ratios should be within a reasonable range for the purposes of actually riding the bike and that the gears should be possible to manufacture for a reasonable price. Configurations that minimize the co-occurrence of the same pair of cycle elements (i.e., gear positions) are optimal in this case because they maximize the wear on the rear tire maximally even, thus maximizing its useful life. So we have a subset of the explanatory core of the cicada explanation, together with an explanation of how the target system realizes the abstract conditions in the explanatory core.¹⁶

This gives us the essential ingredients for an explanation of the optimal gear configurations of fixed-gear bikes that runs parallel to the explanation of the optimal periods of periodical cicadas: for each instance of this schema physically interpreted in terms of fixed-gear bikes, we can run an explanation of why certain gear configurations are optimal by showing that these gear configurations minimize the co-occurrence of the same pair of gear positions by checking each of these gear configurations individually. And if we wish to produce a corresponding scope-general explanation, we can do so by appealing to the coprimeness of the numbers of teeth on the gears in the same deflationary way as the appeal to primeness in the scope-general cicada explanation: no matter which instance of the representation schema

¹⁵Baron [2020] presents a similar explanation of why interlocking gears on machines in general maximize the life of the machine by ensuring even wear. What I say about Baker's example can naturally be adapted to Baron's.

¹⁶Rather than formulate this explanation in terms of the explanatory core of the cicada explanation, construed as a set of propositions, we can equivalently formulate an explanation schema à la Kitcher [1989] or Baron [2020] that unifies the explanations of the cycles of periodical cicadas and of the gear configurations of fixed gear bikes. Whenever we have a target system that is accurately represented by an instance of our representation schema, a formally identical argument explains the parallel explanandum for that target system. The difference between the sort of explanation schema I propose and the one Baron [2020] proposes then parallels the difference between my explanatory core and the one proposed by Baker [2017]. Baron's schema appeals to pure mathematical facts (as opposed to merely mathematically represented physical facts) and in particular to subjunctive conditionals with purely mathematical antecedents. In contrast, my approach allows us to formulate an explanation schema that appeals only to very general, mathematically represented physical facts.

is accurate, coprime periods maximize the life of the rear tire by minimizing the frequency with which the rear gear (and so the rear tire) is configured in the same way when the front gear is in braking position. As in the case of the scope-general cicada explanation, this use of ‘coprime’ does not involve an appeal to mathematical facts, since it serves only to pick out certain patterns of inference (RIC3) common to instances of the representation schema, and these patterns of inference would be licensed by the representations regardless of whether the relevant statements including mathematical vocabulary like ‘coprime’ were true when given their usual mathematical interpretation. And so we can understand the topic-generality of the cicada explanation without ascribing a distinctively explanatory role to the relevant mathematics.

6.7 Explanatory Depth

Baker anticipates something similar to the account I provide in the previous section and objects: “[A]lthough a schema about unit cycles has more topic-generality than [a] schema about life-cycle periods, instantiations of the schema will still end up being treated as disjoint facts, and thus the overall explanation will be less unified and will have less explanatory depth than the full [mathematical explanation of the cicadas’ periods]” [Baker, 2017, p. 12]. So even if the representation schema I discuss at the end of the last section can support topic-general explanations, the resulting explanations lack the *depth* of explanations in which mathematics plays a distinctively explanatory role. The thought would seem to be that each instance of that schema would support explanations that rely on concrete facts that are themselves in need of explanation. If the schema incorporated mathematical facts, as Baker’s does, then these could be used to explain these more concrete facts. But, at first glance, I seem to have no resources to similarly explain the explanatory features of these more particular explanations in terms of “deeper” features of reality that transcend the various domains represented by instances of my representation schema.

But I think this is too hasty. The features in virtue of which a target system is accurately represented by an instance of the representation schema are themselves extremely general, due to the wide range of target systems covered by the schema. In the case of the schema that generalizes the cicada representation, this requires only (1) two entities in the target systems with cycles of x and y in some unit, (2) finite upper and lower bounds on both x and y , (3) that minimizing the frequency of the intersection of these cycles is optimal, and (4) that the mathematical inference patterns used in the explanation (a subset of RIC3) preserve truth under the physical interpretation RIC1.

The first three of these conditions are quite abstract, but distinctly physical; the explanatory cores and explanation schemata put forward by Baron and Baker incorporate similar conditions (for instance, Baker's very general claims about unit cycles). The final condition plays a role closer to that of the mathematical facts in Baker's and Baron's explanation schemata, but it too picks out very high level physical features shared by the relevant target systems. It requires each target system T to be such that certain patterns of inference, which could in principle be spelled out in terms of purely physically interpreted mathematical language, preserve truth when interpreted in terms of T .

So it is not that the explanation schema that I put forward picks out only a very disjoint collection of topic-specific physical properties to serve as explanantia. Rather, it picks out extremely general properties common to the full range of target systems covered by the relevant representation schema, including target systems in very different domains. Arguably such properties are *better* placed to do the right explanatory work than bare mathematical facts, as they are straightforwardly instantiated in each target system accurately represented by an instance of the schema.

But at this point, Baker, Baron, and others might object that I have begged the question against them by insisting that the very general features picked out by mathematical representation schemata of the kind I have considered are not mathematical. Now, it is true that these features often do not seem to be expressible in purely non-mathematical language, and

mathematical language has an important role to play in the account I've presented here. Any mathematical representation schema requires that certain patterns of inference, spelled out in terms of physically interpreted mathematical language,¹⁷ preserve truth when interpreted in terms of the target system of any instance of the schema. But the physical interpretation of this language plays a crucial role, and the result is something quite different from a bare mathematical fact (such as that all and only primes p minimize $\gcd(p, q)$ for all $q < p$), which is given no such interpretation.

The general features I appeal to are highly abstract, mathematically represented features of physical systems, while mathematical facts concern features of abstract objects (or whatever we take the subject matter of pure mathematics to be). These features may coincide if we *already accept* certain versions of mathematical realism—namely, ante rem structuralism [Resnik, 1997, Shapiro, 1997], Aristotelian realism [Franklin, 2014], and some versions of neologicism [Hale & Wright, 2001]—and it is not my aim to show that these views are false. Indeed, if one already accepts one of these views, there is still much to be gained by embracing my account of the metarepresentational role of mathematics, as it then provides a nuanced account of how mathematics contributes to explanatory generality and, in particular, how mathematical objects fit into such explanations.

But if we don't accept one of these views at the outset, then there seems to be no good reason to claim that the general features I appeal to are themselves mathematical. Both scope- and topic-generality are a matter of degree. And, all else equal, the more scope- or topic-generality an explanation possesses, the more abstract the properties in virtue of which a given system is correctly represented by some instance of the representation schema needed to articulate the explanation. To conclude that the general features I appeal to are themselves mathematical, we would need a further reason to think that these differences in degree should, after some threshold, yield a difference in kind.

¹⁷These are those inferences from physically interpreted premises to physically interpreted conclusions that can be arrived at by applying only the inference patterns in (RIC3).

6.8 Conclusion

I have argued that we can capture the high degree of generality and depth possessed by explanations in which mathematics figures prominently without treating the mathematics as playing a distinctively explanatory role—that is, without treating mathematical facts as themselves among the explanantia of the explanation. But doing so requires us to move to a more nuanced account of the role of mathematics in scientific representations, one that recognizes a new, metarepresentational role for mathematics in exploring properties shared by collections of mathematical representations.

The result is an account of the role of mathematics in scientific explanations that takes a middle path between the explanationist approaches of the likes of Baker, Baron, and Colyvan and the representationalist approaches of the likes of Melia and Saatsi, incorporating the best aspects of both approaches. Like existing representationalist accounts, it takes the role of mathematics in scientific explanations to be representational at bottom. But it takes this role to be significantly richer than simply “indexing” physical facts, allowing mathematics to do some (metarepresentational) heavy lifting in scientific explanations. And while the account does not support a distinctively explanatory role for mathematics, as explanationists would have it, and so in particular does not support explanation-based indispensability arguments for mathematical platonism, it fills an important gap in both platonist and nominalist accounts of mathematical explanations in science by explaining how the use of mathematics contributes to explanatory generality.

CHAPTER 7

Inferentialism, Rigor, and Understanding

7.1 Introduction

In the previous chapter, I discussed one significant way in which RIC can help us better understand the contributions of mathematics to scientific explanations. RIC is a natural framework for describing a metarepresentational role for mathematics, which in turn helps to explain how mathematically formulated scientific explanations can achieve a distinctively high degree of generality.

I now turn from the question of the contributions of mathematics to scientific explanations to the closely related question of its contributions to scientific understanding. Existing work on this relationship has largely focused on the role of mathematical abstractions and idealizations in producing understanding. For instance, [Morrison \[2015\]](#) argues that mathematics contributes to scientific understanding by making possible the introduction of certain abstractions, which in turn are indispensable to understanding certain phenomena. For example, in statistical physics, taking the thermodynamic limit (in effect failing to represent the influence of the number of particles in the system¹) makes possible an understanding of phase transitions that cannot be achieved otherwise.

This is certainly an important type of contribution of mathematics to scientific understanding. But in this chapter, I argue that combining RIC with a prominent inferential account of understanding illuminates further connections between mathematics and scientific understanding. Where the advantages of RIC are clearest are cases in which local features of the mathematical reasoning appealed to in using a representation are crucial to extracting information relevant to understanding. These features may include the inferential affordances of the formalism used, as well as context-specific inference restriction strategies used to confine a more fruitful or user-friendly formalism to contexts in which it is known to be well-behaved. Even in cases in which the informational content of a representation—and so also the explanatory information implied by it—is well-described in terms of a mapping account, understanding how human agents can grasp that information via mathematical reasoning re-

¹Cf. Pincock's discussion of the deep-water idealization in the modeling of ocean waves [\[2012, p. 100\]](#).

quires explicit analysis of the inferential capabilities the mathematics affords to those agents in terms of the particular formalism through which any mathematical structure is presented [cf. Vincent *et al.*, 2018]. And, as I have argued in previous chapters, this is an area in which RIC enjoys a marked advantage over even the most liberal versions of the mapping account.

While I think that the contributions of particular mathematical formalisms to the actual inferential capacities of working scientists are more significant in general than their relatively scant treatment in the literature might suggest, these contributions are absolutely central to cases in which the mathematical techniques applied fall significantly short of the ordinary standards of rigor in mathematics. *Ceteris paribus*, as mathematical rigor increases, it becomes more straightforward to connect the mathematical formalism used to express scientists' mathematical reasoning with an underlying structure that can be used to analyze the informational content of the representation that licenses that reasoning. On the other end of this spectrum, where the formalism requires highly local and ad hoc inference strategies, the value of representing such reasoning in terms of a mapping account is dubious indeed.

In light of this, because the thesis is focused on RIC as a tool to make sense of the use of mathematics in science, I have chosen to focus here on the relationship between mathematical rigor and scientific understanding. In particular, there is an apparent tension concerning this relationship. Both unrigorous techniques and their more rigorous alternatives have been supported—with what seems to be good reason in each case—on the grounds that they promote understanding in a way that alternative techniques cannot. I show how RIC can be combined with a prominent inferential account of understanding to shed light on the multifaceted relationship between the degree of mathematical rigor involved in a representational practice and the scientific understanding it produces. In particular, this framework distinguishes several potentially conflicting ways in which mathematical tools can contribute to scientific understanding. Actual conflict between these contributions then helps to explain the tension concerning rigor and understanding considered above.

I focus on two cases discussed previously in the thesis: the early calculus and Heaviside's

operational calculus. Previous chapters were concerned with the use of these techniques before rigorous alternatives existed, with an eye to explaining how practitioners used them to produce successful representations in the absence of more rigorous alternatives. But my focus in this chapter is on their use after the introduction of these more rigorous alternatives. Each of these cases brings out the apparent tension concerning the relationship between mathematical rigor and understanding mentioned above. On one hand, in each case, unrigorous techniques persisted for a significant amount of time after the introduction of rigorous alternatives, and their proponents justified this at least in part on the grounds that they promoted understanding better than more rigorous alternatives. On the other hand, each ultimately fell out of favor, and a very natural explanation of this is that the more rigorous alternatives ultimately did reflect a better understanding of the relevant mathematical and perhaps even physical domains. Indeed, something close to this seems to be presupposed by the very common approach in philosophy of retrospectively explaining the success of less rigorous mathematical techniques in terms of their relation to more rigorous counterparts.

I argue that this apparent tension can be explained by the sometimes conflicting contributions of mathematical representations to scientific understanding, which applications of unrigorous mathematics bring into sharp relief. Unrigorous techniques typically have the advantage of making more salient inferences practically available to scientists either because no rigorous alternative exists or because such alternatives are cumbersome. They may also make inferences more reliable than their rigorous counterparts by simplifying calculations, thereby leaving less room for human error. On the other hand, they may impede understanding by making the accuracy of representations more difficult to evaluate and by introducing new opportunities for human error in implementing the inference restriction strategies they require.

The study of such cases can contribute to a better understanding of the contributions of mathematics to scientific understanding more generally by helping us to distinguish contributions to understanding that might otherwise be conflated or to recognize contributions that

might otherwise be ignored. The resulting framework for understanding the contributions of mathematics to scientific understanding can in turn be used to make sense of some of the complexities of the history of these techniques. The adoption of rigorous alternatives to unrigorous techniques generally does not happen immediately or uniformly, with unrigorous techniques persisting in various forms, to various degrees, in various contexts. The relative value of the contributions to understanding identified by this framework varies significantly depending on one's epistemic and other aims. As a result, very different tradeoffs between the benefits of rigorous and unrigorous techniques will be called for in different contexts. The framework I propose offers a natural way to identify the factors that warrant these tradeoffs in contexts in which rigorous and unrigorous techniques are both available.

In the next section, I present an account of scientific understanding and discuss an initial way of making sense of contributions of mathematics to understanding in terms of this account. I then discuss several general ways in which applications of unrigorous mathematics complicate this picture. In section 7.3, I substantiate these claims by discussing two case studies that exemplify these patterns: the use of intuitive infinitesimals in the calculus and of Heaviside's operational calculus after rigorous counterparts to those techniques had been developed. Finally, in section 7.4, I discuss the implications of this work for a more general account of the contributions of mathematics to scientific understanding as well as for work on how understanding can be promoted by representations that in one way or another get things wrong.

7.2 Understanding and unrigorous mathematics

7.2.1 Scientific understanding

The topic of scientific understanding has recently become the subject of a very active literature. (See, for example, the contributions to [de Regt *et al.*, 2009].) But for present purposes, we can sidestep many of the most contentious issues in this literature, focusing instead on

three features of understanding about which there is widespread agreement.

The first is that understanding is closely related to scientific explanation in the sense that understanding is what is produced when an agent grasps a good scientific explanation. This is, of course, very imprecise, and there is disagreement in the literature, for example, over whether explanation or understanding is prior², what this grasping consists in³, and even whether the presence of an explanation is necessary for understanding⁴. But there is widespread agreement at least that scientific understanding is the sort of thing brought about by grasping a good scientific explanation in the right sort of way.

The second is that understanding essentially involves an understanding agent, their context, and particularly their actual abilities. A scientific representation may encode an explanation of a phenomenon without that explanation being accessible in practice to the scientists making use of that representation. In such a case, we can rightly say that there is an explanation of the phenomenon⁵ without understanding—at least the kind of understanding that that explanation would produce if grasped.⁶ Likewise, even if explanatory information is more easily inferred from a representation, individual scientists working with that representation might still fail to possess the understanding produced by that explanation if they do not actually infer that explanatory information from it.

Finally, understanding a phenomenon requires scientists to get at least something right

²For instance, [Friedman \[1974\]](#), [Salmon \[1984\]](#), [Humphreys \[2000\]](#), [Woodward \[2003\]](#), [Strevens \[2008\]](#), [Wilkenfeld \[2013\]](#), [Potochnik \[2017\]](#), and [Wilkenfeld & Lombrozo \[2020\]](#) all in one way or another use scientific understanding to explain features of explanation. In contrast, [Khalifa \[2012, 2017\]](#) argues that the notion of understanding doesn't add anything to existing discussions in terms of scientific explanation; understanding a phenomenon just amounts to knowing an explanation of it. And [de Regt \[2017\]](#) analyzes scientific understanding in terms of explanation; having scientific understanding of a phenomenon amounts to having an explanation that is empirically adequate, consistent, and based on an intelligible theory.

³For instance, [Kvanvig \[2003, p. 192\]](#) characterizes this grasping in terms of coherence theories of justification: "Understanding requires the grasping of explanatory and other coherence-making relationships in a large and comprehensive body of information." Others characterize it as a kind of knowledge-how—for instance, of how to evaluate explanations [[Khalifa, 2013](#)], of how to reason counterfactually [[Grimm, 2006](#)—or a more general kind of epistemic ability [[Elgin, 2017](#), [Hills, 2016](#)]. (But for reason not to characterize understanding in terms of knowledge-how, see [[Sullivan, 2018](#)].) [Strevens \[2017\]](#) goes so far as to reject the demand to give a philosophical account of grasping, writing that such an account "would be an extraordinary thing" (p. 41).

⁴[Lipton \[2009\]](#) argues that it is not.

⁵The pragmatic account of explanation advocated by [van Fraassen \[1980\]](#) is a notable exception.

⁶For instance, [de Regt \[2017\]](#) emphasizes the importance of intelligibility, and a number of philosophers emphasize grasping an explanation as a distinctive epistemic ability. See footnote 3.

about it. There is disagreement about how stringent this requirement must be⁷, but for the purposes of this chapter it is important that understanding is both consistent with getting some features of the phenomenon wrong (through idealization, for example) and nonetheless still responsive to the way the phenomenon actually is in at least a minimal sense.

In the rest of this chapter, I will assume a particular view of scientific understanding: the *factive inferentialist* view proposed by Kuorikoski & Ylikoski [2015]. The central idea is that understanding consists in an (inferential) ability, rather than a kind of knowledge (of an explanation). The view is *inferentialist* in the sense that scientific understanding of a phenomenon P consists in the ability to make appropriate counterfactual or “what-if” inferences about P . This serves to capture the first two platitudes above. It captures the first because these inferences concern the sort of information that is commonly thought to be involved in good scientific explanations—namely, information about dependencies captured via counterfactual reasoning. It captures the second because it is concerned with scientists’ actual inferential abilities, rather than inferences they could only make in principle. The view is *factive* in the sense that the relevant inferences must be correct. This still allows for getting some features of the phenomenon wrong; the precision and extent of the what-if inferences one can make about a phenomenon determines the *degree* of one’s understanding of it.⁸ And so this captures the final platitude.

While I focus on Kuorikoski and Ylikoski’s exposition of the factive inferentialist view here, due to its explicitness and clarity, something like the factive inferentialist view is arguably implicit in much work on the counterfactual approach to explanation developed by Woodward [2003] on the basis of work by Salmon [1984] and refined by a number of more

⁷This is often framed in terms of the question of whether understanding is “factive” in the sense that knowledge is. (That is, one can know that p only if p is true.) Some say it is [e.g. Grimm, 2006, Hills, 2016, Le Bihan, 2021]. Others say it isn’t [e.g. Elgin, 2017, Riggs, 2009, Zagzebski, 2001]. But this simple binary disguises a spectrum of views concerning how much of a role falsehoods can play in understanding, with some, like Elgin [2017], emphasizing the role of “felicitous falsehoods” in directly contributing to understanding and others requiring such falsehoods to play a more peripheral role [e.g. Kvanvig, 2003].

⁸Kuorikoski and Ylikoski are not clear about whether the precision and extent of these inferences collapse down to a scalar quantity or map onto different aspects of understanding. I suspect the latter is a more useful way of thinking about these things, but this should ultimately make no difference to the discussion in this chapter.

recent authors, particularly to make it compatible with important examples of putative non-causal explanations in science [e.g. Bokulich, 2008a, Jansson & Saatsi, 2019, Reutlinger, 2018, Saatsi & Pexton, 2013]. Accounts of this kind explain scientific explanation in terms of a certain kind of inference (concerning the counterfactual behavior of the target system), which track genuine features of the target system (“dependence relations”). In this way, both the inferential and factive aspects of the view are present, at least in embryonic form, in the counterfactual approach to scientific explanation. And so, while I do frame the rest of the discussion in this chapter in terms of the particular view proposed by Kuorikoski and Ylikoski, their view is one that falls very naturally out of what is arguably the most prominent account of scientific explanation currently on the market. This—together with its emphasis on inference, which makes it a natural ally of RIC—makes it a very natural choice for my purposes here.

Before moving on, it is also worth pausing to distinguish between two kinds of understanding that might be related to mathematical rigor. First, the kind of understanding Kuorikoski and Ylikoski’s factive inferentialist account is intended to capture is understanding of a phenomenon. This sort of understanding typically depends on a model or theory to help produce understanding, but the object of understanding is ultimately not the model or theory itself. Alternatively, we might be interested in how mathematical rigor relates to understanding of mathematically formulated theories or models themselves, or even understanding of the more abstract mathematical tools used to formulate them. In this case, the factive inferentialist account’s restriction to counterfactual inference makes it less plausible. While some, notably Baron and various collaborators, have tried to apply the counterfactual approach to explanation to mathematical explanations, the counterfactual approach is at the very least less immediately plausible when applied to mathematics than when applied to physical phenomena.

There are a few ways around this. The first would be to pitch the discussion in the rest of this chapter solely in terms of understanding of phenomena. And indeed, I think the central

points of this chapter could be framed in this way. However, I think there is more to be said about the relation between mathematical rigor and intramathematical understanding (and understanding of a model or theory more generally), and doing so requires relatively minor modification to the framework sketched so far. Kuorikoski & Ylikoski [2015] and Kuorikoski [2021] assimilate understanding of models and theories into the counterfactual framework by framing such understanding in terms of counterfactual inferences about how the models would behave if they were changed—say, to represent some related but distinct target system.⁹ Alternatively, we might not require the relevant inferences to be counterfactuals in the case of intramathematical understanding, but instead characterize understanding-relevant inferences in other terms. In either case, the basic structure of the factive inferentialist account—in particular, the potential for conflict between and within the inferential and factive aspects of understanding—is preserved. As a result, intramathematical understanding and understanding of models should be expected to raise philosophical issues that are substantially similar to those raised by understanding of phenomena, and these similar issues should be given broadly similar treatment.

The factive inferentialist framework already suggests two broad ways in which mathematics—or indeed any other scientific tool—might contribute to scientific understanding. First, corresponding to the *inferential* aspect of understanding, it might make more inferences practically accessible to scientists. It might do so by making available new representations of a phenomenon, which support new inferences. Or, as Kuorikoski and Ylikoski emphasize, it might do so by making existing representations more tractable, thereby making more inferences available in practice. Second, corresponding to the *factive* aspect of understanding, mathematics might make available inferences more reliable. According to Kuorikoski and Ylikoski, it might do so either by making it easier to evaluate the accuracy of inference-supporting representations (by forcing assumptions to be made explicit in the formalization process) or by reducing human error in carrying out the inferences licensed by a representation (due to

⁹In essence, this is very similar to the move of treating mathematical knowledge (and other knowledge of necessary truths) as metalinguistic knowledge [Rayo, 2009, Stalnaker, 1987].

the existence of well-defined rules for doing so either mathematically or computationally that allow the relevant inferences to be “externalized”).

This is a useful way to start thinking about the contributions of mathematics to scientific understanding, but we will see that the picture must be more complicated than this. In particular, examination of unrigorous mathematical techniques shows how conflicts can arise between the contributions of mathematics to both the inferential and factive aspects of understanding. Most intuitively, a piece of applied mathematics may contribute to the inferential or factive aspect of understanding at the expense of the other. But, as I will try to show, a piece of applied mathematics may also involve tradeoffs between conflicting contributions to a single aspect of understanding—making some range of inferences more reliable at the expense of others or expanding the range of available inferences in one way while restricting it in another.

7.2.2 Unrigorous mathematics

Recall from chapter 5 that applications of unrigorous mathematics require the use of a distinctive, inferentially restrictive methodology, according to which some classically available inferences are not permitted. Problematic concepts are thereby quarantined to contexts in which they exhibit the desired behavior without allowing undesirable results to be derived. The introduction of inferentially restrictive methodologies already complicates the picture of the contributions of mathematics to understanding sketched above.

First, such methodologies are one means by which a piece of applied mathematics might make conflicting contributions to the inferential aspect of understanding. On one hand, such methodologies are only worth using when they expand the range of inferences made practically available by the representation in which they are used—either because no more rigorous alternative exists or because such rigorous alternatives are more difficult to work with and so make fewer inferences accessible in practice. On the other, by their very nature, inferentially restrictive methodologies restrict available inferences. While such restrictions are necessary

to ensure that the allowed inferences are reliable, they may nonetheless prevent the relevant mathematical techniques from being extended to a wider range of cases, to which their more rigorous counterparts can profitably be applied with minimal adjustment. So a piece of mathematics that requires an inferentially restrictive methodology might make more inferences practically available in the limited domain in which it applies than its more rigorous counterpart, while making fewer inferences practically available in a more general domain than that rigorous counterpart.

Which technique better promotes understanding in a given context then must depend both on which range of inferences are of the greatest significance, as well as the degree of fluency the understanding agent has with the two techniques, which will determine how much of an advantage the less rigorous technique enjoys in the restricted domain and the more rigorous technique enjoys in the more general domain. Beyond this, the level of grain of the relevant inferential restrictions matters a great deal. If these restrictions are fine-grained and ad hoc, as in the case of Heaviside's operational calculus, this may make it more difficult to apply unrigorous techniques even in the limited domain to which they apply by making it difficult for those using the technique to know when the inferential restrictions do and do not apply, thereby making fewer inferences practically available even within this restricted domain. (And, in fact, Heaviside's operational calculus was notoriously difficult for students to master.) On the other hand, if the restrictions are relatively straightforward and general, as in the case of the calculus using intuitive infinitesimals, such problems may arise to a lesser extent or even not at all.

Second, such methodologies are also a means by which a piece of applied mathematics might make conflicting contributions to the factive aspect of understanding. By allowing for simplified inferential procedures in restricted domains, they build on the way mathematical techniques in general reduce human error by allowing inference procedures to be externalized. On the other hand, they have the potential to make it less clear which inferences a given representation actually licenses, again to the extent that the required inferential restrictions

are local and ad hoc. This introduces a new source of human error: error about which inferential restrictions are required. This in turn has the potential to make it more difficult to evaluate the representation for accuracy.

In the next section, I discuss two cases that exemplify these features that complicate our picture of the contributions of mathematics to scientific understanding. I try to show that this more complicated picture ultimately helps us better understand these historical episodes—not just why unrigorous techniques persisted after the introduction of more rigorous alternatives, but why they persisted in the forms and contexts in which they did, and why they ultimately fell out of favor (when they did).

7.3 Case studies

7.3.1 Unrigorous infinitesimals

Before the calculus was put on a more rigorous footing by the work of Cauchy, Weierstrass, Dedekind and others in the 19th century, notions central to calculus and real analysis—continuity, differentiability, and so on—were understood in terms of infinitesimals or other similarly hazy notions.¹⁰ In chapter 4, I discussed one simple way this was done in terms of inconsistent, “naïve” infinitesimals. A derivative dy/dx is treated as a ratio of infinitesimally small quantities, understood in such a way that $0 < |dx| < r$ and $0 < |dy| < r$ for all $r \in \mathbb{R}^+$. These infinitesimals can then be manipulated as if they were zero or as if they were non-zero in particular contexts. To prevent the derivation of contradictions or other undesirable results, an inferentially restrictive methodology was adopted. Infinitesimals could be used only in the context of particular algorithms for calculating derivatives and integrals. And at any given point in these algorithm, infinitesimals were to be treated as if they were zero or as if they were non-zero, but never both. These restrictions ensured that no direct contradiction could be derived from the contradictory properties of naïve infinitesimals, but useful results could

¹⁰For a useful summary of much of this history, see [Kline, 1972]. For resources more focused on historical infinitesimals and related work, see [Bair *et al.*, 2013] and [Kleiner, 2002].

still be achieved.

Infinitesimals persisted, albeit in more sophisticated forms, even as mathematical understanding of the calculus became increasingly advanced. For instance, much of Cauchy's foundational work—including his construction of the reals in terms of Cauchy sequences and an important precursor to Weierstrass's epsilon-delta definition of limit, a cornerstone of 19th-century efforts to put the calculus on a rigorous footing—was formulated in terms of infinitesimals, this time understood not as constant quantities but as functions that approached zero as their arguments did [Cauchy, 1821]. This notion of infinitesimal, sometimes called the *dynamic limit concept*, vindicated earlier algorithms for calculating derivatives and integrals, provided the relevant infinitesimals were *proper infinitesimals*—infinitesimals that are non-zero when their argument is close to, but not equal to, zero—as well as in a number of other contexts [Tall, 1981]. But, although it worked in a wide range of contexts and even helped to explain why earlier algorithms worked when they did (and didn't when they didn't), it too was not suitable for the whole of calculus and real analysis.

In contrast, the modern calculus, built on infinitesimal-free foundations largely through the work of Weierstrass, treats a derivative dy/dx not as a ratio at all, but as a *differential operator* d/dx applied to a function y . This too vindicated earlier algorithms for derivatives and integrals, as well as the results of dynamic limit techniques, but could be used to formulate definitions of core concepts like continuity in a way that allowed for the treatment of a much broader range of functions. This generally comes at the cost of making proofs and calculations more complex, and, due to the epsilon-delta definition of limit, these proofs and calculations generally require at least one more level of quantifier nesting than their counterparts appealing instead to infinitesimals. For instance, proving the chain rule

$$\frac{dy}{dz} \frac{dz}{dx} = \frac{dy}{dx}$$

is trivial in terms of naïve infinitesimals and almost trivial in terms of the dynamic limit

concept, but decidedly non-trivial in terms of the modern calculus. Similar inferences abound in physics. For instance, the following step in proving the Work-Energy Theorem (for the special case of a particle moving in a straight line)

$$m \int_{t_1}^{t_2} v \frac{dv}{dt} dt = m \int_{v_1}^{v_2} v dv$$

is similarly trivial in terms of infinitesimals, justified by canceling the dt s, but requires other justification in terms of the modern calculus.

Because such cases were widespread, indeed covering most if not all of the use cases of the early calculus, infinitesimal techniques persisted in several forms and in several contexts well after the core results in the foundations of calculus and real analysis had been established in the mid-19th century. In the rest of this section, I consider two such examples: ongoing heuristic use of unrigorous infinitesimals by physicists and the persistence of unrigorous infinitesimals and related notions in introductory calculus classes well into the 20th century¹¹.

The switch to a modern, “epsilonic” understanding of the calculus seems to have required little change in the practice of working physicists. This is because infinitesimal and dynamic limit methods produce the same results as more rigorous methods in most cases of interest. The functions needed for classical physics were relatively well behaved, and so the increased generality of the modern calculus was unnecessary. On the other hand, 20th-century physics produced a number of integral concepts that were so ill-behaved that they required a radical departure from the integral concept of the modern calculus anyway.¹² What this meant was that physicists could embrace the modern calculus as the standard against which the correctness of proofs should be measured, while continuing to use infinitesimals in a heuristic way—for “back-of-the-envelope” calculations. This seems to be the case even today. Indeed, anecdotally, it seems not at all uncommon for a student to watch their calculus lecturer carry

¹¹Importantly, here I do not have in mind the persistence of infinitesimals in rigorous, twentieth-century theories like non-standard analysis [Robinson, 1966] or smooth infinitesimal analysis [Bell, 2008]

¹²See the discussion of path integrals in §5.3. A closely related, but earlier example is the Wiener integral.

out a rigorous epsilon-delta proof of the chain rule and then watch with consternation (or glee!) as their physics lecturer simply crosses out the dt s to justify the similar step in deriving the Work-Energy theorem. As Steiner puts it,

Even in the twentieth century, physicists persist in thinking of the derivative intuitively, as a quotient of “infinitesimals.” They think of integration as a summation of infinitesimal area (or volume) elements. In line integration, they think of a curve as the union of infinitesimal straight line segments, etc. Similarly, they talk of infinitesimal rotations, rather than tangent spaces of Lie groups. It was only recently that even this talk of infinitesimals was legitimized rigorously by Robinson’s work, and even now, so far as I know, nonstandard analysis has not yet been brought to bear on the infinitesimal rotations. [Steiner, 1992, p. 160]

The framework sketched in the previous section provides a natural way to make sense of why this should be: the mere heuristic use of infinitesimal techniques preserves the benefits with respect to understanding of both approaches. In simple cases, in which the relevant functions are well-behaved, infinitesimal techniques are known to produce the same results as epsilon-delta techniques. In these cases, a heuristic use of infinitesimals saves time and so both makes more inferences practically accessible and reduces the potential for human error due to the simplicity of the calculation procedures relative to epsilon-delta techniques. But treating these uses of infinitesimals as merely heuristic allows one to achieve the benefits of epsilon-delta techniques. Because these techniques apply to a broader class of functions, a wider range of inferences is practically available to scientists when such functions are needed to represent physical phenomena. And because these techniques are the ultimate arbiter of proof correctness, their benefits with regard to reliability of inference are preserved as well. So this heuristic use of infinitesimals in physics is exactly what the framework sketched previously would lead us to expect.

In contrast, controversy over the use of infinitesimals in mathematics education lasted well into the 20th century. For instance, in a 1926 lecture on introductory calculus courses

delivered to the Association of German Scientists and Physicians, Otto Toeplitz characterized the situation as one in which there is “polar opposition” between pro- and anti-rigorist standpoints, reflected in a “colorful diversity of approaches” at the university level [Toeplitz, 2015, p. 297]. At the rigorist end of the spectrum, one begins with a six-week intensive treatment of the foundations of the calculus, from which the concrete rules of the calculus are finally derived and only then used. At the other end of the spectrum, “there is the intuitive path in which the magic of differentials is made to rule, and in which even by the last hour of a two-semester course the fog which arises from indivisibles remains undispelled by the sunshine of a clear limit concept” [Toeplitz, 2015, pp. 297f]. Similar examples abound, particularly concerning the teaching of calculus in secondary schools.¹³ In much of Europe the introduction of epsilontic rigor in secondary-level calculus courses was a project undertaken in earnest only after the Second World War [Zuccheri & Zudini, 2014].

Why should such controversy have persisted for as long and as widely as it did? Again, thinking about the contributions of each approach to student understanding sheds considerable light on the controversy. Here the distinction mentioned previously between understanding a model and understanding a phenomenon becomes quite important, as the various parties to these controversies were concerned not just with the ability of students to use the calculus to better understand phenomena as future scientists and engineers, but also with students’ understanding of the calculus as such, especially in relation to their preparation for courses in more advanced mathematics. (Toeplitz, for example, was primarily concerned with the latter.) In what follows, I will assume that some version of the factive inferentialist approach to understanding applies to intramathematical understanding of the calculus, either formulated in terms of metarepresentational counterfactuals as Kuorikoski and Ylikoski suggest or in terms of some non-counterfactual variety of inference. But if for whatever reason the reader balks at both of these options, a more limited explanation can be provided solely

¹³See, for example, Roquette’s [2010] account of intuitive infinitesimals in his early calculus education, which primed his later interest in Robinson’s work on non-standard analysis, culminating in a collaboration between the two mathematicians.

in terms of the inferences made available by the two approaches to the calculus in applied contexts.

Proponents of less rigorous approaches to calculus education emphasized the importance of intuitive notions that allowed students to work immediately. On the framework sketched above, this naturally contributes to understanding by making more inferences practically available in the limited domain to which the less rigorous techniques apply. And these techniques are perfectly reliable when applied to the sort of problem students are faced with in an introductory calculus course. So the intuitive and unrigorous approach produces greater understanding in students initially.

But such techniques apply only in a limited domain, and they don't prepare students for later courses in which more rigorous techniques are used to solve a broader range of problems about a broader range of functions. In contrast, epsilon-delta techniques, once mastered in the simple case, can naturally be extended to this broader space of problems. At this point in a student's mathematical education, then, epsilon-delta techniques allow them to reliably make a broader range of inferences, despite the advantages of infinitesimal techniques in the more limited domain to which they apply. And so it seems that while infinitesimal techniques indeed promote student understanding better than epsilon-delta techniques early in students' mathematical education, epsilon-delta techniques are required at a certain point for students to advance any further in their understanding. And avoidance of epsilon-delta techniques for too long would seem to draw out the processes by which students come to gain this greater degree of understanding. But it is far from clear at which point the gains in understanding achieved via one of these approaches begin to outweigh the potential gains in understanding via the other approach. We should therefore expect there to be significant and ongoing controversy concerning at what point in students' mathematical careers (and to what degree) to begin introducing epsilon-delta methods.¹⁴

¹⁴Toeplitz's own suggestion, echoed by at least some more recent educators [Can & Aktas, 2019], is that teaching should roughly follow the history of the development of the calculus, both to balance the virtues of each approach and to introduce students to the beauty and drama of working out new mathematics.

Before moving on, it is worth pausing to consider whether it is appropriate to consider students' understanding of the calculus as part of a discussion of scientific understanding. Do we have different types of understanding in pedagogical and research contexts? On the one hand, students' activity in pedagogical contexts and scientists' activity in research contexts certainly are importantly different, and a complete account of scientific understanding should not conflate them. On the other hand, I take it as a strength, not a weakness, of the approach that I propose that it can be used to make philosophical sense of cases drawn from both pedagogical and research practices, regardless of whether we choose to think of these as distinctive kinds of understanding.

7.3.2 Heaviside's operational calculus

Recall from §5.2 that Heaviside's operational calculus was an algebraic technique for solving physical problems, especially concerning electro-magnetism, represented in terms of differential equations. In many cases, this made it possible to find solutions to important problems in applied contexts (such as the design of telegraph lines) that would otherwise have been intractable. Heaviside's techniques were notorious for their lack of rigor, something Heaviside openly embraced. But they were ultimately replaced by more rigorous techniques, such as the use of the Laplace transform, which allowed for a similar algebraic approach by transforming the problem from the time domain to the complex-frequency domain.

Heaviside's operational calculus persisted for several decades after Bromwich introduced the Bromwich integral as a way to calculate the inverse Laplace transform, allowing a rigorous treatment of the problems that the operational calculus had been developed to solve. However, unlike unrigorous infinitesimals, it ultimately fell entirely out of favor. Again, the framework for understanding mathematical contributions to understanding sketched above allows for a very natural explanation of this difference.

Let's begin with why the operational calculus persisted as long as it did. As one of Heaviside's most ardent critics put it,

[A]s a matter of practical convenience there can be no doubt that the operational method is far the best for dealing with the class of problems concerned. It is often said that it will solve no problem that cannot be solved otherwise. Whether this is true would be difficult to say; but it is certain that in a very large class of cases the operational method will give the answer in a page when ordinary methods take five pages, and also that it gives the correct answer when ordinary methods, through human fallibility, are liable to give a wrong one. [Jeffreys, 1927, p. v]

In other words, the operational calculus makes more inferences practically available by greatly simplifying the calculations required to make them, and it improves their reliability when carried out by human beings for the same reason. As in the case of unrigorous infinitesimals, the operational calculus could even be used as a heuristic, with Laplace-transform-based techniques explicitly endorsed as the standard by which inferences should ultimately be judged correct or incorrect. Bromwich himself suggested as much in a letter to Heaviside:

After coming back to these questions after 2½ years of war-work, I found myself able to work more readily with operators than with complex-integration. [...] I at once saw that I must make the operator-method take the leading place: and complex-integrals have accordingly been pushed into footnotes. I still regard the complex-integral as a useful method for convincing the purest of pure mathematicians that the p -method rests on sound foundations: but I am sure that the p -method is the working-way of doing these things. (Letter to Heaviside on 5 April, 1919, quoted in Nahin [2002, p. 229])¹⁵

Heaviside thought his approach had other benefits in relation to understanding. For example, he emphasized the importance of keeping the physical dimension of the problem in

¹⁵Bromwich was not rewarded for his kindness. Heaviside responded,

I rejoice to know that you have seen the simplicity and advantages of my way [...]. Now let the wooden headed rigorists go hang, and stick to differential operators and leave out the rigorous footnotes. It is easy enough if you don't stop to worry. [...] I never could stomach your complex integral method. (Letter to Bromwich, 7 April, 1919, quoted in Nahin [2002, pp. 229f])

mind at all times:

The practice of eliminating the physics by reducing a problem to a purely mathematical exercise should be avoided as much as possible. The physics should be carried on right through, to give life and reality to the problem, and to obtain the great assistance which the physics gives to the mathematics. [...] No mathematical purist could ever do the work involved in Maxwell's treatise. [Heaviside, 1899, §224, pp. 4f]

And he has a point here insofar as the target of understanding is the physical domain rather than the mathematical one. Doing so has the potential to increase one's stock of practically available physical inferences at the cost perhaps of decreasing one's stock of mathematical inferences. Second, Heaviside thought that student understanding was better promoted by experiential familiarity with using mathematical tools than by learning rigorous definitions. The latter, he thought, could hinder students' ability to apply the techniques. For example, he thought it undesirable to try to rigorously state his expansion theorem, one of the pillars of his operational calculus:

Now it would be useless to attempt to state a formal enunciation to meet all circumstances. Even supposing that an absolutely perfect knowledge of the subject made it possible to do so, it would be very unpractical. It would be worse—far worse—than that very lengthy enunciation of a theorem in the 5th Book of Euclid, which may be read and re-read fifty times without properly grasping its meaning, which is not much after all; only something in compound proportion that the modern schoolboy does in a minute or two. It is better to learn the nature and application of the expansion theorem by actual experience and practice. A theorem which has so wide an application is a subject for a treatise rather than a proposition. [Heaviside, 1899, §282, p. 128]

Again, he has a point here insofar as that improves students' practical ability to reason with

the relevant mathematics. (Unfortunately for Heaviside, that does not seem to have been the case. The operational calculus was notoriously difficult for students to learn, as discussed below.)

Why then did electrical engineers stop using Heaviside's operational calculus, even as a heuristic, in favor of Laplace-transform-based techniques? I think we can largely trace this back to a difference in the inferentially restrictive methodologies used in applying the calculus and the operational calculus. In the former case, there are well-defined algorithms that determine how one can reason with the mathematically problematic properties of infinitesimals. On the other hand, the inferential restrictions used in applying the operational calculus were piecemeal, local, and ad hoc. This inferential strategy seems to have been very successful in the hands of Heaviside, who was thereby able to work more flexibly—for instance, by allowing the physical features of a problem to play a heuristic role in the development of the mathematics used to solve it—without deriving undesirable results. But this led to two significant problems.

First, the ad hoc nature of Heaviside's inferentially restrictive methodology made it difficult to know in advance what inferences were actually licensed by a representation formulated in terms of the operational calculus. This has implications for both the inferential and factive aspects of understanding. Regarding the inferential aspect of understanding, making it more difficult to tell whether an inference is licensed by a representation directly increases the cognitive cost of making that inference and so can be expected to reduce the range of inferences that are available in practice. Regarding the factive aspect, if it is unclear what inferences are licensed by a mathematical representation, then it will be difficult for that reason to evaluate the accuracy of that representation. This then undermines one of the usual contributions of mathematics to the factivity of understanding identified by [Kuorikoski & Ylikoski \[2015\]](#): making a model's accuracy easier to evaluate by forcing its assumptions to be made explicit.

Second, this made the operational calculus notoriously difficult for students to learn. As one historian of the operational calculus put it,

Heaviside did not systematically and rigorously develop his calculus from firmly grounded postulates, but rather (as he himself testifies) developed it empirically as he worked, points in doubt to him being settled by checking known exact solutions against corresponding operationally obtained solutions. In consequence, his expositors, lacking Heaviside's intimate and hard-won knowledge of the vicissitudes of his operational calculus, produced books that comprise little other than stated rules of manipulation and examples illustrating use of these rules. Accordingly, *even the most zealous student of these books scarcely can gain real insight to Heaviside's operational calculus; become familiar with the especial limitations, pitfalls, and other shortcomings peculiar to it; or attain marked proficiency in the accurate use of it; especially for the solution of continuous systems characterized by partial differential equations.* In fact, it was realization of these facts plus a desire to establish rigorously certain useful formulas and rules of procedure advanced by Heaviside that sparked much of the early work of those who developed and advanced the other methods of operational calculus. [Higgins, 1949, p. 43, my emphasis]

In contrast, infinitesimal techniques were easy for students to pick up not just because the concepts were intuitive, but because there were a few relatively simple algorithms precisely delineating their correct use in the limited range of contexts that students would encounter in an introductory course. There are no such simple, well-defined rules for the operational calculus, but in the best case complex sets of algebraic rules that couldn't be derived from simpler principles. Again, this has implications for both the inferential and factive aspects of understanding. This again makes fewer inferences practically available to students. But it also makes the available inferences less reliable by introducing a new source of human error, misjudging whether the inferentially restrictive methodology allows for a given inference.

In the first decades after the introduction of techniques based on the Laplace transform and its inverse, these problems did not decisively favor those techniques because they too were

unwieldy and similarly prone to human error. Engineers seem to have turned away from the operational calculus only with the development and popularization of significant shortcuts in applying Laplace-transform-based methods. In particular, engineering textbooks increasingly taught students to calculate Laplace transforms not from scratch—often a daunting task—but by using tables listing the Laplace transforms and inverse Laplace transforms of common functions, from which the transforms of less common functions could be derived.¹⁶ This was the most laborious and error-prone part of the process; once one had moved from the time domain to the complex frequency domain, one could work purely algebraically, just as in the operational calculus. And so the more rigorous techniques were increasingly on a par with the operational calculus in terms of the simplicity of derivations without introducing unclarity about the inferences they licensed and without the pedagogical disadvantages of the operational calculus. As a result, they seem to have gained an advantage with respect to both the inferential and factive aspects of understanding. So, again, on the framework sketched in this chapter, we should expect Heaviside’s operational calculus to have fallen out of favor.

It is worth mentioning here that the case of Heaviside’s operational calculus is atypical, in that Heaviside’s unrigorous techniques were ultimately fully replaced by more rigorous techniques—albeit ones heavily inspired by Heaviside’s work. Much more typical is a case like the continued use of intuitive infinitesimals informally in the practice of physics. For example, the introduction of distribution theory to replace the use of the Dirac delta function resulted in almost no change to the practice of physics. Physicists could continue to use the Dirac delta function as before, sometimes simply adding footnotes referring to Schwartz’s [1945] work on distributions to justify their continued use of the Dirac delta.

¹⁶For a representative example, see [DeCarlo & Lin, 2009, p. 564].

7.4 Implications and further work

So far, I hope to have shown that RIC's account of representation together with an inferentialist account of understanding can shed significant light on the multifaceted relationship between mathematical rigor and scientific understanding both in general and in relation to concrete historical episodes in which rigorous and unrigorous techniques coexisted. By way of conclusion, I will now briefly discuss the broader implications of this work.

First, when rigorous and unrigorous techniques coexist, the tradeoffs between conflicting mathematical contributions to understanding are thrown into sharp relief. But, once we recognize these potentially conflicting contributions, we should expect to find similar tradeoffs even in contexts where the mathematical tools are not particularly unrigorous. In light of this, a promising direction for further work is to explore a more general, understanding-based approach to the phase of applications of mathematics in which the appropriate mathematical tools are chosen or developed.

Second, the account of rigor and understanding presented here is consonant with work on the relation between understanding and other ways of "getting things wrong," such as abstraction and idealization.¹⁷ This provides further support for the framework developed in this chapter; a virtue of my account is its coherence with similar solutions to similar philosophical problems concerning understanding. But this connection also suggests directions for further work. What it suggests is that the rigorous-unrigorous spectrum is another variable of the same kind as the spectrum between idealized and deidealized or abstract and concrete representations to be used in thinking about how particular kinds of models produce understanding. This kind of generalized account could then be used to think about a broader range of cases—including in particular cases in which scientists must make tradeoffs between, say, incorporating explicit idealizations and applying mathematical shortcuts to a less idealized

¹⁷For the relationship between specifically mathematical abstractions and understanding, see [Morrison, 2015]. It has become relatively common to explain the epistemic benefits of idealizations in terms of their contribution to understanding [e.g. Bokulich, 2008b, Elgin, 2017, Potochnik, 2017, Strevens, 2017]. For arguments against the importance of idealization to understanding, see [Sullivan & Khalifa, 2019].

model.

Finally, in this chapter I have assumed that some analogue of the factive inferentialist account of understanding a phenomenon also applies to understanding a model or theory and to intramathematical understanding in particular. While I regrettably have not been able to develop or defend such an account here, the fruitfulness of this assumption, particularly in the case of 20th-century controversies over rigor in calculus education suggests that such an account is well worth pursuing.

CHAPTER 8

Conclusion

8.1 Contributions of the thesis

In the thesis, I have developed and defended a novel, inferentialist approach to mathematical scientific representation, defending it on three grounds. First, it can be applied to a wider range of cases than existing versions of the mapping account, though the extent to which this is so depends on the version of the mapping account in question. Second, it does better than existing versions of the mapping account as a meta-level representational device for representing philosophically salient aspects of scientific practice; it can recover the work done by mapping accounts in cases well treated in those terms, but in a number of cases it can be used to much more perspicuously represent features of scientific practice that are of philosophical interest. Third, it is more manifestly neutral with regard to the nature of mathematics than mapping accounts; it requires only that mathematics involves something that *looks like* inference. (Chapter 3)

I spent most of the thesis substantiating the second of these points, which I take to be the most important, as well as the strongest support for RIC. I argued that RIC is a better meta-level tool for thinking about episodes in which inconsistent (chapter 4) or otherwise unrigorous (chapter 5) mathematical techniques are applied. I extended RIC to make sense of a distinct *metarepresentational* role of mathematics in scientific explanation, which allowed for a middle ground between explanationist and representationalist accounts of mathematical explanations in science, as well as a novel response to indispensability arguments (chapter 6). And I used RIC to help make sense of an apparent tension concerning the relationship between mathematical rigor and scientific understanding (chapter 7).

In so doing, I believe I have made four major contributions to the existing literature on mathematical scientific representations and related issues. First and most importantly, I have developed and defended what I believe to be the first viable, fully worked-out alternative to mapping accounts in the literature on mathematical scientific representations. This then made possible a novel argument against mapping accounts on the grounds that they are less

successful as meta-level representational devices than RIC. Second, I have provided several detailed case studies of cases underappreciated in the literature on mathematical scientific representations, including applications of the inconsistent early calculus, Heaviside’s operational calculus, and path integrals. Third, I have argued for a new approach to mathematical explanation with implications for traditional problems, including non-causal explanations and indispensability arguments. Finally, I have extended work on the relationship between understanding and various features of representations that “get things wrong”—like abstractions and idealizations—to a new such feature: the degree of mathematical rigor of the mathematical tools at work in the representation.

8.2 Directions for further work

I conclude by briefly discussing two promising ways in which the work in the thesis might be extended. The first is to extend RIC to treat a wider range of computational techniques in science, particularly those involved in the recent phenomenon of “data-driven science.” The second is to use RIC to inform an account of the semantics of pure mathematics, driven by the thought that the meaning of mathematical concepts is best explicated in terms of their inferential affordances, rather than in terms of set-theoretic or other structures.

8.2.1 RIC, computational inference tools, and data-driven science

[Napoletani *et al.* \[2011\]](#) characterize data-driven science in terms of the use of general-purpose computational techniques to solve highly circumscribed problems by means of general-purpose computational techniques and large data sets, rather than theory-informed models. About these techniques, they write:

In this new framework, mathematics provides powerful ideas and techniques which then generate broad classes of generic computational tools. [...] While in some cases the computational tools may be useful as an intermediate step

between theories and phenomena, [...] more often the way they are applied to the data is leading to a scenario where it is missing any explicit isomorphism between a mathematical structure and the phenomenon under consideration. The computational tools of data analysis therefore do not truly model the phenomenon. (p. 3)

I suspect that RIC can greatly elucidate the role of mathematics and related computational techniques in these cases, which are of great importance to contemporary science. Moreover, I suspect that, contrary to Napoletani et al.'s conclusion, RIC is particularly well-suited to explain the important respects in which these data-driven techniques are not a radical departure from earlier modeling practices.

Such techniques are typically intended to solve either classification problems or regression problems. Both types of problem are that of finding a function that takes as input numerical representations of certain features of the target system (*input variables* or *predictors*) and outputs numerical representations of other features of that system (*output variables* or *responses*), with the aim of making predictions or inferences about the features of the system represented by the output variables on the basis of the (more easily known) features of the system represented by the input variables. In *classification problems*, the output variables are numerical codes representing qualitative features of the target system, such as whether a given email is spam or non-spam. In *regression problems*, the output variables are representations of quantitative measurements of the target system. For example, [Stamey et al. \[1989\]](#) predict the levels of prostate related antigen (PSA) based on other clinical measures in men with prostate cancer. (Cf. [Hastie et al., 2008](#), §2.2.)

What is common to the techniques considered here is that they aim to solve such problems, understood solely as problems of function approximation. The intended result of such techniques is a function that approximates a function with perfect predictive accuracy for the intended domain. While such a function could in principle be constructed by producing a detailed representation of the structure of the relevant phenomenon, these techniques produce

an approximation of such a function by employing algorithms that make use of existing data and nothing (or almost nothing) more. This is useful not only when an adequate structural understanding of a given phenomenon is elusive, but also in the (extremely common) case that adequate representations of the underlying phenomena are complex enough that the problem of generating these predictions is mathematically or computationally intractable. In this sense, these representations play much the same role as models that incorporate idealizations or approximations to make working with the underlying mathematics more manageable. But they achieve this not by deliberately distorting the representation of particular aspects of the target system, but rather by abstracting away from *all* aspects of the phenomenon of interest other than the features represented by the inputs and outputs of the desired function.

A well-known such technique is the use of artificial neural networks. A supervised feed-forward neural network is a tool to specify and compute a function from one data domain to another. It consists of a layer of input nodes, whose values are specified according to the data about some set of features of the phenomenon. These represent salient features of the target phenomenon from which we would like to be able to make predictions. Data are then passed from each node in the input layer to those in one or more hidden layers via weighted connections. The value of each hidden node is a function of the sum of the weighted inputs to that node. In turn, each of these nodes passes on its value, via another weighted connection, to the nodes in the next layer, and the values of these nodes are determined in the same way. The final layer is an output layer, whose values are to correspond to features of the phenomenon that we would like to predict on the basis of the features represented in our input layer. The weights attached to each connection in the network are set by “training” the network on a set of data that includes the features represented by both the inputs and the outputs. This involves adjusting the weights to reduce the discrepancy between the outputs of the neural net and the corresponding training data via an optimization algorithm such as backpropagation.

For example, [Bailer-Jones \[2000\]](#) constructs an artificial neural network to make predictions about stars’ temperature, surface gravity, and composition on the basis of their optical




Figure 8.1: An artificial neural network with two hidden layers.

spectra. The input nodes then are numerical representations of features of stars’ optical spectra, while the output nodes are numerical representations of its temperature, surface gravity and composition (in particular, its metallicity).

Mathematically, the network is simply a function from an input vector \vec{x} , whose components are the input variables, to an output vector \vec{y} , whose components are the output variables. So each component of \vec{x} represents a features of the target system from which we would like to make a prediction, and each component of \vec{y} represents a feature of the target system about which we would like to make a prediction. The hidden layers are also represented by vectors, and the components v'_j of the vectors for the hidden and output layers are given by a function of $\sum_i w_{i,j}v_i$, where each v_i is a component of the previous layer and $w_{i,j}$ is the weight assigned to the connection between v_i and v'_j . Each particular way of setting up such a network—i.e., a specification of how many hidden layer are included, how many dimensions each vector has, and how the values of each component of the hidden layers and and output layer are to be computed—determines a family of functions. The particular setup chosen determines both how complex the behavior of the resulting functions can be and how computationally efficient the training algorithm can be, so the choice of setup depends on how one wishes to balance these competing considerations. Finally, the training algorithm (in the best case) serves to pick out the function from this family of functions that most reliably predicts the right values

for the output variables (and so most reliably predicts the features of the phenomenon one wants to predict). It does this by changing the w_i to reduce the error each time the function is computed for training data, which include the correct values of both the input and output variables.

On the one hand (contra Napoletani et al.), there is a clear sense in which we can understand this in terms of an isomorphism between a mathematical structure and a target phenomenon. The mathematical structure consists solely of the function picked out by the neural network and an appropriate domain—viz. $\mathbb{R}^i \times \mathbb{R}^j$ where i is the dimensionality of the input and j that of the output. Our morphism takes the numerical representations of stellar spectra and other properties to the corresponding properties, and it takes the function to a physical relation R such that Rxy holds iff x is a stellar spectrum, y is a collection of output properties, and all stars of the relevant type with spectrum x have properties y . The assumption that this morphism is, say, an isomorphic embedding could then justify the use of the representation to make inferences about the output properties on the basis of the input properties.

On the other hand, this mapping-based approach does little to illuminate the epistemic features central to and distinctive of such cases. Given the extremely impoverished nature of the physical structure mapped to the neural net, the mapping does no more than encode the inferences we would like to make on the basis of the neural net. Supposing that our mapping is an embedding ultimately amounts to the same thing as supposing that the neural net will yield accurate predictions when we feed it new data of the appropriate kind. The more interesting question is why we should take ourselves to be justified in making these inferences—why we should expect the neural net’s predictions to be accurate. To explain this, we must consider not just the final neural net, construed as a single function, but the way in which it is constructed. Given the prominence of mathematical techniques in this construction, we should expect the mapping account to explain this in terms of further mappings between mathematical structures. For example, [Pincock \[2012\]](#) emphasizes the value of bringing in mathematics that is extrinsic to a given representation, largely by relating the intrinsic mathematics to the

extrinsic mathematics via morphisms. Likewise, on the inferential version of the mapping account proposed by [Bueno & Colyvan \[2011\]](#), the processes of immersion and interpretation, both underwritten by mappings, are allowed to be iterated, and the purpose of this is partially to help explain the role of mathematics that is not (initially at least) part of the representation itself. Certainly, we can establish some such mappings by, for example, embedding the impoverished structure described above within a structure containing all functions that could be generated by the relevant neural network, the relevant error function, and so on. But it is not clear how much that would actually get us. After all, part of the reason to use artificial neural networks is that it is difficult to determine what we really wish to know about such structures—particularly which function from that family minimizes the error function. This is why we need optimization algorithms like backpropagation to effectively use neural network techniques. So the picture we get from both the default and the inferential versions of the mapping account, according to which mappings between this larger structure and the narrower structure we arrive at after training the neural network underwrite inferences about properties of the former from properties of the latter seems to be at least incomplete.

I believe RIC allows for a more informative approach by allowing for a more fine-grained treatment of the setup of the relevant models in terms of reasoning about computational inference procedures. Training neural networks amounts to a reliable way to choose among a family of such procedures. Since the aim is only to produce a computationally tractable and reliable enough inference procedure, there is no need for the artificial neural network itself or the mathematics used to generate it to be appropriately morphic to the system the neural network represents—beyond the minimal sense described above, so that only the final function as a whole gets mapped to the structure of the target system. As a result, we can better understand a number of scientists more ad hoc seeming choices, which are difficult to understand if we construe them as reasoning about structures. For example, features like the dimensionality and specific functions for calculating hidden layers are largely chosen experimentally, with final values chosen to balance computational tractability with accurate enough results.

This situation is strikingly similar to the “experimental” mathematical techniques advocated for Heaviside (see §5.2.3), and I suspect a similar treatment would be useful here.

Examining neural network models in this way allows us to better see continuities between them and previous modeling practices. While they differ in that they are largely disconnected from theory as such, they are continuous with paradigmatic applications of mathematical models in that they are chosen and developed to balance reliable inference based on scientists’ prior understanding with practical features required by scientists’ cognitive limitations. Mathematical models often distort or abstract away from features of their target systems, and are sometimes shaped by ad hoc maneuvers in the service of tractability (as, for example, in the case of Heaviside). Neural network models differ from ordinary mathematical models in the degree to which this is the case, but they are arguably not different in kind.

Neural networks are just one case of the general-purpose computational inference techniques at work in recent, data-driven science. Others include simulated annealing algorithms [Bailer-Jones & Bailer-Jones, 2002], applied, for example, to vapor-liquid equilibrium models [Bonilla-Petriciolet *et al.*, 2007]; clustering algorithms [Ben-Dor *et al.*, 2000], applied, for example, in DNA microarray-based cancer research [Lu & Han, 2003]; and boosting techniques [Freund & Schapire, 1997; Hastie, Tibshirami & Friedman, 2008, ch. 10], used to combine multiple classifier functions that are each only slightly better than chance to produce a reliable classifier. In each case, I suspect that there is much to be gained by an analysis in terms of RIC, which permits a finer degree of grain than mapping accounts.

Moreover, this suggests the possibility of using RIC to unify a broader class of computational phenomena, including both explicit applications of mathematics and computational techniques that aren’t explicitly cashed out in mathematical language. In particular, I suspect that an analysis of computer modeling techniques in terms of RIC would allow the more ad hoc techniques used in computer simulations (as discussed, e.g., by Weisberg [2013]) to be treated similarly to techniques to secure computational tractability in applications of unrigorous mathematics and data-driven techniques.

8.2.2 RIC and mathematics as such

Finally, I suspect that RIC can be used to formulate an account of mathematics as such, particularly its semantics, that has a number of interesting properties. In particular, I think such an account has the potential to be a useful tool for thinking about mathematics prior to the move toward greater rigor in the nineteenth century and the twentieth-century project of set-theoretic foundations. In addition, I believe this account better accommodates the connection between the sort of mathematics physicists do and “pure” mathematics in a way that supports a better understanding not just of applications of the latter to the former, but also of the former to the latter. Finally, such an account opens up the intriguing possibility of providing a decidedly non-Fregean way to satisfy Frege’s application constraint, according to which the possibility for applying a mathematical concept must be built into its very meaning. While such a view must inevitably rely on some controversial assumptions, I believe it is well worth exploring for these reasons.

The central idea is to supplement RIC with a further inferentialist thesis in the tradition of [Brandom \[1994, 2000, 2008\]](#) and [Sellars \[1953\]](#): in explaining the meaning of a mathematical concept or expression, its inferential role is explanatorily prior to its reference. This doesn’t preclude a parallel treatment of all or part of mathematics in terms of mathematical objects or structures. As [Steinberger & Murzi \[2017\]](#) observe, inferentialism of this kind is a metase-mantic thesis, and is perfectly compatible with a referential semantics. But it does mean that the inferential behavior of mathematical concepts and expressions is explanatorily prior to a treatment in terms of objects or structures. In particular, it means that we should take the inference patterns in the RIC3 part of a mathematical scientific representation to exhaust the semantic meaning of the mathematics being applied. As a result, the high degree of flexibility afforded by RIC in reasoning about applications of mathematics can be extended to reasoning about mathematics independently of its applications.

While Brandomian inferentialism faces a number of well-known objections¹, the combin-

¹See [\[Steinberger & Murzi, 2017\]](#) for a useful summary.

ation of RIC with inferentialism about mathematical meaning has a number of benefits that make the view worth exploring.

The first of these advantages is its ability to account for features of historical mathematics. Since the project of set-theoretic foundations reached maturity in the twentieth century, standards of rigor in pure mathematics have dictated that proofs should be straightforward in principle to translate into proofs in axiomatic set theory, particularly ZFC.² In such cases, we have a straightforward connection to set-theoretic structure, and so it is natural to think of the relevant mathematics as in some sense the study of such structure. But it is less clear that it is useful to think of mathematicians prior to the introduction of these standards of rigor as reasoning about structures in the same way. In extreme cases, such as the early, inconsistent calculus (chapter 4), it is not clear that we can associate a particular structure with the relevant mathematics at all. More central to understanding the early calculus is understanding the inference strategies employed to safely calculate derivatives and integrals, and this is more usefully done directly in terms of inference rules and restrictions than indirectly in terms of structure. In essence, this involves understanding the calculus, even in contexts independent of physical applications, in terms of the inference patterns that would figure into RIC3 in the context of applications of the calculus.³

Second, for related reasons, such an approach would allow for a straightforward account of the applicability of the mathematical tools developed by physicists to “pure” mathematics. Often, the highly unrigorous tools developed by physicists have been useful both for solving problems in pure mathematics and as the inspiration for new, more rigorous mathematics [Urquhart, 2008a,b]. The account I have sketched here, unlike traditional accounts in the philosophy of mathematics, can treat the physicists’ home-brew mathematics as of fundamentally the same kind as the mathematics done by “pure” mathematicians, despite its failure

²But see [Tanswell, 2015] for reason to doubt that we should base an account of informal proof on formal proof, set-theoretic or otherwise.

³The same goes for a wide range of less well-appreciated cases from historical mathematics. For instance, Cauchy, well known for his foundational work in real analysis, made extensive use of the Dirac delta function *avant la lettre* [Laugwitz, 1987, 1992].

in many cases (e.g., path integrals, as discussed in §5.3) to pick out a well-defined structure. This allows for a much more natural account of the interface between physics (and other sciences) and mathematics.

Finally, the combination of RIC and inferentialism about mathematical meaning has the intriguing property of making applicability intrinsic to the semantic content of mathematical expressions, even in “pure mathematics.” If RIC3 is an exhaustive account of the semantic meaning of the mathematics being applied, then there is a sense in which the conditions for the applicability of that piece of mathematics is built into its semantic meaning. As a result, unlike traditional accounts in the philosophy of mathematics, this account meets Frege’s famous constraint [Wright, 2000]—used in his notoriously devastating critiques of mathematical formalism [Frege, 2013]—that the meanings of mathematical expressions must include the conditions for their application, albeit in a way of which Frege would have strongly disapproved. Notably, this semantic relationship goes the other way as well. Because mathematical representations directly build in the inferential behavior of the relevant mathematics via RIC3, a consequence of the view is that certain scientific concepts build mathematical concepts directly into their semantic meaning (rather than indirectly via a mapping).

Taken together, I think these benefits make a view combining RIC with inferentialism about the meaning of mathematical concepts and expressions well worth exploring.

REFERENCES

- ALBEVERIO, S.A., HØEGH-KROHN, R.J. & MAZZUCCHI, S. (2008). *Mathematical Theory of Feynman Path Integrals*. Springer, Berlin, 2nd edn.
- ARMOUR-GARB, B., STOLJAR, D. & WOODBRIDGE, J. (2022). Deflationism About Truth. In E.N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, Summer 2022 edn.
- ASENJO, F.G. (1966). A calculus of antinomies. *Notre Dame Journal of Formal Logic*, 7, 103–5.
- BAILER-JONES, C.A.L. (2000). Stellar parameters from very low resolution spectra and medium band filters: T_{eff} , $\log g$ and $[M/H]$ using neural networks. *Astronomy and Astrophysics*, 357, 197–205.
- BAILER-JONES, D.M. & BAILER-JONES, C.A.L. (2002). Modeling data: Analogies in neural networks, simulated annealing, and genetic algorithms. In L. Magnani, ed., *Model-Based Reasoning: Science, Technology, Values*, 147–65, Kluwer Academic Publishers, Dordrecht.
- BAIR, J., BŁASZCZYK, P., ELY, R., HENRY, V., KANOVEI, V., KATZ, K.U., KATZ, M.G., KUTATELADZE, S.S., MCGAFFEY, T., SCHAPS, D.M., SHERRY, D. & SHNIDER, S. (2013). Is mathematical history written by the victors? *Notices of the American Mathematical Society*, 60, 886–904.
- BAKER, A. (2003). The indispensability argument and multiple foundations for mathematics. *Philosophical Quarterly*, 53, 49–67.

- BAKER, A. (2005). Are there genuine mathematical explanations of physical phenomena? *Mind*, **114**, 223–38.
- BAKER, A. (2017). Mathematics and explanatory generality. *Philosophia Mathematica*, **25**, 194–209.
- BAKER, A. & COLYVAN, M. (2011). Indexing and mathematical explanation. *Philosophia Mathematica*, **19**, 323–334.
- BALAGUER, M. (1998). *Platonism and Anti-Platonism in Mathematics*. Oxford University Press, Oxford.
- BARON, S. (2020). Counterfactual scheming. *Mind*, **129**, 535–62.
- BARON, S. (forthcoming). Mathematical explanation: A pythagorean proposal. *British Journal for the Philosophy of Science*.
- BARTELS, A. (2006). Defending the structural concept of representation. *Theoria*, **55**, 7–19.
- BATTERMAN, R.W. (2010). On the explanatory role of mathematics in empirical science. *British Journal for the Philosophy of Science*, **61**, 1–25.
- BELL, J.L. (2008). *A Primer of Infinitesimal Analysis*. Cambridge University Press, Cambridge, 2nd edn.
- BEN-DOR, A., BRUHN, L., FRIEDMAN, N., NACHMAN, I., SCHUMMER, M. & YAKHINI, Z. (2000). Tissue classification with gene expression profiles. In *Proceedings of the Fourth Annual International Conference on Computational Molecular Biology*.
- BENHAM, R., MORTENSEN, C. & PRIEST, G. (2014). Chunk and permeate III: The Dirac delta function. *Synthese*, **191**, 3057–62.
- BOKULICH, A. (2008a). Can classical structures explain quantum phenomena? *British Journal for the Philosophy of Science*, **59**, 217–35.

-
- BOKULICH, A. (2008b). *Reexamining the Quantum-Classical Relation*. Cambridge University Press, Cambridge.
- BONILLA-PETRICIOLET, A., BRAVO-SÁNCHEZ, U.I., CASTILLO-BORJA, F., ZAPIAIN-SALINAS, J.G. & SOTO-BERNAL, J.J. (2007). The performance of simulated annealing in parameter estimation for vapor-liquid equilibrium modeling. *Brazilian Journal of Chemical Engineering*, 24, 151–62.
- BRANDOM, R. (1994). *Making It Explicit*. Harvard University Press, Cambridge, Mass.
- BRANDOM, R. (2000). *Articulating Reasons*. Harvard University Press, Cambridge, Mass.
- BRANDOM, R. (2008). *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford University Press, Oxford.
- BROMWICH, T.J. (1928). Some solutions of the electromagnetic equations, and of the elastic equations, with applications to the problem of secondary waves. *Proceedings of the London Mathematical Society*, 28, 438–75.
- BROMWICH, T.J.I. (1916). Normal coordinates in dynamical systems. *Proceedings of the London Mathematical Society*, 15, 401–48.
- BROWN, B. & PRIEST, G. (2004). Chunk and permeate, a paraconsistent inference strategy. Part I: The infinitesimal calculus. *Journal of Philosophical Logic*, 33, 379–88.
- BUENO, O. (1997). Empirical adequacy: A partial structures approach. *Studies in History and Philosophy of Science*, 28, 585–610.
- BUENO, O. & COLYVAN, M. (2011). An inferential conception of the application of mathematics. *Noûs*, 45, 345–374.
- BUENO, O. & FRENCH, S. (2011). How theories represent. *British Journal for the Philosophy of Science*, 62, 857–94.

-
- BUENO, O. & FRENCH, S. (2012). Can mathematics explain physical phenomena? *British Journal for the Philosophy of Science*, **63**, 85–113.
- BUENO, O. & FRENCH, S. (2018). *Applying Mathematics: Immersion, Inference, Interpretation*. Oxford University Press, Oxford.
- CAMERON, R.H. (1960). A family of integrals serving to connect the Wiener and Feynman integrals. *Journal of Mathematics and Physics*, **39**, 126–40.
- CAN, C. & AKTAS, M.E. (2019). “Derivative makes more sense with differentials”: How primary historical sources informed a university mathematics instructor’s teaching of derivative. In S. Brown, G. Karakok, K. Roh & M. Oehrtman, eds., *Proceedings of the 22nd Annual Conference for Research in Undergraduate Mathematics Education*, 866–71, SIGMAA-RUME, Oklahoma City.
- CARSON, J.R. (1926). *Electric Circuit Theory and the Operational Calculus*. McGraw-Hill, New York.
- CAUCHY, A.L. (1821). *Cours d’analyse de l’École Royale Polytechnique*. L’Imprimerie Royale, Paris.
- CHAKRAVARTTY, A. (2010). Informational versus functional theories of scientific representation. *Synthese*, **172**, 197–213.
- COLYVAN, M. (2002). Mathematics and aesthetic considerations in science. *Mind*, **111**, 69–74.
- COLYVAN, M. (2008a). The ontological commitments of inconsistent theories. *Philosophical Studies*, **141**, 115–23.
- COLYVAN, M. (2008b). Who’s afraid of inconsistent mathematics. *ProtoSociology*, **25**, 24–35.
- COLYVAN, M. (2009). Applying inconsistent mathematics. In O. Bueno & Ø. Linnebo, eds., *New Waves in Philosophy of Mathematics*, 160–72, Palgrave Macmillan, New York.

- COLYVAN, M. (2013). Road work ahead: Heavy machinery on the easy road. *Mind*, 121, 1031–46.
- CONTESSA, G. (2007). Scientific representation, interpretation, and surrogative reasoning. *Philosophy of Science*, 74, 48–68.
- CONTESSA, G. (2011). Scientific models and representation. In S. French & J. Saatsi, eds., *The Bloomsbury Companion to the Philosophy of Science*, 120–37, Bloomsbury, New York.
- COOPER, J.L.B. (1952). Heaviside and the operational calculus. *The Mathematical Gazette*, 36, 5–19.
- DA COSTA, N.C.A. & FRENCH, S. (1990). The model-theoretic approach to the philosophy of science. *Philosophy of Science*, 57, 248–65.
- DA COSTA, N.C.A. & FRENCH, S. (2003). *Science and Partial Truth: A Unitary Approach to Models and Scientific Reasoning*. Oxford University Press, Oxford.
- DALY, C. & LANGFORD, S. (2009). Mathematical explanation and indispensability arguments. *Philosophical Quarterly*, 59, 641–58.
- DAVEY, K. (2003). Is mathematical rigor necessary in physics? *British Journal for the Philosophy of Science*, 54, 439–63.
- DE REGT, H.W. (2017). *Understanding Scientific Understanding*. Oxford University Press, Oxford.
- DE REGT, H.W., LEONELLI, S. & EIGNER, K., eds. (2009). *Scientific Understanding: Philosophical Perspectives*. University of Pittsburgh Press, Pittsburgh, PA.
- DECARLO, R.A. & LIN, P.M. (2009). *Linear Circuits: Time Domain, Phasor, and Laplace Transform Approaches*. Kendall Hunt, Dubuque, Iowa, 3rd edn.

- DEWITT, B.S. (1964). Theory of radiative corrections for non-abelian gauge fields. *Physical Review Letters*, **12**, 742.
- DIRAC, P. (1933). The Lagrangian in quantum mechanics. *Physikalische Zeitschrift der Sowjetunion*, **3**, 64–72.
- DIRAC, P. (1967). *The Principles of Quantum Mechanics*. Clarendon Press, Oxford, revised 4th edn., first edition published in 1930, 4th edition first published in 1958.
- DOWKER, F. (2012). The path integral interpretation of quantum mechanics. Public lecture at the Perimeter Institute, University of Western Ontario. Video available at <https://pirsa.org/12070001>.
- DOWKER, F., JOHNSTON, S. & SORKIN, R.D. (2010). Hilbert spaces from path integrals. *Journal of Physics A: Mathematical and Theoretical*, **43**, 275–302.
- ELGIN, C. (2017). *True Enough*. MIT Press, Cambridge, Mass.
- FADDEEV, L.D. & POPOV, V.N. (1967). Feynman diagrams for the Yang-Mills field. *Physics Letters B*, **25**, 29–30.
- FEYNMAN, R.P. (1948). Space-time approach to non-relativistic quantum mechanics. *Reviews of Modern Physics*, **20**, 367–387.
- FEYNMAN, R.P. (1949). Space-time approach to quantum electrodynamics. *Physical Review*, **76**, 769–89.
- FEYNMAN, R.P. (2005). The principle of least action in quantum mechanics. In L.M. Brown, ed., *Feynman's Thesis: A New Approach to Quantum Theory*, 1–70, World Scientific, Singapore.
- FEYNMAN, R.P. & HIBBS, A.R. (1965). *Quantum Mechanics and Path Integrals*. Dover, Mineola, N.Y., emended edition by Daniel F. Styer published in 2010.
- FOURIER, M. (1822). *Théorie analytique de la chaleur*. Firmin Didot, Paris.

-
- FRANKLIN, J. (2014). *An Aristotelian Realist Philosophy of Mathematics: Mathematics as the Science of Quantity and Structure*. Palgrave Macmillan, London.
- FREGE, G. (2013). *Basic Laws of Arithmetic: Derived using concept script*. Oxford University Press, Oxford, ed. and trans. P. A. Ebert and M. Rossberg.
- FRENCH, S. (2003). A model-theoretic account of representation (or, I don't know much about art... but I know it involves isomorphism). *Philosophy of Science*, **70**, 1472–83.
- FRENCH, S. (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press, Oxford.
- FRENCH, S. & LADYMAN, J. (1999). Reinflating the semantic approach. *International Studies in the Philosophy of Science*, **13**, 103–19.
- FREUND, Y. & SCHAPIRE, R. (1997). A decision-theoretic generalization of online learning and an application to boosting. *Journal of Computer and System Sciences*, **55**, 119–139.
- FRIEDMAN, M. (1974). Explanation and scientific understanding. *Journal of Philosophy*, **71**, 5–19.
- FRIGG, R. (2006). Scientific representation and the semantic view of theories. *Theoria*, **21**, 49–65.
- FRIGG, R. & NGUYEN, J. (2016). The fiction view of models reloaded. *The Monist*, **99**, 225–42.
- FRIGG, R. & NGUYEN, J. (2017). Scientific representation is representation-as. In H.K. Chao & J. Reiss, eds., *Philosophy of Science in Practice: Nancy Cartwright and the Nature of Scientific Reasoning*, 149–79, Springer, Berlin.
- GATTRINGER, C. & LANG, C.B. (2010). *Quantum Chromodynamics on the Lattice: An Introductory Presentation*. Springer, Berlin.
- GIERE, R. (2004). How models are used to represent reality. *Philosophy of Science*, **71**, 742–52.

-
- GIERE, R. (2010). An agent-based conception of models and scientific representation. *Synthese*, 172, 269–81.
- GLIMM, J. & JAFFE, A. (1981). *Quantum Physics: A Functional Integral Point of View*. Springer-Verlag, New York.
- GOLES, E., SCHULZ, O. & MARKUS, M. (2001). Prime number selection of cycles in a predator-prey model. *Complexity*, 6, 33–8.
- GRIMM, S. (2006). Is understanding a species of knowledge? *British Journal for the Philosophy of Science*, 57, 515–35.
- HALE, B. & WRIGHT, C. (2001). *The Reason's Proper Study: Essays towards a Neo-Fregean Philosophy of Mathematics*. Oxford University Press, Oxford.
- HALL, B.C. (2013). *Quantum Theory for Mathematicians*. Springer, New York.
- HASTIE, T., TIBSHIRAMI, R. & FRIEDMAN, J. (2008). *The Elements of Statistical Learning*. Springer, New York, 2nd edn., corrected 12th printing, 13 January 2017.
- HEAVISIDE, O. (1892). *Electrical Papers*, vol. 1. Macmillan, London.
- HEAVISIDE, O. (1894). *Electrical Papers*, vol. 2. Macmillan, London.
- HEAVISIDE, O. (1899). *Electromagnetic Theory*, vol. 2. Chelsea Publishing Company, New York, 3rd edn., third edition published in 1971.
- HERON, J. (2020). Representational indispensability and ontological commitment. *Thought: A Journal of Philosophy*, 9.
- HIGGINS, T.J. (1949). History of the operational calculus as used in electric circuit analysis. *Electrical Engineering*, 68, 42–5.
- HILLS, A. (2016). Understanding why. *Noûs*, 49, 661–88.

-
- HUGHES, R.I.G. (1997). Models and representation. *Philosophy of Science*, **64**, S325–36.
- HUMPHREYS, P. (2000). Analytic versus synthetic understanding. In *Science, Explanation, and Rationality: The Philosophy of Carl G. Hempel*, 267–86, Oxford University Press, Oxford.
- JANSSON, L. & SAATSI, J. (2019). Explanatory abstractions. *British Journal for the Philosophy of Science*, **70**, 817–44.
- JEFFREYS, H. (1927). *Operational Methods in Mathematical Physics*. Cambridge University Press, Cambridge.
- JOHNSON, G.W. & LAPIDUS, M.L. (2000). *The Feynman Integral and Feynman's Operational Calculus*. Oxford University Press, Oxford.
- KHALIFA, K. (2012). Inaugurating understanding or repackaging explanation. *Philosophy of Science*, **79**, 15–37.
- KHALIFA, K. (2013). Understanding, grasping, and luck. *Episteme*, **10**, 1–17.
- KHALIFA, K. (2017). *Understanding, Explanation, and Scientific Knowledge*. Cambridge University Press, Cambridge.
- KITCHER, P. (1981). Explanatory unification. *Philosophy of Science*, **48**, 507–31.
- KITCHER, P. (1989). Explanatory unification and the causal structure of the world. In K. Philip & W.C. Salmon, eds., *Scientific Explanation*, 410–505, University of Minnesota Press, Minneapolis.
- KLEINER, I. (2002). History of the infinitely small and the infinitely large in calculus. *Educational Studies in Mathematics*, **48**, 137–74.
- KLINE, M. (1972). *Mathematical Thought from Ancient to Modern Times*. Oxford University Press, Oxford.

-
- KLINE, R.R. (1992). *Steinmetz: Engineer and Socialist*. Johns Hopkins University Press, Baltimore.
- KNOWLES, R. & SAATSI, J. (2021). Mathematics and explanatory generality: Nothing but cognitive salience. *Erkenntnis*, **86**, 1119–37.
- KOPPELMAN, E. (1971). The calculus of operations and the rise of abstract algebra. *Archive for History of Exact Sciences*, **8**, 155–242.
- KUORIKOSKI, J. (2021). There are no mathematical explanations. *Philosophy of Science*, **88**, 189–212.
- KUORIKOSKI, J. & YLIKOSKI, P. (2015). External representation and scientific understanding. *Synthese*, **192**, 3817–37.
- KVANVIG, J. (2003). *The Nature and Value of Knowledge*. Cambridge University Press, Cambridge.
- LACROIX, S.F. (1819). *Traité du calcul différentiel et du calcul intégral*. Courcier, Paris, 2nd edn.
- LANGE, M. (2013). What makes a scientific explanation distinctively mathematical? *British Journal for the Philosophy of Science*, **64**, 485–511.
- LAUGWITZ, D. (1987). Hidden lemmas in the early history of infinite series. *Aequationes Mathematicae*, **34**, 264–76.
- LAUGWITZ, D. (1992). Early delta functions and the use of infinitesimals in research. *Revue d'histoire des sciences*, **45**, 115–28.
- LE BIHAN, S. (2021). Partial truth versus felicitous falsehoods. *Synthese*, **198**, 5415–36.
- LENG, M. (2002). What's wrong with indispensability? (or the case for recreational mathematics). *Synthese*, **131**, 395–417.
- LIGGINS, D. (2016). Grounding and the indispensability argument. *Synthese*, **193**, 531–48.

-
- LIPTON, P. (2009). Understanding without explanation. In H.W. de Regt, S. Leonelli & K. Eigner, eds., *Scientific Understanding: Philosophical Perspectives*, 43–63, University of Pittsburgh Press, Pittsburgh.
- LLOYD, E. (1984). A semantic approach to the structure of population genetics. *Philosophy of Science*, 51, 242–64.
- LLOYD, E.A. (1988). *The Structure and Confirmation of Evolutionary Theory*. Princeton University Press, Princeton, N.J.
- LU, Y. & HAN, J. (2003). Cancer classification using gene expression data. *Information Systems*, 28, 243–68.
- LÜTZEN, J. (1979). Heaviside’s operational calculus and the attempts to rigorise it. *Archive for History of Exact Sciences*, 21, 161–200.
- LYNCH, M.P. (2009). *Truth as One and Many*. Oxford University Press, Oxford.
- LYON, A. (2012). Mathematical explanations of empirical facts, and mathematical realism. *Australasian Journal of Philosophy*, 90, 559–78.
- MANCOSU, P. (2018). Explanation in Mathematics. In E.N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, Summer 2018 edn.
- MANIN, Y.I. (1981). *Mathematics and Physics*. Birkhäuser, Boston, translated by A. Koblitz and N. Koblitz.
- MANIN, Y.I. (1989). Strings. *The Mathematical Intelligencer*, 11, 59–65, reprinted in [Wilson & Gray, 2001, pp. 229–37].
- MAZZUCCHI, S. (2009). *Mathematical Feynman Path Integrals and Their Applications*. World Scientific, Singapore.

- McCULLOUGH-BENNER, C. (2019). Representing the world with inconsistent mathematics. *British Journal for the Philosophy of Science*, **71**, 1331–58.
- McCULLOUGH-BENNER, C. (2022a). Applying unrigorous mathematics: Heaviside’s operational calculus. *Studies in History and Philosophy of Science Part A*, **91**, 113–24.
- McCULLOUGH-BENNER, C. (2022b). The metarepresentational role of mathematics in scientific explanations. *Philosophy of Science*, **89**, 742–60.
- MELIA, J. (2000). Weaseling away the indispensability argument. *Mind*, **109**, 458–79.
- MELIA, J. (2002). Response to Colyvan. *Mind*, **111**, 75–9.
- MESSIAH, A. (1961). *Quantum Mechanics*, vol. 1. North-Holland Publishing Company, Amsterdam.
- MOERDIJK, I. & REYES, G.E. (1991). *Models for Smooth Infinitesimal Analysis*. Springer, New York.
- MONTALDI, J. & SMOLYANOV, O.G. (2017). Feynman path integrals and Lebesgue-Feynman measures. *Doklady Mathematics*, **96**, 368–72.
- MORRISON, M. (2015). *Reconstructing Reality: Models, Mathematics and Simulations*. Oxford University Press, Oxford.
- MORTENSEN, C. (1995). *Inconsistent Mathematics*. Kluwer Academic Publishers, Dordrecht.
- MUNDY, B. (1986). On the general theory of meaningful representation. *Synthese*, **67**, 391–437.
- NAHIN, P.J. (2002). *Oliver Heaviside: The Life, Work, and Times of an Electrical Genius of the Victorian Age*. Johns Hopkins University Press, Baltimore, Maryland, 2nd edn., originally published in 1988 by the Institute of Electrical and Electronics Engineers.
- NAPOLETANI, D., PANZA, M. & STRUPPA, D.C. (2011). Agnostic science. toward a philosophy of data analysis. *Foundations of Science*, **16**, 1–20.

- NEWTON, I. (1999). *The Principia: Mathematical Principles of Natural Philosophy*. University of California Press, Oakland, Cal., trans. I. Bernard Cohen and Anne Whitman, assisted by Julia Budenz.
- NGUYEN, J. & FRIGG, R. (2021). Mathematics is not the only language in the book of nature. *Synthese*, **198 (Suppl 24)**, 5941–62.
- NGUYEN, T. (2016). The perturbative approach to path integrals: A succinct mathematical treatment. *Journal of Mathematical Physics*, **57**.
- OBERBECK, A. (1882). Ueber elektrische Schwingungen mit besonderer Berücksichtigung ihrer Phasen. *Annalen der Physik und Chemie*, **17**, 816–41.
- PEDERSEN, N.J.L.L. & WRIGHT, C. (2018). Pluralist Theories of Truth. In E.N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, Winter 2018 edn.
- PETROVA, S.S. (1987). Heaviside and the development of the symbolic calculus. *Archive for History of Exact Sciences*, **37**, 1–23.
- PINCOCK, C. (2004). A new perspective on the problem of applying mathematics. *Philosophia Mathematica*, **3**, 135–61.
- PINCOCK, C. (2012). *Mathematics and Scientific Representation*. Oxford University Press, Oxford.
- PLEBANI, M. (2016). Nominalistic content, grounding, and covering generalizations: Reply to ‘Grounding and the indispensability argument’. *Synthese*, **193**, 549–558.
- POTOCHNIK, A. (2017). *Idealization and the Aims of Science*. University of Chicago Press, Chicago.
- PRIEST, G. (1979). The logic of paradox. *Journal of Philosophical Logic*, **8**, 219–41.

-
- LORD RAYLEIGH (1886a). On the self-induction and resistance of straight conductors. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, Fifth Series*, 21, 381–94.
- LORD RAYLEIGH (1886b). The reaction upon the driving-point of a system executing forced harmonic oscillations of various periods, with applications to electricity. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, Fifth Series*, 21, 369–81.
- LORD RAYLEIGH (1891). On the sensitiveness of the bridge method in its application to periodic electric currents. *Proceedings of the Royal Society*, 49, 203–17.
- RAYO, A. (2009). Toward a trivialist account of mathematics. In O. Bueno & Ø. Linnebo, eds., *New Waves in the Philosophy of Mathematics*, 239–260, Palgrave Macmillan, New York.
- RÄZ, T. & SAUER, T. (2015). Outline of a dynamical inferential conception of the application of mathematics. *Studies in History and Philosophy of Modern Physics*, 49, 57–72.
- REDHEAD, M.L.G. (1975). Symmetry in intertheory relations. *Synthese*, 32, 77–112.
- RESNIK, M. (1997). *Mathematics as a Science of Patterns*. Clarendon Press, Oxford.
- REUTLINGER, A. (2018). Extending the counterfactual theory of explanation. In A. Reutlinger & J. Saatsi, eds., *Explanation Beyond Causation: Philosophical Perspectives on Non-Causal Explanations*, 74–95, Oxford University Press, Oxford.
- REUTLINGER, A. & SAATSI, J., eds. (2018). *Explanation Beyond Causation: Philosophical Perspectives on Non-Causal Explanations*. Oxford University Press, Oxford.
- RIGGS, W. (2009). Understanding, knowledge, and the Meno requirement. In A. Haddock, A. Millar & D. Pritchard, eds., *Epistemic Value*, 331–38, Oxford University Press, Oxford.
- RIVERS, R.J. (1987). *Path Integral Methods in Quantum Field Theory*. Cambridge University Press, Cambridge.

-
- RIZZA, D. (2013). The applicability of mathematics: Beyond mapping-based accounts. *Philosophy of Science*, **80**, 398–412.
- ROBINSON, A. (1966). *Non-Standard Analysis*. North-Holland Publishing Co., Amsterdam.
- ROQUETTE, P. (2010). Numbers and models, standard and nonstandard. *Mathematische Semesterberichte*, **57**, 185–99.
- ROSS, B. (1977). The development of fractional calculus 1695–1900. *Historia Mathematica*, **4**, 75–89.
- SAATSI, J. (2011). The enhanced indispensability argument: Representational versus explanatory role of mathematics in science. *British Journal for the Philosophy of Science*, **62**, 143–154.
- SAATSI, J. (2016). On the ‘indispensable explanatory role’ of mathematics. *Mind*, **125**.
- SAATSI, J. & PEXTON, M. (2013). Reassessing Woodward’s account of explanation: Regularities, counterfactuals, and noncausal explanations. *Philosophy of Science*, **80**, 613–24.
- SALIVAHANAN, S., VALLAVARAJ, A. & GNANAPRIYA, C. (2000). *Digital Signal Processing*. McGraw-Hill, New Delhi.
- SALMON, W.C. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton.
- SCHULMAN, L.S. (1988). Introduction to the path integral. In S. Lundqvist, A. Ragnagni, V. Sa-yakanit & L.S. Schulman, eds., *Path Summation: Achievements and Goals*, 3–46, World Scientific, Singapore, page numbers from preprint available at <https://people.clarkson.edu/~lschulma/1988TriestePathIntegralLectures.pdf>.
- SCHULMAN, L.S. (2005). *Techniques and Applications of Path Integration*. Dover, Mineola, N.Y., 2nd edn.

- SCHWARTZ, L. (1945). Généralisation de la notion de fonction, de dérivation, de transformation de Fourier et applications mathématiques et physiques. *Annales de l'Université de Grenoble*, 21, 57–74.
- SELLARS, W. (1953). Inference and meaning. *Mind*, 62, 313–38.
- SHAPIRO, S. (1997). *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press, Oxford.
- STALNAKER, R.C. (1987). *Inquiry*. MIT Press, Cambridge, Mass.
- STAMEY, T.A., KABALIN, J.N., MCNEAL, J.E., JOHNSTONE, I.M., FREIHA, F., REDWINE, E.A. & N, Y. (1989). Prostate specific antigen in the diagnosis and treatment of adenocarcinoma of the prostate. II. Radical prostatectomy treated patients. *Journal of Urology*, 141, 1076–83.
- STEINBERGER, F. & MURZI, J. (2017). Inferentialism. In *Blackwell Companion to Philosophy of Language*, 197–224, Wiley Blackwell.
- STEINER, M. (1992). Mathematical rigor in physics. In M. Detlefsen, ed., *Proof and Knowledge in Mathematics*, 158–70, Routledge, London.
- STEINER, M. (1998). *The Applicability of Mathematics as a Philosophical Problem*. Harvard University Press, Cambridge, Mass.
- STREVENS, M. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, Mass.
- STREVENS, M. (2017). How idealizations provide understanding. In S. Grimm, C. Baumberger & S. Ammon, eds., *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*, 37–49, Routledge, New York.
- SUÁREZ, M. (2003). Scientific representation: Against similarity and isomorphism. *International Studies in the Philosophy of Science*, 17, 225–44.

- SUÁREZ, M. (2004). An inferential conception of scientific representation. *Philosophy of Science*, 71, 767–79.
- SUÁREZ, M. (2015). Deflationary representation, inference, and practice. *Studies in History and Philosophy of Science*, 49, 36–47.
- SUÁREZ, M. & CARTWRIGHT, N. (2008). Theories: Tools versus models. *Studies in History and Philosophy of Modern Physics*, 39, 62–81.
- SULLIVAN, E. (2018). Understanding: Not know-how. *Philosophical Studies*, 175, 221–40.
- SULLIVAN, E. & KHALIFA, K. (2019). Idealizations and understanding: Much ado about nothing? *Australasian Journal of Philosophy*, 97, 673–689.
- SWANSON, M.S. (1992). *Path Integrals and Quantum Processes*. Academic Press, San Diego, Calif.
- SWOYER, C. (1991). Structural representation and surrogative reasoning. *Synthese*, 87, 449–508.
- 'T HOOFT, G. (1971). Renormalizable Lagrangians for massive Yang-Mills fields. *Nuclear Physics B*, 35, 167–88.
- TALL, D. (1981). Comments on the difficulty and validity of various approaches to the calculus. *For the Learning of Mathematics*, 2, 16–21.
- TANSWELL, F. (2015). A problem with the dependence of informal proofs on formal proofs. *Philosophia Mathematica*, 23, 295–310.
- TOEPLITZ, O. (2015). The problem of university courses on infinitesimal calculus and their demarcation from infinitesimal calculus in high schools. *Science in Context*, 28, lecture originally presented in 1926. Translated by M.N. Fried and H.N. Jahnke.
- URQUHART, A. (2008a). The boundary between mathematics and physics. In P. Mancosu, ed., *The Philosophy of Mathematical Practice*, 407–416, Oxford University Press, Oxford.

- URQUHART, A. (2008b). Mathematics and physics: Strategies of assimilation. In P. Mancosu, ed., *The Philosophy of Mathematical Practice*, 417–40, Oxford University Press, Oxford.
- VAN FRAASSEN, B.C. (1980). *The Scientific Image*. Oxford University Press, Oxford.
- VAN FRAASSEN, B.C. (2008). *Scientific Representation: Paradoxes of Perspective*. Oxford University Press, Oxford.
- VICKERS, P. (2009). Can partial structures accommodate inconsistent science? *Principia*, 13, 233–50.
- VICKERS, P. (2013). *Understanding Inconsistent Science*. Oxford University Press, Oxford.
- VINCENT, A., BARBEROUSSE, A. & IMBERT, C. (2018). Inferential power, formalisms, and scientific models, preprint of a paper given at the 2018 Philosophy of Science Association meeting. Archived at <http://philsci-archive.pitt.edu/15165/>.
- WAKIL, S. & JUSTUS, J. (2017). Mathematical explanation and the biological optimization fallacy. *Philosophy of Science*, 84, 916–30.
- WEINBERG, S. (1995). *The Quantum Theory of Fields: Volume I, Foundations*. Cambridge University Press, Cambridge.
- WEISBERG, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford University Press, Oxford.
- WIETLISBACH, V. (1879). Ueber die Anwendung die Telephons zu elektrischen und galvanischen Messungen. *Monatsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*, 278–83.
- WIGNER, E.P. (1960). The unreasonable effectiveness of mathematics in the natural sciences. *Communications on Pure and Applied Mathematics*, 13, 1–14.

- WILKENFELD, D.A. (2013). Understanding as representation manipulability. *Synthese*, **190**, 997–1016.
- WILKENFELD, D.A. & LOMBROZO, T. (2020). Explanation classification depends on understanding: Extending the epistemic side-effect effect. *Synthese*, **197**, 2565–92.
- WILSON, K.G. (1974). Confinement of quarks. *Physical Review D*, **10**, 2445–59.
- WILSON, M. (2006). *Wandering Significance: An Essay on Conceptual Behavior*. Oxford University Press, Oxford.
- WILSON, R. & GRAY, J., eds. (2001). *Mathematical Conversations: Selections from The Mathematical Intelligencer*. Springer-Verlag, New York.
- WINSBERG, E.B. (2010). *Science in the Age of Computer Simulation*. University of Chicago Press, Chicago.
- WOODWARD, J. (2003). *Making Things Happen: A Causal Theory of Explanations*. Oxford University Press, Oxford.
- WRIGHT, C. (1992). *Truth and Objectivity*. Harvard University Press, Cambridge, Mass.
- WRIGHT, C. (2000). Neo-Fregean foundations for real analysis: Some reflections on Frege's constraint. *Notre Dame Journal of Formal Logic*, **41**, 317–334.
- YAVETZ, I. (1995). *From Obscurity to Enigma: The Work of Oliver Heaviside, 1872–1889*. Birkhäuser, Basel.
- ZAGZEBSKI, L. (2001). Recovering understanding. In M. Steup, ed., *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, 235–52, Oxford University Press, Oxford.
- ZUCCHERI, L. & ZUDINI, V. (2014). History of teaching calculus. In A. Karp & G. Schubring, eds., *Handbook on the History of Mathematics Education*, 493–514, Springer, New York.