

Mathematics and Scientific Representation

Christopher Pincock

December 31, 2010

Contents

1	Introduction	1
1.1	A Problem	1
1.2	Classifying Contributions	4
1.3	An Epistemic Solution	11
1.4	Explanatory Contributions	19
1.5	Other Approaches	23
1.6	Interpretative Flexibility	29
1.7	Key Claims	33
I	Epistemic Contributions	38
2	Content and Confirmation	39
2.1	Concepts	39
2.2	Basic Contents	42
2.3	Enriched Contents	47
2.4	Schematic and Genuine Contents	51
2.5	Inference	53
2.6	Core Conceptions	57
2.7	Intrinsic and Extrinsic	60

2.8	Confirmation Theory	63
2.9	Prior Probabilities	69
3	Causes	75
3.1	Accounts of Causation	75
3.2	A Causal Representation	81
3.3	Some Acausal Representations	86
3.4	The Value of Acausal Representations	101
3.5	Batterman and Wilson	107
4	Varying Interpretations	112
4.1	Abstraction as Variation	112
4.2	Irrotational Fluids and Electrostatics	116
4.3	Shock Waves	126
4.4	The Value of Varying Interpretations	131
4.5	Varying Interpretations and Discovery	136
4.6	The Toolkit of Applied Mathematics	139
5	Scale Matters	146
5.1	Scale and Scientific Representation	146
5.2	Scale Separation	148
5.3	Scale Similarity	157
5.4	Scale and Idealization	162
5.5	Perturbation Theory	175
5.6	Multiple Scales	178
5.7	Interpreting Multiscale Representations	190
5.8	Summary	200

6	Constitutive Frameworks	203
6.1	A Different Kind of Contribution	203
6.2	Carnap's Linguistic Frameworks	206
6.3	Kuhn's Paradigms	213
6.4	Friedman on the Relative A Priori	222
6.5	The Need for Constitutive Representations	232
6.6	The Need for the Absolute A Priori	235
7	Failures	239
7.1	Mathematics and Scientific Failure	239
7.2	Completeness and Segmentation Illusions	241
7.3	The Parameter Illusion	247
7.4	Illusions of Scale	260
7.5	Illusions of Traction	264
7.6	Causal Illusions	273
7.7	Finding the Scope of a Representation	277
II	Other Contributions	283
8	Discovery	284
8.1	Semantic and Metaphysical Problems	284
8.2	A Descriptive Problem	295
8.3	Description and Discovery	303
8.4	Defending Naturalism	310
8.5	Natural Kinds	317
9	Indispensability	321
9.1	Descriptive Contributions and Pure Mathematics	321

9.2	Quine and Putnam	322
9.3	Against the Platonist Conclusion	332
9.4	Colyvan	339
10	Explanation	345
10.1	Explanatory Contributions	345
10.2	Inference to the Best Mathematical Explanation	359
10.3	Belief and Understanding	370
11	The Rainbow	376
11.1	Asymptotic Explanation	376
11.2	Angle and Color	381
11.3	Explanatory Power	388
11.4	Supernumerary Bows	391
11.5	Interpretation and Scope	401
11.6	Batterman and Belot	406
11.7	Looking Ahead	411
12	Fictionalism	413
12.1	Motivations	413
12.2	Literary Fiction	416
12.3	Mathematics	427
12.4	Models	437
12.5	Understanding and Truth	446
13	Facades	450
13.1	Physical and Mathematical Concepts	450
13.2	Against Semantic Finality	452
13.3	Developing and Connecting Patches	457

13.4 A New Approach to Content	470
13.5 Azzouni and Rayo	474
14 Conclusion: Pure Mathematics	477
14.1 Taking Stock	477
14.2 Metaphysics	480
14.3 Structuralism	487
14.4 Epistemology	489
14.5 Peacocke and Jenkins	497
14.6 Historical Extensions	503
14.7 Non-conceptual Justification	506
14.8 Past and Future	511
Appendices	515
A Method of Characteristics	515
B Black-Scholes Model	517
C Speed of Sound	519
D Two Proofs of Euler's Formula	521

List of Figures

3.1	The Bridges of Königsberg	87
3.2	A Non-Eulerian Graph	88
3.3	An Eulerian Graph	88
3.4	Velocity and Flux (Illner et. al. 2005, p. 155)	93
4.1	Source (Kundu & Cohen 2008, p. 173)	122
4.2	Flow Around an Object (Kundu & Cohen 2008, p. 177)	123
4.3	Characteristic Base Curves (Illner et. al. 2005, p. 187)	128
5.1	Normal vs. ENSO Conditions (courtesy of NOAA/PMEL/TAO)	152
5.2	Temperature vs. Distance (Batterman 2002 <i>b</i> , p. 48)	158
5.3	Self-Similar Coordinates (Batterman 2002 <i>b</i> , p. 49)	158
5.4	Surface Waves (Kundu & Cohen 2008, p. 219)	165
5.5	Wave Dispersion (Kundu & Cohen 2008, p. 243)	168
5.6	Boundary Layer Flow (Kundu & Cohen 2008, p. 48)	180
5.7	Ideal Fluid Flow (Kundu & Cohen 2008, p. 166)	184
5.8	The Failure of a Single Scale (McGivern 2008, p. 62)	193
5.9	The Success of Two Scales (McGivern 2008, p. 63)	193
5.10	Bénard Cells	195
7.1	Column Storage	242

7.2	Two Support System	242
7.3	Three Support System	244
7.4	The Failure of the Three Support System	244
7.5	LTCM Return on Investment	252
7.6	A Fold Catastrophe (Sussman & Zahler 1978, p. 146)	267
7.7	A Cusp Catastrophe (Sussman & Zahler 1978, p. 149)	269
7.8	Divergent Trajectories (Sussman & Zahler 1978, p. 125)	271
11.1	A Rainbow	381
11.2	Ray Path (Nahin 2004, p. 181)	383
11.3	Rainbow with Supernumerary Bows	391
11.4	Impact Parameter (Nussenzweig 1992, p. 16)	396
11.5	Critical Points (Nussenzweig 1992, p. 104)	397
11.6	Comparing Predictions (Nussenzweig 1992, p. 113)	400
13.1	Polar Coordinates	465
13.2	A Riemann Surface	467

Chapter 1

Introduction

1.1 A Problem

The success of science is undeniable, but the nature of that success remains opaque. To make a start on characterizing the success of science, I will stipulate that science is in the business of producing representations of the physical world. These representations will be accurate just in case some region of the world is a certain way. When approached from this angle, it seems obvious that the success of science should be characterized by saying that our scientific representations are, for the most part, accurate. That is, they present this or that region to be a certain way, and in fact the world is that way.

This natural form of scientific realism is opposed on many fronts. In the history and sociology of science many would argue that the success of science is not best characterized by the accuracy of its representations. Instead, the success is defined in terms of the agreement of the community of scientists. The main question, for a given historical or sociological episode, is how a

representation came to be widely accepted by a given community, and there is no easy inference from widespread agreement to realistic correctness. Call this the *social conception* of the success of science. Within the philosophy of science, the social conception is largely ignored, but there are still several apparently viable alternatives to realism. Each agrees that the success of science has something to do with the accuracy of scientific representations, but qualifies the respects in which a given representation has to be accurate in order for it to be successful. Instrumentalism and constructive empiricism, in very different ways, restrict the success of science to what is observable. Others focus on the existence of entities or the structural relations between entities. Cartwright offers an even more qualified picture of the success of science by restricting the scope of the regularities offered by our scientific representations (Cartwright 1999).

All of these differing characterizations of the success of science need to face the problem that I will articulate here. Put most generally, the problem is to account for the presence of mathematics in our scientific representations. The central issue is to see to what extent the central place of mathematics in science contributes to the success of science. More carefully, for a given physical situation, context, and mathematical representation of the situation, (i) what does the mathematics contribute to the representation, (ii) how does it make this contribution and (iii) what must be in place for this contribution to occur? The parameters of this more careful formulation of the problem make clear that the way a given piece of mathematics contributes to the success of science in one representation can differ from the contribution from some other piece in some other context.

Advocates of the social conception of the success of science are likely to emphasize the respects in which mathematics contributes to the convergence

of scientists on a given representation. Porter's *Trust in Numbers* is an especially thorough exploration of this approach (Porter 1995). While this strategy for tackling our problem cannot be dismissed out of hand, it surely would be a disappointing result for the realist, or anybody else who thinks scientific representations aim to capture some aspect of the real world, if this social conception of the contribution of mathematics was the full story. For if this was the whole account of the contributions of mathematics, then it would be hard to resist the social conception of science more generally.

Still, it is my experience that philosophers otherwise resistant to the social conception of science are tempted to take this way out when it comes to the role of mathematics in science. One reason for this seems to be that several exaggerated claims have been made on behalf of the role of mathematics in the success of science, and the only alternative to these claims seems to be a social conception. The position I have in mind here might be termed the *metaphysical conception* of the way in which mathematics contributes to the success of scientific representations. According to the metaphysical conception, mathematics contributes to accurate scientific representations because the physical world is itself essentially mathematical. An inventory of the genuine constituents of the physical world reveals that these constituents include mathematical entities. So, it is no surprise that our best scientific representations include a lot of mathematics. For this mathematics is needed to characterize this or that part of reality accurately.¹ A similar attitude has been held for scientific representations that include causation. A metaphysical account of the causal character of these representations is simply that the world itself contains causal relationships and these should be reflected in

¹This appears to be the conclusion of Franklin 1989 and Franklin 2008, but it is not clear how widely he intends to extend this metaphysical approach.

our most accurate representations.

As with the social conception, it is hard to deny that there are certain cases that fit the metaphysical conception of the contribution of mathematics to the success of science. At the same time, there is good reason to think that this cannot be whole story. In the vast majority of cases in which we find mathematics being deployed, there is little inclination for the scientist to take the mathematics to reflect the underlying reality of the phenomenon in question. We see this, for example, whenever a representation employs a mathematical idealization. An idealized representation, for my purposes, is a representation that involves assumptions that the agent deploying the representation takes to be false. When a fluid is represented as continuous, even though the agent working with the representation believes that the system is discrete, then we have some useful mathematics that resists any simple metaphysical interpretation. And it is precisely these sorts of cases that tend to provoke a slide into a social conception of their contribution. For if the mathematics does not reflect any underlying physical constituents of the system, then what other contribution could it be making besides helping to generate agreement among scientists?

1.2 Classifying Contributions

To move beyond this unsatisfying dilemma, we need a more careful inventory of the contributions that mathematics seems to be making in our successful scientific representations. With this in hand, we can see if there are any other general strategies for answering our question. I will start by presenting four different dimensions that we can consider for a given representation. For each dimension, mathematics seems to be doing something different as

we move from one extreme to the other. While somewhat elaborate, this classification scheme will allow us to isolate more or less pure cases of how mathematics contributes in a given way. Other, more difficult, cases can then be understood as impure mixtures of these different sorts of contributions.

Before we explore these four dimensions we should first distinguish between the mathematical content of the representation itself and the mathematical techniques which were used to arrive at that representation. We will dub the mathematics of a representation *intrinsic* when it figures in the content of the representation. But there are usually additional *extrinsic* mathematical elements that contribute to the representation beyond those that figure in the content. Suppose, for example, we have a representation of the dynamics of a particular physical system like a cannon ball. Here the intrinsic mathematics of the representation includes the differential equations with which the trajectory of the cannon ball is specified. These equations are used directly to fix how the world must be for the representation to be correct. At the same time, there may very well be different elements of mathematics relevant to the representation, both in its derivation from other representations and its further application in allowing the derivation of other representations. Understanding the contribution of this extrinsic mathematics is often just as important as understanding what the intrinsic mathematics is doing.

Our first two dimensions relate to the degree to which a given representation is abstract. One sense of abstractness which I will pursue is the degree to which a given representation purports to capture the genuine causal structure underlying a given phenomenon. A dynamical representation aims to include the causal details of a series of events, where causal effects propagate along the series in accordance with the mechanisms deployed in the representation.

Mathematics can play a central role here in allowing the different parts of the representation to track the genuine causal structure. For example, we may aim to represent the causal structure of the development of a traffic jam. The mathematical representation could accomplish this by having a series of real numbers standing for positions of the road. The representation would also include functions and mathematical rules or equations governing those functions corresponding to the causal laws about the propagation of cars along the road. If everything worked properly, these mathematical transformations would mimic the causal effects of the cars on one another, e.g. how the slowing down of one car would change the velocity of the cars behind it. We would then have a *concrete causal* representation of the system.

Here we have the sort of role for mathematics envisaged by the metaphysical conception. Its general viability depends in the end on how generous we are in populating the world with genuine causal relationships. Even the traffic case could be questioned by someone who restricted causes to relationships between fundamental physical entities. On this sparse approach to causes, a given car's slowing down is not the cause of its neighbor slowing down and the real cause of this event is a complicated chain of events involving the fundamental particles making up the cars and the drivers. There is little motivation to restrict causal relationships so severely and the most popular accounts of causes are much more liberal. Still, even on the most profligate picture of causes, there are mathematical representations where causes are absent. What I have in mind are representations which aim to represent the system at some designated equilibrium point or the respects in which the system is stable around such a point. I term such cases *abstract acasual* because they abstract away from the causal relationships found in concrete causal representations. In our traffic case, we might imagine a representation

of a given system of cars that simply presented what the system would look like after all the cars had achieved some fixed velocity for a given density of cars. Here we have two respects in which this abstract acausal representation could be mathematical corresponding to our intrinsic/extrinsic distinction. The specification of the velocity for the given density exhausts the intrinsic mathematics of the representation. But there are additional extrinsic mathematical elements that contribute to the representation such as the mathematical transformations that produce the abstract acausal representation. These might involve a study of the concrete causal representation of the traffic, but there is no general assurance that an abstract acausal representation must result from an associated concrete causal representation.

This is one sense of abstractness. A second dimension also involves the abstractness of the representation, but in a different sense of “abstract”. It sometimes turns out that there is a family of representations with a constant core of intrinsic mathematics, but where each representation in the family aims at a different kind of physical system. Such *abstract varying* representations can be contrasted with the *concrete fixed* representations that aim to capture the features of just one kind of physical system. With an abstract varying representation we can bring together a collection of representations in virtue of the mathematics that they share, and it can turn out that the success of this family of representations is partly due to this mathematical similarity. As an example of this we can take the mathematics associated with the harmonic oscillator model. The physical systems studied using this model range from simple systems like a pendulum and a spring through to domains like acoustic systems, electromagnetic fields and even nuclear phenomena (Marion & Thornton 1995, §3.7). What unites our representations of these different systems is the constant mathematical contribution. Exactly

how this mathematics might provide for the success of this sort of representation is not initially clear. But we should not make the mistake of assuming that mathematics helps with abstract varying cases in the same way that it helps with abstract acausal cases.

A third dimension along which we can consider the contribution of mathematics is *scale*. The scale of a representation is given by the size of the magnitudes involved in the representation. Spatial representations vary in scale from the very small, as with many quantum mechanical representations, through to the extremely large, as with much of general relativity. Time and energy are other magnitudes that can be represented at a wide range of different scales. The intriguing aspects of the scale of a representation often arise when we have successful representations of a given phenomenon at different scales and the features of these representations fail to fit together in a transparent way. For example, at the microscale we can represent a substance like water or air as a collection of molecules flying around and hitting each other. A macroscale representation of the same substance may treat it very differently, i.e. as a fluid that occupies space continuously. Even some of the magnitudes that the representations assign to the substance can differ. While both will assign a mass to the water, the macroscopic representation may also include the fluid's temperature and pressure distribution. Neither magnitude is a fundamental feature of the substance in the microscopic representation. A difference in scale, then, can involve a host of other changes in the content of what is represented.

In some sense the opposite sort of case where scale is important is where we represent a phenomenon as “self-similar” or scale invariant. We can use small-scale model ships, for example, to decide how to design large-scale ships because we take the representation of the relevant aspects of the motion of the

ship through water to be scale invariant. This is a special case of a family of abstract varying representations united by scale-invariant mathematics. But the success of this sort of abstract varying representation need not have any immediate implications for other sorts of abstract varying representations.

The relationship between representations of different scale is often specifiable using mathematics. Here, then, we have a contribution of extrinsic mathematics to the success of this or that representation. It can also happen that a single representation combines two or more representations where the links between these representations are mathematically significant. As we will see in chapter 5, in the boundary layer representation of fluid flow around an object like an airplane wing, one set of equations is used for most of the flow and another set of equations is used to represent the flow in the boundary layer close to the wing. The success of this hybrid representation depends on delicate considerations of scale. In effect, the two sets of equations are derived by different rescalings of a more realistic representation of the fluid flow. An understanding of scale and these rescaling techniques is essential, then, to seeing how the mathematics is successful in these sorts of cases. Exactly what the system has to be like in order for these manipulations to be successful is a delicate topic that I will return to.

Another sort of hybrid representation that involves considerations of scale is where two or more representations at different scales are tied together so that they mutually inform each other about the represented system's features. An intuitive example of this multi-scale representation is a representation of the Earth's climate. Climatologists are interested in the year-to-year behavior of the Earth's climate and the large-scale patterns of things like precipitation and temperature. Still, there is a complex interaction between these large-scale phenomena and small-scale phenomena like vegetation and

hurricanes. These interactions find their place via sophisticated representations that incorporate effects between these phenomena at different scales. Again, mathematical aspects of these representations are often central to their success, and so here is another place where we need to consider what the mathematics is doing and how it is doing it.

The final dimension along which to consider the contribution of mathematics to the success of a scientific representation is its constitutive vs. derivative character. Several philosophers of science, from Carnap and Kuhn to Michael Friedman, have argued for the need for a *constitutive framework* as a backdrop against which a given area of science with its *derivative* representations is possible. There is no clear agreement on exactly why these frameworks are needed or what role they perform, but the frameworks often reserve a prominent place for mathematics. For example, general relativity theory operates under the assumption that space-time can be accurately represented by the mathematics of differential geometry. Friedman presents this framework as a necessary presupposition for the rest of general relativity theory, most notably the more particular claims about the relationship between mass and the geometry of space-time encapsulated in Einstein's field equations (Friedman 2001). The claims of the constitutive framework, then, have a different status from the more derivative representations that are made from within the framework. We should be open to the possibility that the contribution of mathematics to the framework is different in kind from what we find with the more usual contributions to ordinary representations.

1.3 An Epistemic Solution

For a given mathematical representation we can start by asking, then, whether the mathematics in question is intrinsic or extrinsic. Then it is possible to determine how the mathematics contributes to (i) concrete causal or abstract acausal content, (ii) a concrete fixed or abstract varying content, (iii) issues of scale or (iv) a constitutive or derivative content. Some cases will combine different features from each of these choices while others will be best understood as more pure instances of a single element of our classification scheme. The hope is that we can say something informative about what mathematics is doing as we move through these components of our framework and that understanding each of these components will undermine the unsatisfactory social and metaphysical conceptions. The main conclusion that I will argue for in this book is that mathematics makes an *epistemic* contribution to the success of our scientific representations. Epistemic contributions include aiding in the confirmation of the accuracy of a given representation through prediction and experimentation. But they extend further into considerations of calibrating the content of a given representation to the evidence available, making an otherwise irresolvable problem tractable, and offering crucial insights into the nature of physical systems. So, even when mathematics is not playing the metaphysical role of isolating fundamentally mathematical structures inherent in the physical world, it can still be making an essential contribution to the success of science. For part of this success is the fact that we take our evidence to confirm the accuracy of our best scientific representations. And it is here that the mathematical character of these representations makes its decisive mark.

The vindication of this proposal depends on the careful consideration of a

range of examples which fit into the various parts of the classification scheme developed so far. Still, a *prima facie* case for this epistemic approach can be made by considering a few kinds of cases and what their epistemic contributions might be. To start, I return to the abstract acausal representations. Recall that in one instance of this sort of case the mathematics involves a representation of the system at some equilibrium point and may also include claims about the stability properties of such a point. In the traffic case, we consider the representation of the cars after they have reached some fixed velocity for a given density. There are two initially clear epistemic aspects of this acausal representation. Most obviously, if the acausal representation has been derived from a causal representation, then we can get some information about the system that was not explicit in the causal representation. We might learn, for example, at what density the maximum velocity is 50 km/h. The mathematics of the acausal representation allows us to know something about the traffic system that we might otherwise be ignorant of.

A second epistemic benefit of the acausal representation is somewhat less obvious. This is the highly restricted content of the acausal representation when compared to a causal representation of the same system. At an intuitive level, the acausal representation says less, and so takes fewer risks when representing the system. This impoverished content has an immediate benefit when we think about how we might confirm the acausal representation. Because it says less about the system, it will be easier to confirm. This aspect comes through clearly in our traffic case. If we had doubts about the accuracy of the acausal representation, we could remedy these doubts by carrying out the following experiment: set up a traffic system so that it corresponded to features of the causal model and see how the system stabilized as more cars were added to the road. Agreement on the density required to produce

a velocity of 50 km/h is then good evidence that the acausal representation is correct.

It might seem like the same point would hold for a causal representation of the traffic system. Why would the same experiment not confirm the causal representation to the same degree that it confirmed the acausal representation? A detailed answer would depend on an account of the causal claims, but it is not too controversial to assume that claims about causal chains involve more than stable features of the eventual result of the causal chain. Many accounts of causation emphasize counterfactual considerations, i.e. what would have happened had the system been different in this or that respect. Simply examining an actual traffic system does little to confirm these features of the causal representation and so to this extent it will be less well confirmed by the sort of experiment I have described.

There is another kind of case that also illustrates the relative ease of confirmation of acausal over causal representations. This is the case where we lack a causal representation of the development of the system, but are still able to formulate an acausal representation of what the system might look like at equilibrium. In a trivial sense, this sort of situation is one where the acausal representation can be confirmed and the causal representation cannot. While it might seem that such cases would be rare, they are easily found in the practice of science. Typically, causal representations of a given representation cannot be formulated because of their daunting complexity. And even when they can be presented, they may be untestable because the mathematics associated with the complex causal interactions proves intractable. As a result, the scientist or applied mathematician shifts to an acausal representation. It need not be a wholly acausal representation, as abstracting away from one family of causes need not involve a complete abstraction from

all causal representation. An example of this that we will return to is the representation of steady flow in fluid mechanics. Steady flow involves functions that do not vary over time. As a result, steady flow representations are acausal with respect to the causal features of the flow that depend on time. In many cases, we can confirm aspects of the steady flow representation of some fluid flow situation and we remain totally unable to formulate, let alone confirm, a representation of the time-dependent causes of the phenomena in question.

A related kind of case has received some discussion recently by philosophers under the name of “robustness analysis”.² Here we have a causal model, but we do not know how to set its parameters to produce an accurate representation. A second acausal model is then developed based on the features of the original model that are preserved under variations in the unknown parameters. This can lead to greater confidence in the acausal model than in any of the various causal models.

Mathematics can contribute to the epistemic goals of science, then, by contributing to the formulation of abstract acausal representations. The same result can be seen for the contributions of mathematics to abstract varying representations. Recall that an abstract varying representation was a family of representations of different kinds of physical systems with a common mathematical core. An important difference with the earlier abstract acausal representation is that an abstract varying representation has more content than an associated concrete fixed representation. This is because the varying representation includes representations of a wide variety of physical systems, while the fixed representation targets just one. Nevertheless, the varying representation can have epistemic advantages over the fixed representation

²Weisberg 2006 and Weisberg 2008.

and these advantages can be traced back to the mathematical character of the varying representation.

One sort of case involves a transfer of evidential support from one member of the family to another where the transfer is mediated by the mathematical similarities between the parts of the abstract varying representation. Suppose, for example, that we have thoroughly explored the simple harmonic oscillator representation of the pendulum. The core features of this representation are that the pendulum has an equilibrium point, when the bob is pointing straight down, and the force restoring it to this point is proportional to the displacement from equilibrium. These features allow us to classify different kinds of pendulums and accurately predict the motion of the bobs across a wide variety of changes in mass, initial velocity and angular displacement. Using experiments, we can confirm the accuracy of this representation in a wide variety of respects. Suppose further that we have arrived at a representation of a certain kind of spring where the spring representation is mathematically similar in the relevant aspects to the successful pendulum representation. That is, there is an equilibrium point and a linear restoring force. With the spring the equilibrium and the linear restoring force are determined by the materials making up the spring. Still, the mathematical similarities between the representations allow us to transfer our confidence in the pendulum representation over to the spring representation. Confirmation flows, then, within the family of representations making up an abstract varying representation. We can come to be more confident in the members of the family in virtue of their mathematical similarities than would have been possible via the consideration of the available data and each representation individually.

This process of abstraction and widening of the scope of a family of repre-

sentations can be extended in a much more radical direction. So far we have only considered the relatively modest move of varying the restoring force from the force restoring the pendulum to equilibrium to the force restoring the spring to equilibrium. An even more abstract perspective involves considering the family of representations where the restoring force is linear, no matter what its physical basis is. It is this fully abstract representation that is typically described as the simple harmonic oscillator model. Once it has been mathematically isolated, it can be investigated in a rigorous fashion completely independently of its physical manifestations. That is, we can develop a wholly mathematical theory of this model in its own right. This provides us with a host of claims that we can reasonably believe describe any physical system once we determine, either experimentally or otherwise, that it meets the minimal conditions to be a member of the family of representations. A model, mathematically identified, can lead to a huge range of well-confirmed representations of vastly different physical systems. It is this abstract unifying power which allows the harmonic oscillator to encompass the wide range of physical systems mentioned earlier.³

Issues connected to the scale of a representation are more difficult to survey as the way scale is deployed can vary so much from case to case. As we have already seen, sometimes a representation is of a single scale, while in other cases the representation mixes scales or purports to be scale invariant. The main point to emphasize at this stage is that mathematics can play a central role in both the intrinsic specification of the scale or scales of the representation and in the extrinsic derivation of a representation where scale is significant. As with the two sorts of abstract representations already con-

³The value of this approach to modeling has also been emphasized by Weisberg at, e.g., Weisberg 2007, §4.

sidered, epistemic benefits may result. In the simplest case, the mathematics may lead to a representation that is highly restricted in content because it represents the system only at a single scale. When this occurs, we have the same benefits that we saw for the abstract acausal representation. Because the content is so restricted, it is easier to confirm that the representation is accurate. At the other extreme, we may derive a mathematical representation that purports to be scale invariant. Successful scale invariant representations clarify under which conditions the phenomenon will be unaffected by scale. As a result, scale invariant representations take on all the positive features of the abstract varying representations just considered. We can transfer the evidence for one representation in the family to another representation in the family. This is, after all, just what happens when we decide how to design a ship based on experiments with a small-scale ship. What varies in such a family is not the interpretation of the elements of the mathematical representation, but instead the scale at which the magnitudes associated with the elements are interrelated. Other multi-scale representations involve further epistemic benefits which we will return to later.

Finally, we come to the constitutive vs. derivative dimension of a mathematical representation. As with scale considerations, here the issues are complex. It is not even clear if a different kind of framework representation is necessary in addition to ordinary representations, and even if such representations are needed, there is no consensus on why they are needed. I will argue, though, that there are decisive reasons in favor of requiring constitutive frameworks and that these reasons involve issues of the confirmation of derivative representations. This strand is prominent in Kuhn and Friedman, but not always as clearly articulated as one might like. Still, here is one summary of the position I will be defending: “A constitutive framework . . .

defines a space of empirical possibilities . . . , and the procedure of empirical testing against the background of such a framework then functions to indicate which empirical possibilities are actually realized” (Friedman 2001, p. 84). Very roughly, the mathematics deployed in a framework for a given domain has the effect of representing a determinate range of states of affairs as physically possible. Taking this framework for granted, more derivative representations can then be confirmed by comparing data against these physical possibilities. A central example for Friedman is the framework representation of the general theory of relativity according to which the structure of space-time is identified with a member of the family of a certain kind of manifold. Delicate issues remain, though, about how mathematical these frameworks really need to be, and also the basis for the confirmation of the frameworks themselves.

While none of the discussion so far is conclusive, it at least holds out the hope of understanding how a given mathematical representation contributes to the overall success of science. The general outlines of the account make clear that mathematics has a more important role than just facilitating the consensus of scientists, as the social conception suggests. It is also consistent with a more modest place for mathematics than the metaphysical conception requires. Even when the mathematics fails to track fundamental features of the world, it still affords us access to successful scientific representations. The contribution of the mathematics is in helping us to formulate and investigate representations that we can actually confirm. This epistemic role puts mathematics at the center of science without the appeal to dubious metaphysical assumptions.

1.4 Explanatory Contributions

There is a growing debate on the existence of mathematical explanations of physical phenomena. Advocates of the explanatory power of mathematics in science often present their examples as part of the so-called indispensability argument for mathematical platonism. Critics object either that the alleged explanations are not especially mathematical, or else that non-platonists about mathematics can also account for the explanatory power in the examples. The issue is pressing for the account I aim to develop because of the complex interactions between issues of explanation and confirmation. After surveying some of these links, I turn to the place of explanatory issues in the epistemic account of the contributions of mathematics to scientific representation. In the next section I consider the links between this project and the indispensability argument.

Accounts of explanation can be divided into pragmatic and objectivist approaches. A pragmatic approach conceives of explanation as something that we do with representations that we have already elected to accept based on other virtues. As a result, explanatory power is not itself a consideration that should be used to decide which representation to accept. The most prominent example of a pragmatic approach is van Fraassen (van Fraassen 1980). He argues that explanations are best conceived as answers to why-questions. Various different why-questions receive different scientific answers using the representations that we have already adopted. But just because a representation can contribute to an answer to a why-question, we should not conclude that it is likely to be true, or more likely to be true than its competitors. Van Fraassen's pragmatic conception of explanation, then, undercuts inference to the best explanation.

Opposed to these pragmatic accounts, we find objectivists about explanation. They place additional constraints on an acceptable explanation and so are able to draw stronger conclusions than van Fraassen from the existence of an explanation of a phenomenon. The most ambitious versions of an objectivist approach insist that our best available explanation is likely to be true. This allows one to start with more or less uncontroversial cases of successful explanations and conclude that the ingredients of the explanation receive a boost in confirmation simply in virtue of appearing in the best explanation available. Two strategies for reaching this more substantial conception of explanation are especially popular. The first insists that a good explanation must give the cause of what is being explained, while the second requires that a good explanation be part of the best scheme of unifying a wide range of phenomena. Causal and unification approaches to explanation consequently disagree on which examples are cases of genuine explanation, but both vindicate some form of inference to the best explanation for their respective explanations.

The three different approaches to explanation surveyed here will have quite different reactions to the classification scheme developed in the previous sections. A pragmatic approach can concede that mathematics contributes to successful explanations, but insist that this is only because mathematics appears in our best representations as determined by criteria besides explanatory power. There is no deep connection between the mathematics and the explanatory power on the pragmatic approach in line with the more general resistance to draw any substantial consequences from explanatory power. Advocates of the causal and unification approach will likely to be tempted to draw more interesting conclusions. If we insist that all explanations involve causes, and grant that there are cases where the mathematics contributes

to explanation, then it seems inevitable that we conclude that there are mathematical causes of physical phenomena. From this perspective, we have independent confirmation of the metaphysical conception. That is, the physical world is essentially mathematical, and so when we explain phenomena, we will need to invoke the mathematical entities that make up, in part, the system in question.

The puzzling interpretation of this kind of conclusion has prompted most advocates of mathematical explanation to go beyond causal explanation. One strategy is to emphasize the unifying powers of mathematics and claim that mathematics can contribute to explanation precisely because it can unify diverse scientific phenomena.⁴ We can place this position in our scheme by aligning it with the abstract varying representations such as the simple harmonic oscillator model. These representations bring together representations of vastly different physical systems. In virtue of linking these representations, the abstract varying representation is able to represent the similarities between these systems. Presumably, it is these similarities that form the core appeal of the claim that unification is the key to explanation. We have also seen in outline how an abstract varying representation can mediate the confirmation of its different components. If this picture can be sustained, then we can understand the alleged link between explanatory unification and a boost in confirmation. This would allow the justification of a limited form of inference to the best explanation.

There are two problems with this happy reconciliation, however. First, there seem to be mathematical explanations that resist incorporation into the unification framework. I turn to these shortly. Second, there are *prima facie* convincing cases of unified mathematical representations that fail to be

⁴See Kitcher 1989, discussed in Mancosu 2008*a*.

explanatory (Morrison 2000). The suspicion here is that bringing together a family of representations via some shared mathematics may take us away from the relevant details of the different systems, and so actually decrease the overall explanatory power of the different representations. This may be because all explanation is really causal, or else because the mathematics unifies the phenomena in a misleading way. Either way, if these examples stand up to scrutiny, it will not be possible to trace the explanatory power of the mathematics in all cases to its unifying power.

What sorts of cases resist the unification framework? One example is the topological explanation for why there is no non-overlapping return path across the bridges of Königsberg.⁵ It is quite artificial to force this explanation into the unification model, but at the same time it seems fairly clear that there is something mathematical and non-causal going on. Here we can appeal to another slot of our classification scheme, namely abstract acausal representations. With the bridges we have abstracted away from the material causes of the bridges being in the shape that they are in, and represented them as they are independently of the details of their construction. This can be thought of as a representation of the system at equilibrium in an extended sense where the structure of the system is stable under changes in its material construction. As with the abstract varying cases, we have seen how abstract acausal representations can lead to an increase in confirmation. If that story can be applied to these explanatory cases as well, then we have a different way to vindicate inference to the best explanation. Abstract acausal representations can figure in explanations and at the same time receive a boost in confirmation.

There are, then, some prospects for linking the picture of mathematical

⁵This example is discussed later in chapter 3.

representation that I have sketched to discussions of mathematical explanation and perhaps even using this approach to ground some form of inference to the best explanation. In what follows, though, I will proceed somewhat cautiously. As there is no agreement on cases or accounts of explanation, I will not assume that mathematics contributes explanatory power in this or that example. Instead my focus will be on confirmation. Here there is also disagreement, but the question of confirmation seems more tractable and amenable to a naturalistic treatment via a careful study of the actual practice of scientists. If, as I hope will be the case, there are substantial links that can be drawn between boosts in confirmation and apparent mathematical explanations, then we will have a basis for further discussion about the explanatory contributions of mathematics. Still, if it turns out that there are no such links, the point about the contribution of mathematics to confirmation can be independently defended.

1.5 Other Approaches

There are a variety of other ways to approach the problem of mathematics in science, and it is worth setting out how my strategy promises to link up with these different approaches. The most prominent alternative approach focuses on the correct metaphysical interpretation of pure mathematics. Is pure mathematics best thought of as involving abstract objects, as the platonist argues? Many nominalist accounts of mathematics aim to interpret pure mathematics so that it does not require the existence of any abstract objects. Most nominalist interpretations of mathematics preserve the intuitive objective truth-values associated with mathematical statements, but there are some more recent non-standard approaches that I will group under the

label of “fictionalist” that do not. Fictionalists provide an account of pure mathematics that preserves some features of ordinary mathematical practice, but give up the standard truth-values we would assign to some statements. That is, they are not truth-value realists, while the standard nominalists remain realists about truth-values, even if they are not realists about abstract objects. Platonists are realists both about truth-values and abstract objects.

The indispensability argument for platonism argues that we should accept the existence of abstract objects based on their essential or indispensable contribution to our best scientific theories (Colyvan 2001). This contribution can be characterized in different ways. Initially with Quine and Putnam the contribution was specified in terms of the formulation of our best scientific theories. That is, it was not possible to formulate acceptable versions of these theories that did not involve mathematical terms. More recently, the focus has shifted to the explanatory power of mathematics. That is, there are superior explanations available from within our mathematical scientific theories. Either way, the indispensable contribution of mathematics is supposed to require that we believe in the existence of abstract mathematical objects.

From the perspective of this project, the problem with the positions of both advocates and critics of these indispensability arguments is that they have not considered the prior question that I have put at center stage: what is the contribution that mathematics makes to this or that scientific representation? A proper evaluation of the role of mathematics in science in determining the interpretation of pure mathematics turns on delicate considerations of the content of mathematical scientific representations. The platonist picture seems to be that it is only by interpreting mathematics in terms of abstract objects that we can make sense of these contents. Unfortu-

nately, the platonist advocates of indispensability have only recently turned to a clarification of what these contents are (Bueno & Colyvan 2010). Critics of indispensability have done little better. There was initially a laudable attempt to provide apparently non-mathematical formulations of some of our best scientific theories (Field 1980). Although most critics of indispensability arguments would grant that these attempts failed, there is little attempt to diagnose why they failed. Instead, many anti-platonists have gone on to insist that they do not need to reformulate these theories, but can accept them without giving any account of the content of the claims of the theories. In certain cases, some discussion of the contents has been given, but only for simple examples with no indication of how to extend this to more interesting cases.⁶

The benefit of formulating the issue as I have presented it, then, is that the content of these scientific representations takes center stage. We have already seen in broad outline an argument that the contribution of mathematics is epistemic based on the different ways in which it contributes to the content of our representations. As I will argue in the next chapter, the content can be further characterized in terms of broadly structural relationships between physical situations and mathematical structures. While there are many details to work out, this picture will support some criticisms of the indispensability argument. For if the content can be understood in this structural way, the accuracy of these representations will not require much from the mathematical objects. In particular, applications will turn on extrinsic features of the mathematical entities, and so they will be unable to pin down the intrinsic features of these entities, including whether they are abstract. Such a conclusion does not support all forms of anti-platonism, however. For

⁶See, e.g., Balaguer 1998. I return to this issue in chapter 12.

I will argue that fictionalists are not able to recover the contents of our most successful scientific representations, at least on the account of their content that I will develop in chapter 2. Until they respond to these concerns about content, then, we have a strong reason to reject fictionalist interpretations of pure mathematics. It remains possible that some non-fictionalist, anti-platonist interpretation can be sustained. Whether or not this is possible seems to me to turn largely on questions internal to pure mathematics and its epistemology. I conclude this book with an all too brief discussion of these questions and their independence from reflection on applications of mathematics.

Beyond the indispensability argument, the most well-known philosophical debate surrounding mathematical scientific representations concerns what is deemed the unreasonable effectiveness of the contribution of mathematics to science. Wigner famously concluded from a survey of impressive scientific discoveries involving mathematics 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). It is not clear exactly what Wigner finds mysterious when it comes to the interactions between mathematics and physics, but the examples he emphasizes turn on the development of sophisticated mathematics which was only later to prove crucial in the formulation of successful scientific theories. Unlike the calculus, for example, which was discovered as part of attempts to understand physical phenomena like motion, the twentieth century saw the application of domains of mathematics which were developed based on internal mathematical considerations. A simple example of this is the theory of complex numbers. They were introduced into mathematics based on reflections on the real numbers. Rounding out the real numbers to include

complex numbers contributed to a superior mathematical theory. The complex numbers lacked any intuitive physical significance, but in spite of this they provide a crucial part of the mathematics deployed in science. The mystery, then, is that mathematics developed independently of applications should prove so effective when it is eventually applied.

Part of the problem with Wigner's conclusions, and the resulting debate, is that it is often not clear exactly what sort of contribution is supposed to be surprising or unreasonable. Steiner has rectified many of the flaws of Wigner's original discussion by focusing on the role of mathematics in suggesting how to formulate new scientific theories (Steiner 1998). The conclusion of Steiner's argument is correspondingly much more definite and controversial than any of the vague suggestions by Wigner. Steiner concludes that this role for mathematics undermines certain forms of naturalism. Naturalism, in Steiner's argument, is the view that the universe is not constituted in a way that fits with our cognitive faculties. Anti-naturalists, then, include not only theists who believe that the human mind has been designed to understand the natural world, but also those whom Steiner calls "Pythagoreans" who insist that the universe is formed in line with the mathematics that our cognitive faculties find of interest.

The main premises of Steiner's anti-naturalist argument involve historical cases where mathematical considerations were crucial in getting a scientist to formulate a theory that she deemed worth testing. These mathematical considerations, in turn, are linked to species-specific aesthetic judgments. The combination of these two premises entails for Steiner that the scientist acted as if species-specific judgments would give good insight into the way the physical world works. This challenges the scientists' own claims to be naturalists. More seriously, Steiner argues that the overall success of this

strategy for discovering new scientific theories should undermine our own confidence in naturalism.

On a first pass, there is little connection between the question that I have asked about the contributions of mathematics to successful scientific representations and Steiner's discovery problem. Most obviously, the discovery issue concerns how or why a scientist came to formulate a representation that she judged worth testing. As with other questions of innovation and discovery, we often have little understanding of how this occurs and there is no reason to think that the ultimate story of a given discovery would not rely on subjective psychological considerations like intuition or "gut feeling". There should then be no mystery about why a scientist employed methods of discovery without a rational foundation or how these methods were successful in some cases. By contrast, my focus has been on the way mathematics helps a scientist to confirm a representation through ordinary experimental testing. This confirmation or justification is independent of the origins of the representation or why the scientist decided to test it.

Still, this division between discovery and justification is hard to enforce when we turn to the history of science and actual scientific practice. It may turn out that reflection on how mathematics makes its epistemic contribution to science helps us to understand the sorts of cases discussed by Steiner. To see how this might work, recall the discussion of abstract varying representations like the simple harmonic oscillator. Once it is validated for certain pendulum systems, we can transfer the confirmation from this domain to a new domain like spring systems based on the claimed mathematical similarities. Now, independently of the actual historical details of the scientific understanding of pendulum and spring systems, it certainly could have happened that a scientist sought to represent springs as harmonic oscillators

based on this kind of analogical reasoning. If we can understand how confirmation can be transferred in this sort of way, then this discovery procedure becomes far less puzzling. For the scientist is in a position to recognize that if this way of representing the spring is successful, then the family of well confirmed harmonic oscillator systems will be expanded.

While this is just a quick sketch of how the epistemic conception might bear on Steiner's argument, it is enough to show that an answer to my representation question might provide a rationale for some of the discovery procedures that Steiner considers. If this is right, and the epistemic approach does not violate naturalism, then there is something wrong with Steiner's argument. The crucial flaw, I will argue, is the link between mathematical theories and species-specific aesthetic judgments. If this link were correct, then mathematics would be ill-suited to contribute to successful scientific representations. The connection between mathematics and confirmation established by the epistemic conception requires that mathematics be objective, and further reflection on the practice of pure mathematics reinforces this objectivity. So, we can block Steiner's anti-naturalist conclusion by focusing on mathematical representations and clarifying the link between pure mathematics and aesthetic judgment.

1.6 Interpretative Flexibility

Discussion of Wigner's and Steiner's worries about the success of appealing to mathematics leads one naturally to wonder just how successful mathematics has really been in aiding science. In fact, it might seem like the whole way in which I have formulated my main topic has been biased towards the positive role of mathematics in science. Perhaps in addition to the contributions that

mathematics makes to the success of science, there are also many ways in which mathematics hinders the development of accurate scientific representations. We have seen this worry already in a specific case, where it seems that unifying diverse representations using some single mathematical framework can undermine the usefulness of the overall representation. The thought is that mathematics is a tool that sometimes works and sometimes does not, and to understand its contribution in any realistic way we must take the bad along with the good. The failures of mathematics to help science are many and even though these failures are not always celebrated in textbooks and histories along with the successes, a philosophical understanding of our issue cannot ignore them.

I will argue that these failures are to be expected. The features of mathematics that make it possible for mathematics to make a positive contribution also allow its negative contributions. Above all, the *interpretative flexibility* of mathematical representations creates the opportunities exploited by the epistemic conception as well as the sorts of failures that result from an overly confident reliance on mathematics to guide our scientific representations. This can happen in many ways, but for now notice only that the abstract nature of these representations, in both the acausal and varying senses discussed so far, permit these representations to be interpreted in ways that outstrip their well-confirmed core. There is a risk in taking the well-confirmed aspects of a given representation and extrapolating our confidence to other aspects that are completely unconfirmed. To consider two simple examples, we might start with a causal concrete representation and obtain an associated abstract acausal representation. If the abstract acausal representation is confirmed, then there is a temptation to take this as evidence of the correctness of the causal concrete representation. Here

the mathematical link between the acausal and the causal encourages us to carry over the confirmation even if the link makes it clear that the acausal representation is easier to confirm. A second kind of simple case involves concrete fixed and abstract varying representations like a pendulum and the harmonic oscillator. When we transfer our confidence to the harmonic oscillator representation from the pendulum to the spring, we must be careful not to transfer too much of the pendulum interpretation over to the spring representation. For example, we cannot assume that the force restoring the bob to the pendulum's equilibrium is the same kind of force as that restoring the spring to its equilibrium. A scientist that understands both systems well will of course not make this mistake, but in other sorts of cases the error is easy to make. We over interpret the members of the family of the representation and wind up thinking that the systems have greater physical similarities underlying their mathematical similarities.

Mathematics contributes flexibility, then, and this flexibility creates the opportunity for scientific failures. Often these failures will escape detection for some time because they involve features that are remote from our experimental access. It follows that there is no simple path to well-confirmed scientific representations using mathematics. As with everything else, it all turns on considerations arising from eventual experimental testing, and the realist's hope must be that no matter what interpretative errors have been made at a given stage of scientific theorizing, there will be a later stage when this error can be isolated and rectified.

A final caveat concerning the project undertaken here concerns the approach to confirmation that I will appeal to. There can be a whiff of paradox in asserting that mathematics contributes to the confirmation of our best scientific theories, just as we saw there was some mystery with the role of

mathematics in discovery with Wigner and Steiner. On the discovery side, the puzzle came from the way in which mathematicians allegedly anticipated the needs of scientists. With confirmation, it might seem impossible for mathematical beliefs to play any positive role in supporting scientific claims. For mathematical beliefs are not about the physical world, and it is commonly assumed that mathematical truths are necessary, i.e. true no matter how the physical world happens to be. If this is right, then how can mathematical truths have any relevance to the truth about the physical world and how can believing some mathematics help to tell us which way our evidence bears on a given physical possibility?

As we will see in some detail in chapter 2, the problem is not remedied by drawing on any of the main contenders for a comprehensive account of confirmation. These theories try to explain when boosts in confirmation occur, but they are ill-equipped to factor in a role for mathematical beliefs in this process. To take a simple case, a scientist may assign some prior probability to a mathematical scientific representation T in advance of experimental testing. But suppose that the scientist must know another mathematical theory M in order to extract any predictions from this representation. There is little to no discussion in the confirmation literature about how coming to believe M can allow the scientist to determine the predictions and actually test the theory. The main culprit here is the assumption of logical omniscience, i.e. that scientists should assign the same probability to logically equivalent propositions. As mathematical truths are metaphysically necessary, it seems to follow that they are all assigned the probabilities of logical truths. But this obscures an important place where mathematics is making its mark, namely in helping us to see what implies what. Questions of logical entailment need to be distinguished from the mathematical analysis of the

content of representations.

A new approach is needed, then, to understand how mathematics can help us to confirm scientific representations. My approach will not involve a new comprehensive theory of scientific confirmation. Instead, I will rely on various principles that relate the content of a representation and the evidence in different sorts of cases. These principles have some prima facie plausibility and will be illustrated in the case studies I develop in this book. Some accord well with certain extant theories of confirmation, but when there is a conflict with these theories, I will argue that it is the theories that are wrong, not the principles. It is only in this way that we can isolate when mathematics makes its contribution to confirmation and come to understand what philosophers of science seem to have overlooked in their reflections on scientific confirmation to date.

1.7 Key Claims

The goal of this introductory chapter has been to set out the main claims that I will argue for in the rest of this book and to situate my project with respect to some other topics associated with mathematics and scientific representation. Here are twelve claims which I am to clarify and defend in the course of this book:

1. A promising way to make sense of the way in which mathematics contributes to the success of science is by distinguishing several different contributions (ch. 1).
2. These contributions can be individuated in terms of the contents of mathematical scientific representations (ch. 2).

3. A list of these contributions should include at least the following five: concrete causal, abstract acausal, abstract varying, scaling and constitutive (ch. 3-6).
4. For each of these contributions a case can be made that the contribution to the content provides an epistemic aid to the scientist. These epistemic contributions come in several kinds (ch. 3-6).
5. At the same time, these very contributions can lead to scientific failures, thus complicating any general form of scientific realism for representations which deploy mathematics (ch. 7).
6. Mathematics does not play any mysterious role in the discovery of new scientific theories. This point is consistent with a family of abstract varying representations having some limited benefits in suggesting new representations which are worthy of testing (ch. 8).
7. The strongest form of indispensability argument considers the contributions I have emphasized and argues for realism of truth-value for mathematical claims (ch. 9).
8. These contributions can be linked to explanatory power, so we can articulate an explanatory indispensability argument for mathematical realism (ch. 10).
9. However, even such an argument based on explanatory contributions faces the challenge of articulating a plausible form of inference to the best explanation (IBE) which can support mathematical claims (ch. 10).

10. This challenge to IBE for mathematical claims is consistent with mathematics contributing to successful IBE for non-mathematical claims, as in the extended example of the rainbow (ch. 11).
11. Fictionalist approaches to mathematics and scientific models face challenges which undermine their main motivations (ch. 12).
12. The way in which our physical and mathematical concepts relate to their referents suggests that our representations depend for their contents both on our grasp of concepts and our beliefs (ch. 13).

I end this chapter with an outline of the remaining chapters. In chapter 2 I develop further two of the central areas of my project: content and confirmation. The content of a given scientific representation turns on a variety of considerations that need to be clarified before we can investigate the contribution of mathematics. A main aim of this discussion is to sharpen the intrinsic/extrinsic distinction that I have worked with informally so far. Another priority is confirmation. This is a topic of intense discussion although, as I have just noted, little of it is directly relevant to my main concerns. Still, a survey of some views on confirmation and the necessary principles needed to make sense of mathematics in science must be on the table to vindicate the epistemic conception.

After these preliminaries, I turn to four chapters exploring the four dimensions of mathematical contributions described above. Chapter 3, on the concrete causal/abstract acausal distinction, considers several examples along this spectrum as well as some influential accounts of causation and its representation. In chapter 4 I bring in the other sense of abstractness with the concrete fixed/abstract varying dimension. Again, examples are used to illustrate the differences between these sorts of representations and

the epistemic benefits of the mathematics. Chapter 5 continues this theme with a survey of examples that involve scale in different ways. The constitutive/derivative distinction in chapter 6 rounds out this positive survey of how mathematics makes its positive contribution. A cautionary chapter 7 emphasizes the respects in which things can go wrong by discussing several cases where mathematics seems to have contributed to the failure of a given representation. These failures are diagnosed as a consequence of the interpretative flexibility which goes hand in hand with using mathematics in science.

The next six chapters link up the epistemic conception with other discussions of the role of mathematics of science. I start, in chapter 8, by surveying the worries about discovery raised by Wigner and Steiner. Chapter 9 turns in earnest to the debates about the indispensability argument for platonism. The more recent focus on mathematical explanations in science and their significance is considered in chapter 10. An extended case study involving inference to the best explanation makes up chapter 11. Chapter 12 considers the anti-platonist views I group under the label of “fictionalism” and their prospects for adopting the epistemic conception of the contribution of mathematics to science. Finally, in chapter 13 I consider relaxing the assumptions introduced in chapter 2 about how the content of a representation is determined. It will turn out that these assumptions are somewhat unrealistic, but I will argue that this does not affect the lessons drawn from the examples discussed in earlier chapters. Here I engage with Wilson’s picture of scientific representations as facades and conclude that even if Wilson is correct, we can still understand the contribution of mathematics in epistemic terms (Wilson 2006).

The book concludes with chapter 14 on the implications of the previous chapters for our understanding of pure mathematics. It will turn out that the role of mathematics in science has little to no bearing on our understanding of pure mathematics, but even this negative conclusion gives indirect support to the approach to pure mathematics known as structuralism. On the epistemic side, I will argue that much of pure mathematics must be a priori if it is to make its contributions to the success of science.