Who is a Modeler?*

Michael Weisberg

Final Draft to appear in *British Journal for Philosophy of Science*

Abstract

Many standard philosophical accounts of scientific practice fail to distinguish between modeling and other types of theory construction. This failure is unfortunate because there are important contrasts among the goals, procedures, and representations employed by modelers and other kinds of theorists. We can see some of these difference intuitively when we reflect on the methods of theorists such as Vito Volterra and Linus Pauling on one hand, and Charles Darwin and Dimitri Mendeleev on the other. Much of Volterra's and Pauling's work involved modeling; much of Darwin's and Mendeleev's did not. In order to capture this distinction, I consider two examples of theory construction in detail: Volterra's treatment of post-WWI fishery dynamics and Mendeleev's construction of the periodic system. I argue that modeling can be distinguished from other forms of theorizing by the *procedures* modelers use to represent and to study real-world phenomena: *indirect* representation and analysis. This differentiation between modelers and non-modelers is one component of the larger project of understanding the practice of modeling, its distinctive features, and the the strategies of abstraction and idealization it employs.

*I have benefited tremendously from the helpful comments of Jerome Berson, Barry Carpenter, Zolton Domotor, Edward Epsen, Ben Escoto, Melinda Fagan, Ron Giere, Steven Gross, Gary Hatfield, Roald Hoffmann, Paul Humphreys, Jason Rheins, Tania Lombrozo, Jay Odenbaugh, Angela Potochnik, Scott Weinstein, Rasmus Winther, and especially Peter Godfrey-Smith, Elisabeth Lloyd, and Deena Skolnick, as well as two anonymous referees. Earlier versions of this paper were presented at the Institute for Research in Cognitive Science (Penn), Indiana University, and the Center for Philosophy of Science (Pitt). I thank those audiences for vigorous discussion and helpful suggestions.

Contents

1 Introduction

After the first World War, there was an unusual shortage of certain types of fish in the Adriatic sea. This seemed especially strange because during the war, fishing had dropped off considerably. Most Italians believed that this should have given the fish a chance to increase their numbers. The well-known Italian biologist Umberto D'Ancona was on the case. After carefully analyzing the statistics of fish markets he discovered a very interesting fact: the population of sharks, rays, and other predators had increased during the war while the population of squid, several types of cod, and Norwegian lobster had decreased. How could this be? Why did the small amount of fishing associated with the war favor the sharks?

D'Ancona brought this question to his father-in-law, Senator Vito Volterra, who held a chair of Mathematical Physics at Rome. Volterra approached the problem with what I will call the *modeling strategy* or more simply *modeling*. He imagined a simple biological system composed of one population of predators and one population of prey. He attributed to these populations just a few properties and wrote down mathematical expressions describing them. After carefully studying the dynamics of his model populations, Volterra knew why war seemed to favor the sharks and he had good, yet surprising news: Resumption of heavy fishing would cause the populations to return to their pre-war proportions. (Volterra, 1926a)

The strategy employed by Volterra is a common one found in scientific disciplines that face the difficulty of describing, explaining, and making predictions about complex phenomena. When faced with such complexity, theorists can employ one of several strategies: They can try to include as much complexity as possible in their theoretical representations. They can make strategic decisions about which aspects of a phenomenon can be legitimately excluded from a representation. Or they can model, studying a complex phenomenon in the real world by first constructing and then studying a model of the phenomenon. Characterizing modeling as a distinct kind of theorizing is the main subject of this paper.

There are many insightful discussions in the philosophical literature about the nature of models and the structure of theories, where theories are understood to be composed of models. Although less has been written explicitly about the *practice* of theorizing, substantial discussions of these issues can be found in the writings of Hesse (1966), Wimsatt (1987), Cartwright (1983, 1989), Giere (1988), Morgan and Morrison (1999), Winther (forthcoming), and others. These discussions concern many different aspects of models and modeling. Hesse gives us an account of the role models can play in analogical reasoning; Wimsatt explains the role of false models in building more accurate theories; Cartwrght and Giere describe how models help ground scientific explanation and prediction; Morgan and Morrison discuss how models often serve as autonomous instruments and as mediators between theoretical structures and the world. This paper considers a somewhat different issue and, as such, defends a different sort of thesis. I contend that there are important distinctions to be made among the different *strategies of theorizing*. One of these strategies involves modeling, but there is at least one other as well: what I will call *abstract direct representation* (ADR). The primary aim of this paper is to articulate some of the differences between these strategies, so as to give an account of what is distinct about modeling.

2 The Essential Contrast

In order to characterize the practice of modeling and its distinctive features, a natural place to begin is with the rich literature about the structure of models and theories. Beginning with the writings of Patrick Suppes and colleagues, philosophers committed to the *semantic view* have reconstructed physical, psychological, and biological theories as sets of models, rather than as sets of axiomatic sentences. (Suppe, 1977; Domotor, 2001) There is a good deal of variety among proponents of the semantic view as to the nature of scientific models. Suppes and his colleagues have argued that scientific models are simply logician's set-theoretic models. (Suppes, 1960) Later philosophers in this tradition sought to describe models in terms closer to those actually used by scientists conceiving of models either as sets of trajectories through a state space (van Fraassen, 1991; Lloyd, 1994) or as systems that would be concrete if they were real (Giere, 1988).[1]

Proponents of the semantic view are primarily giving accounts of the structure of theories. However, insofar as these accounts make implicit claims about the practice of theorizing, they treat this practice as univocal, focusing exclusively on model-based representation. If we set aside the question of whether all theories are best reconstructed as sets of models, we can still ask how scientists go about theory construction and whether or not this theory construction explicitly depends on models. In arguing that modeling is just one kind, albeit an important kind, of theorizing, I am arguing that some theoretical practices depend explicitly on the construction and analysis of models, while others do not.

Modeling, I will argue, is the indirect theoretical investigation of a real world phenomenon using the model. This happens in three stages. In the first stage, a theorist constructs a model. In the second, she analyzes, refines, and further articulates the properties and dynamics of the model. Finally, in the third stage, she assesses the relationship between the model and the world if such an assessment is appropriate. If the model is sufficiently similar to the world, then the analysis of the model is also, indirectly, an analysis of the properties of the real world phenomenon. Hence modeling involves indirect representation and analysis of real world phenomena via the mediation of models.

This is not the only way to construct a theory. Phenomena can also be represented and analyzed without the mediation of a model. I will refer to this non-model based form of theorizing as *abstract direct representation* or ADR. The similarities and differences in these two forms of theorizing are best appreciated by looking at examples in detail. Volterra's explanation of the cause of the post-WWI fish shortage will serve as our primary example of modeling, and Dimitri Mendeleev's explanation of chemical properties via their periodic dependence on atomic mass will serve as our primary example

[1] For a review of the diverse accounts of models defended by proponents of the semantic view see Downes (1992) and Odenbaugh (forthcoming).

of ADR.

2.1 Modeling

Volterra was a modeler. He studied the dynamics of real Adriatic fish by first studying the properties of two model populations. Unlike the myriad properties possessed by two real populations of organisms, Volterra's model organisms possessed just a few properties such as an intrinsic exponential growth rate for the prey in the absence of predators and a constant death rate for the predators. (Roughgarden, 1979, 434) He did not arrive at these model populations by abstracting away properties of real fish, he *constructed* these model populations by stipulating certain properties.

Volterra gave a mathematical description of these model populations in the form of two coupled differential equations:

$$\frac{dV}{dt} = rV - (aV)P \tag{1}$$

$$\frac{dP}{dt} = b(aV)P - mP \tag{2}$$

In these equations, V is the size of the prey population and P is the size of the predator population. The parameter r stands for the intrinsic growth rate of the prey population and m, the intrinsic death rate of the predators. The other parameters (b and a) have to do with the functional response and correspond to the prey capture rate and the rate at which each predator converts captured prey into more predator births. (Roughgarden, 1979, 432)

Equations (1) and (2) describe a model in which the population of both predators and prey exhibit periodic oscillations in their sizes. Since there is no stable equilibrium described by these equations, we know that the size of the model populations will oscillate indefinitely. The equations do admit of one unstable equilibrium solution which has a useful property: it corresponds to the time-averaged size of the two model populations. (Hofbauer and Sigmund, 1998, 16) Thus the first step to solving D'Ancona's puzzle is to solve the equations for this equilibrium. We can do this by setting each differential equation to zero. After some algebra, we find that the equilibrium values are:

$$\hat{P} = \frac{r}{a} \tag{3}$$

$$\hat{V} = \frac{m}{ab} \tag{4}$$

If we define ρ as the ratio of the average size of the predator population to the average size of the prey population ($\frac{\hat{P}}{\hat{V}}$), then decreases in ρ will correspond to increases in the relative size of the prey population. From equations (3) and (4) we can see that

$$\rho = \frac{rb}{m} \tag{5}$$

The next step is to consider how fishing affects the model populations. We can represent heavy fishing as corresponding to changes or perturbations of r and m. Specifically, heavy fishing decreases the prey growth rate (r) and increases the predator death rate (m). Inspecting ρ, the expression for the ratio of average densities, we can see that $\rho(\text{heavy fishing}) < \rho(\text{normal})$. (May, 2001; Roughgarden, 1979, 439) Since smaller values for ρ mean a larger relative size of the prey population, the population of prey will increase relative to the number of predators during times of heavy fishing.

These results led Volterra to his solution. His model predicted that heavy fishing favors the prey and light fishing favors the predator. Because WWI had slowed Adriatic fishing considerably, his model suggested that the shark population would be especially prosperous. This is not something that Volterra or anyone else would have expected a priori. However, armed with the dynamics of the model, he believed that he had a solution to this perplexing problem.

What remained was a careful analysis of the actual fisheries data to determine the adequacy of the model. Working with Volterra, D'Ancona examined records detailing the sale of fish between 1910 and 1924 in three of the upper Adriatic's markets. His conclusion was that Volterra's model made the correct qualitative predictions. The onset of WWI caused an initial increase in the number of prey (cod, squid, Norwegian lobster), but that this initial increase was followed by a decline in the abundance of both species and a relative increase in the number of predators (sharks).

Not everyone was convinced by Volterra's argument. Egon Pearson argued that D'Ancona had not correctly analyzed the fishery data, failing to take into account changes in fishing techniques. (Pearson, 1927) But Volterra and D'Ancona continued to defend the model, arguing that it captured the core factors which accounted for the fluctuating predator and prey populations in the Adriatic. (Kingsland, 1995)

Recognizing that his model was extremely simple and highly idealized with respect to any real world phenomenon, Volterra engaged in a lengthy study of predator-prey models of increasing complexity. (Volterra, 1926b) These models were initially studied as mathematical objects, but were subsequently deployed to understand more complex interactions in real populations. Volterra's work on population dynamics was one of the earliest applications of the modeling strategy in population biology and since then, modeling has been employed by nearly every population biologist involved in theorizing.[2]

2.2 Abstract Direct Representation

The story of Mendeleev's construction of the periodic system has a humble beginning. When assigned to teach courses on inorganic chemistry at the University of St. Petersburg, Mendeleev found that there was no good inorganic chemistry textbook available. Inorganic texts lacked an organized and coherent structure from which to characterize the known elements and inorganic reactions. In order to deepen his and his students' understanding of the elements, Mendeleev wanted to develop a classification system that elucidated their underlying properties. This would allow for a more systematic understanding of the properties of each element, the reactions each element could participate in, and trends underlying these properties.

Mendeleev faced a daunting theoretical challenge: samples of the pure elements had many chemically important properties, any of which might form the basis of a classification system. One might sort elements by color, density, conductivity, ductility, melting point, or by their affinity to react with various reagents. In the end, Mendeleev decided to focus his attention on finding trends in the properties of valency, isomorphism, and, most importantly, atomic weight, abstracting away from all of the other properties. Atomic weight is a familiar concept, but valency and what 19th century chemists called 'isomorphism' may not be. Elements are said to be *isomorphic* when families of salts containing chemically similar but distinct metals form similar crystal shapes. (Brock, 1992, 158) 'Valency' refers to the combining ratio of an element. For example, carbon is tetravalent, meaning that it can combine with four equivalents of hydrogen.

Mendeleev's first step was to organize the elements by atomic weight. This gave him a one dimensional ordering of the elements which served as an initial organizational device, but did not reveal any information about the elements' underlying structure or unity. Focusing next on valency and isomorphism, Mendeleev tried to find other dimensions along which to organize the elements. In modern terms, we can think of his next step as trying to figure out where each *period* or row on the Periodic Table ended.

[2]Other early modelers in population biology include Alfred Lotka and the major figures of the neo-Darwinian synthesis: R.A. Fisher, J.B.S. Haldane, and Sewall Wright.

In some accounts, Mendeleev is said to have put the names and properties of elements on cards and played "chemical solitaire" on long train journeys until he found a satisfactory ordering of the known elements. (Brock, 1992, 320)

In 1869, Mendeleev announced his ordering of the elements according to their weight and properties. This ordering, which later became know as the Periodic Table of the Elements, organized the elements in order of atomic weight and then in columns or groups in virtue of their chemical properties.[3] When the elements were properly ordered, Mendeleev argued, one could see the periodic dependence of elemental properties on their atomic weight. This principle, which Mendeleev called *The Periodic Law*, is one of the bedrock principles which organizes chemistry. It is still recognized as one of the most basic patterns among chemical phenomenon, although we now think of it as periodic dependence of elemental properties on atomic number or effective nuclear charge.

Mendeleev's theoretical achievements are sometimes overlooked because of the suspicion that the Periodic Table is *merely* a classification device. It makes certain trends explicit, but, it has been argued, the Table does not actually explain anything. The Library of Congress did a service to humanity by developing a relatively rational system for organizing books in our libraries, but surely we would not want to treat this as a theoretical achievement. Similarly, it has been argued that Mendeleev articulated an important classification system, but not a theory. For example, Shapere claimed that what Mendeleev discovered was an *ordered domain* and that "[o]rderings of domains are themselves suggestive of several different sorts of lines of further research" but not themselves theories. (Shapere, 1977, 534)

I believe this view to be mistaken for several reasons. The first reason involves the remarkable predictions that Mendeleev was able to make on the basis of his Periodic System. In 1869, he noted that there were gaps in his Table for three elements. On the basis of information about chemical trends encoded on the Table, he hypothesized the existence of what he called eka-aluminium, eka-silicon, and eka-boron. The properties of these novel elements are listed in table 1. Just a few years later, the elements gallium, scandium, and germanium were discovered and, as indicated on Table 1, their properties were in remarkable agreement with Mendeleev's predictions. (Scerri, 2001, 2006)

[3] In several instances, Mendeleev inverted the atomic weight ordering for the sake of chemical consistency. For example, by atomic weight alone, beryllium would have been in in the nitrogen group. However, its behavior is much more like that of magnesium, so Mendeleev placed it in the magnesium group.

	Predicted	Discovered
	Eka-aluminium (1871)	**Gallium (1875)**
Atomic Weight	68	69.9
Specific Gravity	6.0	5.96
Atomic Volume	11.5	11.7
	Eka-boron (1871)	**Scandium (1879)**
Atomic Weight	44	43.79
Specific Gravity	3.5	3.86
	Eka-silicon (1871)	**Germanium (1886)**
Atomic Weight	72	72.3
Specific Gravity	5.5	5.47
Boiling Point	<100 deg.C	86 deg.C
Density	1.9	1.89

Figure 1: Mendeleev's Predictions (after Brock, 1992)

Mendeleev's predictions might look like trivial exercises, making inferences about missing "books on the shelf" or filling empty slots. This underestimates the significance of the achievement: Mendeleev had no empirical knowledge that there were any empty slots to be filled. His task was thus not as simple as interpolating the properties of unknown elements on the basis of known elements. He first needed to hypothesize the existence of the missing elements by analyzing the theoretical structure he had created. Then he was able to use the trends posited by the Periodic Table to make predictions about the properties of the "missing" elements. This prediction was a theoretical, not merely classificatory, achievement.

While I believe that Mendeleev's remarkable predictions are one sign that he had developed an important theoretical structure, this view is not uncontroversial. Shapere argued that Mendeleev's predictions are not reasons to count the periodic system as a theory. He writes: "Nor does the fact that the ordering sometimes allows predictions to be made . . . turn such orderings into theories. (In particular, the periodic table is not 'explanatory' even though predictions can be made on its basis alone.)" (1977, 535) Even if we grant Shapere's argument that prediction alone was not enough to make Mendeleev a theorist, we should note that Mendeleev *did* give explanations on the basis of his periodic system.

One example of a trend Mendeleev explained using his system involves the oxides of the main group elements. For example, Mendeleev showed that the quantity of oxygen in the oxides was a periodic function of the element's group (column) on the Periodic Table. Group I elements (the alkali metals) formed oxides with structure R_2O, where R is a generic symbol of element. Group II elements (the alkaline earth metals) formed oxides as R_2O_2 and so on up to the halogens (R_2O_7). (Gordin, 2004, 31) This can be accounted for by the Periodic Law, but would have remained mysterious otherwise. Although accounting for this trend did not give Mendeleev causal or mechanistic knowledge about the formation of oxides, it certainly allowed him to make contrastive explanations about the reactivity of different metals. It also brought a number of things previously accepted as brute facts in to systematic unity. By the lights of many philosophers of science, these achievements count as explanatory ones and hence by Shapere's own standard, Mendeleev's system is a theory.

As Shapere further argued, Mendeleev's discovery of periodicity gives us a new fact that requires a further theoretical explanation. Periodicity is a phenomenon which is still not completely derivable from quantum mechanics (Scerri, 2004), although using semi-empirical methods, the trend can be derived (Levine, 1991). This is another reason that some, including Shapere, have questioned whether the periodic system is a theory or that Mendeleev made a theoretical, as opposed to a classificatory, contribution to chemistry. While it is true that Mendeleev's periodic system is in need of further theoretical explanation, the same could be said of any theory that is not a fundamental physical one. Everyone accepts classical thermodynamics as a theory, yet many would argue that core parts of it such as the Second Law themselves cry out for deeper theoretical explanation. Theories allow us to unify, make predictions, and frame explanations. It should not be required that they need no further explanation themselves. Mendeleev's system clearly unifies, allows us to make predictions, and can serve as the basic for chemical explanations. Thus it ought to be considered a theory, and Mendeleev considered a theorist.

Although Mendeleev is a theorist, his methods and style of theorizing were quite distinct from Volterra's. Mendeleev examined elemental properties, worked out which properties were essential and which ones could be abstracted away, and then constructed a representational system that elucidated important patterns and structure among the elements. This scientific activity constitutes theory construction, but not modelling. Mendeleev represented chemical phenomena *directly*, without the mediation of a model.[4] Although his theoretical descriptions of elemental properties and trends were abstract, they were descriptions of properties of the elements themselves.

Volterra's achievement was quite different. He engaged in *indirect* representation and analysis of predator-prey phenomena via the construction and analysis of a model. His equations described mathematical models of biological populations and these models were similar in certain respects to real biological systems, but the equations were not direct representations of any real system. It was only in virtue of the similarity between the models he had characterized and real populations of fish in the Adriatic that Volterra could answer D'Ancona's query. His characterization of the population dynamics of the Adriatic were made indirectly.

In comparing the cases of Mendeleev and Volterra, we can see that a central contrast between their theoretical styles involves their approach to representing and analyzing real world phenomena. Mendeleev created and studied a representation of real elemental properties, while Volterra created and studied representations of mathematical models that were similar to real phenomena. This difference is the essential contrast between modeling and ADR. To clarify this contrast we must take a closer look at the various stages of both modeling and ADR. As the construction and analysis of models are key steps of modeling, we must first consider some properties of scientific models.

3 Scientific Models

There are many characterizations of scientific models that have been offered in the philosophical literature. My own view of scientific models is discussed in (Weisberg, 2003), but as the focus of this paper is on the practice of modeling, much of what I will say is compatible with a wide range of accounts of models. In the first part of this section, I will give a description of the nature of models and their relationship to the world. This sketch is based on or compatible with accounts offered by Suppes (1960), Suppe (1977, 1989), Cartwright (1983), and Lloyd (1994). There are also some affinities between my sketch and the account offered by Giere (1988). The second half of the

[4] My claim here is that Mendeleev did not construct a model and analyze it. He worked directly with abstractions from data, focusing on the key factors that account for chemical behavior. According to some accounts of models, the Periodic Table itself might be considered a model. Even if that is true, it is irrelevant to my thesis because Mendeleev's intention was to represent trends in real chemical reactivity, not trends in a model system. These issues will be discussed in greater detail in §4.

section considers the role of the theorist's intentions in the evaluation of the model-world relationship. These features have been discussed in connection with models of experiments and models of data (Suppes, 1962) and the inferential conception of scientific representation (Suárez, 2004), but have not, for the most part, played a major role in discussions of scientific theories and theorizing.

Models are abstract or physical structures that can potentially represent real world phenomena. Many different things can serve as models including physically constructed scale models, model organisms, and mathematical objects such as sets of trajectories through a state space. For the purposes of this paper, I will restrict my attention to abstract, mathematical models, for they are the ones of primary importance in model-based theorizing.

When employing mathematical models, one studies the model by studying representations of the models, which I call *model descriptions*. For abstract models, model descriptions usually take the form of equations, but graphs and other kinds of representations can also serve as model descriptions. In the predator-prey model previously discussed, the two differential equations are the model description. In the model of molecular structure which I will discuss in §4.2, the model description takes the form of a computer program.

In the original semantic view (Suppes, 1960), scientific models are equated with logicians' models and are said to satisfy a set of axiomatic statements, what I have called the model description. In another account, model descriptions are taken to *define* models. (Giere, 1988, 83) Both of these accounts of the relationship between a model and its description are too strict because model descriptions often lack the precision to pick out a single model and this vagueness or partial specification is actually a benefit in the early stages of theorizing.[5] Rather than characterize the relationship between models and their descriptions as one of definition, let us characterize the relationship between the description and the model as one of *specification*. This highlights the fact that the relationship can be weaker than definition or satisfaction, but that models are picked out by their descriptions.

Some philosophical accounts collapse the distinction between models and their descriptions. For example, Orzack and Sober treat differential equations as models. (Orzack and Sober, 1993)[6]. This is a mistake for several

[5]An anonymous referee helpfully pointed out that in Suppes' discussion of the semantic view, this issue is partially addressed by giving set theoretical predicates for model *types* that will satisfy the axioms of the theory.

[6]It should be noted that, in conversation, Sober has said he believes that equations are like sentences and theoretical structures are like the propositions expressed by those

reasons. One of the most important insights behind the semantic view and other attempts to reconstruct theories as sets of models is that a theory should not depend on a particular linguistic formulation. More importantly for understanding the practice of modeling, a modeler often conceives of a model in a vague way, writes down some equations to describe the model she thought she had in mind, studies the model actually specified by the equations, and determines whether or not they pick out the right model. Situations can arise where the modeler's imagination picks out some set of models and her model description picks out a different set of models, necessitating a refinement either to her imagination or to her model description.

Modelers often use models in order to learn about real world phenomena. In these cases, the model must be *similar* to a real world phenomenon in certain appropriate respects. As Quine pointed out (1969), similarity is a vague notion and we therefore should not be content with such a simple formulation of the model–world relationship. One of the most active and contentious areas of the structure of theories literature concerns the question of how to give a more precise and detailed formulation of the model–world relation. Some philosophers, such as Giere (1988), argue that the appropriate relationship between models and the world is one of *structural similarity.* On his view, models are imaginary structures that would be concrete if real. Their similarity to real world phenomena lies in some parts of the imaginary structure literally having similar properties to parts of the real world phenomenon. Other theorists such as Suppes (1960), van Fraassen (1980), and Lloyd (1994) conceive of similarity more abstractly, describing it as a relationship between mathematical properties of the model and of the real world phenomenon described mathematically. A third view, related to the second, holds that models have *partial isomorphism* to their intended target phenomena via a series of models that are ultimately tied back to data. (da Costa and French, 2003, 49, 73) This means that some substructure of the model, for example the relations between its properties, stand in a one-to-one correspondence to properties of another model, which can be a model of the data.

None of these views has clearly emerged as dominant in the structure of theories literature; each has its critics and its supporters. For the main project of this paper—explicating the distinction between modeling and ADR—I can abstract away from many of the details of this debate. Naturally, a complete and final account of modeling will require this issue to be

sentences. This view brings him closer, but not all the way, to the perspective of the semantic view and the one that I am advocating.

settled.

Some accounts of models treat their relationship to the world as determinable simply by knowing the structure of the model and of the real world phenomenon being represented. For the purposes of understanding the practice of modeling, this view is too restrictive. Models do not have a single, automatically determinable relationship to the world. (Suppes, 1962; Weisberg, 2003; Suárez, 2004; Godfrey-Smith, forthcoming) Different modelers employing the same model may intend different parts of it to correspond with different parts of a real world phenomenon. Some modelers may require the model to faithfully represent the causal structure of the relevant phenomenon as well as make quantitatively accurate predictions. Others may only require that the model make accurate predictions. Still others may only require predictions in qualitative agreement with the properties of real world phenomena.

Volterra's model is a good example of these more nuanced properties of the modeler's intentions about the model-world relationship. Volterra believed that his model captured the essential causal relationships that gave rise to the unusual fishery data following the war. Modern ecologists think of Volterra's model as a *minimal model*: a template for building models of greater complexity. Thus if a modern ecologist deploys Volterra's model to study a real ecosystem, she does so with a much lower standard of fidelity. Her use of the model is only intended to give a first approximation to the most important dynamics of the system.

The relevant intentions of the modelers are included in what I will call the *construal* of the model.[7] The construal of a model is composed of four parts: an *assignment*, the modeler's intended *scope*, and two kinds of *fidelity criteria*. The assignment and scope determine and help us evaluate the relationship between parts of the model and parts of the real world phenomenon. The fidelity criteria are the standards theorists use to evaluate a model's ability to represent real phenomena.

The first aspect of a model's construal is its assignment, which is the specification of the phenomenon in the world to be studied and the explicit coordination of parts of the model with parts of the real-world phenomenon, itself described mathematically according to some accounts. This explicit coordination is important for two reasons. Firstly, although the parts of some models seem naturally to coordinate with parts of real-world phenomena, this is often not the case. For example, harmonic oscillator models were

[7] I have greatly benefitted from discussing modelers' intentions with Peter Godfrey-Smith, who coined the term 'construal' to refer to these intentions.

first developed to make predictions about the periodic motion of physical systems, but as mathematical models, they remain abstract objects without obvious analogs to the properties of springs, molecules, or even pendula. Chemists use harmonic oscillators to model vibrations in bonds. Thus they need to represent atomic positions as points in a coordinate system and treat the periodic offset of these points, which corresponds to molecular vibration, as the behavior described by the dynamics of the harmonic oscillator model.

Models typically have structure not present in the real-world phenomena they are being used to study. This brings us to the second role of the assignment: to specify which parts of the model are to be ignored. This aspect of the assignment is well illustrated by a practical, everyday model. In my kitchen I have an unusual kind of egg timer, a model egg. The model is a red, plastic egg-shaped object that has the words hard, medium, and soft printed on it. If you want to make a medium-boiled egg, you drop the plastic egg in simmering water along with your real egg. As the plastic egg heats up, it gradually changes from red to black, starting from the outside and working inwards. This mirrors what is happening inside your boiling egg — heat is slowly being transferred from the outside of the egg to the inside. When the egg timer has been sufficiently heated, the black color reaches the word 'medium' and you can remove the real egg from the heat. Apparently the egg timer works because the plastic out of which it is made has similar heat transference properties to a real egg. So the plastic egg provides a high fidelity model of heat diffusion, the most important property associated with getting a well cooked egg.

Although this is not a mathematical model and not a scientifically interesting case, the plastic egg model demonstrates this second important role of the model's construal. Many aspects of the model egg are irrelevant for the purpose to which it is being used — we do not care that the model egg is red, made of plastic, a bit lighter than a real egg, completely homogeneous, and printed with words. We do care that the model egg has approximately the same heat diffusion properties as the real egg and that the color change on the model egg represents heat-induced coagulation on the real egg. The construal tells us which parts of the model correspond to parts of the real phenomenon and which parts can be ignored.

No one would ever assume that the real egg is made of plastic or is red; however, this issue often arises in a more subtle way when one considers mathematical models. Volterra's predator-prey model is described by two coupled differential equations. These equations and the mathematical model they describe are defined for any real valued number of predators and prey. However, Volterra certainly did not intend the fact that his model could describe the dynamics of non-integer numbers of predator and prey to correspond to any real or possible population of fish. Thus in his construal of the model, Volterra only assigned the integer values of the model (and probably only certain ranges of these integer values) to the population in the Adriatic and other possible populations.

The second component of a model's construal is the model's intended scope, which tells us the aspects of phenomena intended to be represented by the model. (Suppe, 1977, 223) This is related, but somewhat different from the assignment, which tells us about how to coordinate particular parts of the model to particular parts of the target phenomenon and which parts of the model should not be taken to represent anything.

Scope is best illustrated by example, so let us turn once again to Volterra's predator-prey model. The model itself only describes the size of the predator and the prey population, the natural birth and death rates for these species, the prey capture rate, and the number of prey captures required to produce the birth of a predator. (Roughgarden, 1979, 267) It contains no information about spatial relations, density dependence, climate and microclimate, or interactions with other species. If the scope is such that we intended to represent those features, Volterra's model does a poor job because it would indicate that there is no density dependence, no relevant spatial structure, etc. By choosing a very restrictive scope, we indicate that Volterra's model is not intended to represent these features.

The third and fourth aspects of a model's construal are its fidelity criteria. While the assignment and scope describe how the real world phenomenon is intended to be represented with the model, fidelity criteria describe how similar the model must be to the world in order to be considered an adequate representation. There are two types of fidelity criteria: *dynamical fidelity criteria* and *representational fidelity criteria*.

Dynamical fidelity criteria tell us how close the output of the model must be to the output of the real world phenomenon. It is often specified as an error tolerance. For example, a dynamical fidelity criterion for a predator–prey model might state that the population size of the predators and prey in the model must be $\pm 10\%$ of the actual values before we will accept the model.

Dynamical fidelity criteria only deal with the output of the model, its predictions about how a real world phenomenon will behave. Representational fidelity criteria are more complex and give us standards for evaluating whether the model makes the right predictions for the right reasons. These criteria usually specify how closely the model's internal structure must match the causal structure of the real world phenomenon to be considered an adequate representation.

In summary, an account of models adequate for characterizing the practice of modeling must have the following characteristics:

1. Models can be physical or abstract.

2. Descriptions of models should be distinguished from models themselves.

3. Models and model descriptions have a many-many relationship.

4. Different accounts of the relationship between models and the world are possible, but they can all loosely be described as relations of similarity, as opposed to relations like truth or reference.

5. The model–world relationship is partially determined by the construal, which depends on the intentions of the model user.

6. The construal along with the world determines whether or not any real phenomenon is represented by the model.

4 Distinguishing Modeling from ADR

Returning to the question of what distinguishes modeling from other kinds of theorizing, we need to consider more carefully the nature of theoretical representation and what role the model is playing in Volterra's analysis of the Adriatic. Continuing the discussion from §2.1, I will argue that the most distinctive characteristic of the practice of modeling is the analysis of a real world phenomenon by first analyzing a model. Hence the process of representing a model becomes the first step in representing a real world phenomenon. This is distinct from ADR because, in ADR, there is no intermediary stage. Theorists engaged in ADR analyze and represent the properties of a real world phenomenon, suitably abstracted, in the first instance. Their aim is to represent a real phenomenon and analyze their representation of this phenomenon.

Even this preliminary discussion of the difference between modeling and ADR reveals a central feature of these practices. The practices are distinguished by the actions and intentions of theorists, not by the outcome of the process of theorizing. This means that to judge whether or not a particular theorist is a modeler, it will not be sufficient to determine whether or not her theory can be represented as a model or cluster of models. We will actually need to know something about how the theory was developed and how the modeler set about trying to represent the world.

4.1 The First and Second Stages of Modeling

Volterra began his investigation of Adriatic fish not by looking directly at these fish or even the statistics gathered from the fish markets, but by constructing a model. This is characteristic of the first stage of modeling. He imagined a population of predators and a population of prey, each with only two properties. Setting this idea to paper, he wrote down equations specifying the model that he had imagined. Theorists do not often record the details of this process, so we do not know how satisfied Volterra was with the initial model. Perhaps it did not match the model he had imagined and so he refined the model. Or perhaps he had correctly specified the model he was imagining and was able to proceed to analyze the model. In either case, the first stage of Volterra's analysis involved him constructing something. He did not start by looking for patterns in data; he began by constructing a model.[8]

Once Volterra had constructed his mathematical model, he could go on to perform an analysis of it. This is the second stage of modeling. As I discussed in §2.1, Volterra's analysis involved studying the effect of a perturbation on the ratio of the average abundance of predators and prey. The results of Volterra's analysis proved useful for determining the behavior of the Adriatic populations of fish, but it is important to see how these first two stages of description and analysis are distinct from any application to the real population. Although guided by what he knew was happening in the Adriatic, the second stage of modeling involved finding out some very general properties of predator-prey models, ones that apply far more widely than the particular case he had in mind.

For example, Volterra's discovery about the effect of fishing on the predator/prey balance was merely an instance of a more general principle he discovered in the analysis of his model. He discovered what is now called the *Volterra Principle*: general pesticides (some intervention that kills both predator and prey) increase the relative proportion of the prey. (Roughgarden, 1979) This principle was first discovered using the model described in his original predator-prey model, but Volterra and subsequent theorists

[8]In less path-breaking investigations, modelers often use "off-the-shelf" models, structures that have already been applied to other phenomena of interest. In such cases, the first stage of modeling involves identifying the appropriate model, rather than explicitly constructing it.

showed that it emerges from many other predator-prey models as well. (Weisberg, forthcoming) This discovery is an example of how analyzing a model and determining what is true of a model can take place autonomously from any particular system being studied. Even where the model is inspired by a real world system, what the theorist finds out about it is distinct and usually more general from the system which inspired it.

There are other types of cases that are even more striking examples of the autonomy of the first and second stages of modeling. Some modelers construct models simply to explore or illuminate a hypothesis. For example, population biologists often examine very simple models of sexual and asexual reproduction in order to better understand the evolution of sex. (Roughgarden, 1997, x) These models are not intended to describe any actual organism; they are far too simple for that. Their importance lies in helping us to understand very general facts about the differences between sexual and asexual reproductive systems. In this type of case, the model itself is clearly the object of study.

It is also possible to study a model of a phenomenon that is known not to exist. A. S. Eddington once wrote: "We need scarcely add that the contemplation in natural science of a wider domain than the actual leads to a far better understanding of the actual." (1929, 266) Agreeing with Eddington, R. A. Fisher explained that the only way to understand why there are always two sexes involved in sexual reproduction is to construct a model of a three-sexed sexually reproducing population of organisms. (Fisher, 1930) Constructing a model of such a phenomenon is the only way to study it because, by stipulation, the phenomenon does not exist. Modelers are often interested in phenomena such as three-sex biology, perpetual motion machines, or non-aromatic cyclohexatriene because, insofar as we can understand why these phenomena do not exist, we will have gained a better of understanding of phenomena that do exist. Again in this case, it is clear that the model and only the model is the object of study.

As these cases make clear, modeling can terminate after the second stage and involve no assessment of the model-world relationship. When it does terminate here, any analytical results gained in the analysis apply only to the model, not to the world.[9] But even when modeling does not terminate after the second stage, the first two stages involve the study of a model as an autonomous object. This is true even when these steps are carried out with a particular real-world phenomenon in mind, such as in Volterra's analysis

[9]Insights gained from the analysis, however, may be useful in understanding real phenomena.

of the predator-prey model.

4.2 Third Stage of Modeling

Modelers sometimes proceed no further than the second stage of modeling, but much of the time theorists construct and analyze models in order to study real world phenomena. This requires a third stage in which the theorist attempts to coordinate a model with a real world phenomenon. As I discussed in §3, the model-world relationship is still the subject of much debate in the literature and a detailed analysis of the third stage of modeling will depend crucially on how this debate is resolved. Thus I will only describe the third stage in outline form, which I believe to be more-or-less compatible with all of the major accounts of the model-world relation currently under discussion.

Comparing a model to a real world phenomenon involves first preparing the real world phenomenon using what Suppes calls a model of the data and what I have called a parameterized target system (Weisberg, 2003). Essentially, the theorist has to subject the phenomenon in the real world to a process of abstraction, deciding which aspects of the phenomenon will actually be considered. The result of this process of abstraction has variously been called a physical phenomenon (Suppe, 1977, 224) or a target system. The next step is for the modeler to represent the target system mathematically, assigning variables and parameters to the properties of the target system.[10]

How to articulate the next step of the process depends on one's account of the nature of model-world relationship. In the state-space version of the semantic view, the next step involves comparing the states of the target system to the trajectories associated with the model. In van Fraassen's version (1980), the curve which best fits the data must be isomorphic to one of the model's trajectories. Given the unrealistic demands of this requirement, several authors have weekend the requirement considerably. (e.g., Lloyd, 1994; da Costa and French, 2003) In light of the discussion in §3, I think the relationship should be made even more flexible. The assessment of the adequacy of a model can only be made with respect to the fidelity criteria a theorist has chosen. Sometimes this will involve something as strong as isomorphism, but other times only qualitative agreement between measurements and the model will be required. This is settled by the interests of the

[10] Although this is an accurate description of theoretical practice, it is not required by the accounts of the model-world relationship offered by Giere (1988) and da Costa & French (2003).

theorist, not universal criteria.

I have already discussed Volterra and D'Ancona's evaluation of their model using fishery data. So let's consider another example of modeling, paying particular attention to the third stage. Molecular models are often used in the course of explaining three-dimensional molecular structure. Say a chemist is interested in explaining the major conformational change caused by a tertiary butyl substitution in cyclohexane. She might employ a classical molecular model in this analysis. In such a model, bonds are treated as rotating harmonic oscillators and atoms as the masses being oscillated. The model further builds in information about the spatial position of the atoms and the potential energy associated with these atoms at different distances from one another. For a very simple model such as this one, the tertiary butyl substitution would be simply treated as a change in mass and in the amount of space filled by the substituent.

Using a set of techniques collectively referred to as *molecular mechanics*, the theorist can calculate the minimum energy conformation of substituted and non-substituted cyclohexane models. She can then analyze the model, examining the factors which predict and explain structural changes such as torsional strain, van der Waals strain, or steric hindrance. For this particular substitution, the conformational analysis of molecular models suggests that the drastic change in conformation is a result of a steric interaction between the tertiary butyl group and the axial hydrogen which is *syn* to it.[11]

So far the chemist is still at the second stage of her analysis, having accounted for the properties of the model system, but not the real system. However, if the model is similar to the real molecule both with respect to aspects of its structure and the effect that substitution has on the model's conformation, then the chemist can move to the third stage of her analysis. If the model molecule is similar to the real molecule, then she can indirectly account for the conformational change as being due to the interaction between the t-butyl group and the *syn*-axial hydrogen in the real molecule. The chemist has indirectly represented and analyzed the structure and properties of the real molecule by describing and analyzing a model.

The chemist's description of her model can serve as a description of the real phenomenon because of the phenomenon's similarity to the model. Had this theorist engaged in ADR, she would have attempted to describe the actual phenomenon without the mediation of the model. Therein lies the key difference between modeling and ADR.

[11] Molecular mechanics and other techniques of conformational analysis are reviewed in Carroll (1998), Carey and Sundberg (2000), and Lowry and Richardson (1997).

In §2.2, I argued that Mendeleev was not a modeler, but it is instructive to think about the counterfactual history of science in which he was one. How would Mendeleev-as-modeler construct the Periodic System? The first step would have involved creating some simplified system of the elements, perhaps a system in which the elements really only have a few properties and some kind of structured dependence on one another. He would then have proceeded to write down a description of this model and analyze its logical consequences: Does it say anything about trends in reactivity? About the color of salts of the elements? About elemental properties such as valence and isomorphism? After analyzing such a model, Mendeleev would then turn to the question of the real elements asking: Do the real elements behave, largely, like my model behaves? If the answer was yes, then Mendeleev's representation of his model as well as the analysis he performed on the model would be, indirectly, a representation of the real elements and their properties.

This was not the historical Mendeleev. There was no point at which Mendeleev constructed a model, studied this model as an independent object, discovered the phenomenon of periodicity in the model, and then coordinated the model to the real world phenomenon of periodicity. He analyzed a direct representation of periodicity, not periodicity in a constructed model. However, imagining Mendeleev as a modeler helps us to see the important differences between modeling and the kind of theorizing actually practiced by Mendeelev. To this alternative form of theorizing, what I have called ADR, we now turn.

4.3 ADR

Like modeling, ADR takes place in multiple stages. In the first stage, a theorist subjects the phenomenon of interest to a process of abstraction, making decisions about which properties to focus on and which to consign to a less important role. She then constructs a representation of the relevant properties and relationships between those properties which can be determined. This can take the form of equations, graphs, pictures, etc.: the same sorts of things that can be used to describe models. The difference here is that the representation is being used to describe the real world phenomenon, not a model. This procedure is analogous to the beginning of the *third* stage of modeling. In both cases, the phenomenon of interest is subject to a process of abstraction, and parameterization follows yielding what I call a parameterized target system or model of data[12].

ADR's second stage is much like modeling's second stage. The theorist engages in an analysis of her representation, which in this case is a representation of the real world phenomenon. Because the theorist is analyzing a representation that is directly related to a real phenomenon, anything she discovers in her analysis of the representation is a discovery about the phenomenon itself, assuming that it was represented properly. There is no extra stage where the theorist must coordinate the model to a real phenomenon.

Take Mendeleev's construction of the Periodic System. His construction began with a process of abstraction, focusing on the elemental properties he considered to be the most relevant. He went on to build up a theoretical representation of properties of the elements themselves. There were no intermediate steps involving the construction or analysis of a model, which had independent existence. Mendeleev's intention was to represent real elements and their properties and then to use this representation to make novel predictions and explain reactivity and structure with it.

Darwin's geological work provides another good example of theorizing without the use of models. Struck by the many atolls, coral islands with a ring-like configuration, in the Pacific Ocean, Darwin set out to construct a theory explaining their origin and distribution, and more generally the origin and distribution of coral reefs of all types. His theory of atoll formation is relatively simple, although it was constructed after painstaking observation of the structure and distribution of the three major types of coral formations.

Darwin's atoll theory involved several components. He theorized that the first stage involved the formation of a fringing reef around the perimeter of an existing volcanic island, creating a submerged ring of coral growth around the island. If the island were to start sinking, the coral would continue growing upwards, for coral can only grow in relatively shallow water. If the existing volcanic island ultimately receded below the surface of the water and the coral continued growing upwards, we would be left with a ring-shaped coral island containing a central lagoon, where the volcanic island had previously been located. (Darwin, 1842/1984; Ghiselin, 1969)

Darwin's theory seems quite simple, but coupled to other simple geological theories, it can be used to make specific predictions about the distribution of different types of coral formation in the Pacific. For example, Darwin predicted that in a given area, the type of reef formation would be fairly uniform. Atolls can only form in specific circumstances — where there was

[12]Like Suppes, I am distinguishing here between a theoretical model, which I just call a 'model', and a model of data

volcanic activity leading to volcanic islands and where there was ultimately geological subsidence, resulting in the islands receding below the surface of the water while the coral continued to grow upwards. Since Darwin had made numerous observations about the geology of the South Pacific, he was able to make explicit predictions about where atolls would be found. The final chapter of *The Structure and Distribution of Coral Reefs* contains an extensive analysis of the distribution of reefs in the Pacific and explains how these observations are predicted on the basis of his theory.

Nowhere in Darwin's work on coral reefs does he construct and analyze a model. He intends his theoretical representations to apply, in the first instance, to the data he has collected in the Pacific. He organizes this data, considers other geological information that bears on it, and constructs a theory to explain this data. He then goes on to make further predictions about the distribution of reefs and uses many different data sources to confirm his theory. But at all times, Darwin was talking about the actual atolls in the Pacific. There was no analysis of a constructed model.

Thus Mendeleev and Darwin approached theorizing very differently than Volterra; their representations and analyses were directly aimed at real world phenomena. This does not mean that the final product they produced — Mendeleev's Table and Periodic Law and Darwin's theory of atoll formation — cannot be reconstructed using models as the semantic theorists would urge us to do. It may be possible to recast Mendeleev's and Darwin's theoretical representations into model-based representations. Similarly, it may be possible to take the equations that describe Volterra's model and treat them as approximate, direct representations of Adriatic predator and prey populations. That these transformations may be possible should not change our analysis of their theoretical practice. The contrast between modeling and ADR is about the practice, not the products of theorizing.

5 Who is Not a Modeler?

I want to close this paper by discussing three tempting, but misleading ways to characterize modeling. These ideas are often found in the scientific literature, especially in textbooks, and some philosophers have also discussed modeling in these ways. Models are sometimes equated with approximately accurate representations; hence theorists who make abstractions, who approximate, and who idealize are called modelers. This designation cuts across many of the distinctions that I think are important for characterizing theoretical practice and I believe it should be avoided.

The first mistake is to equate modeling with abstraction. All theoretical representation involves abstraction, the process of systematically ignoring aspects of the phenomenon of interest. This is just the difference between theorizing and simply giving a report about the raw data. Almost every scientific paper that is published involves some form of abstraction, since this allows raw data to be examined for patterns. Take Mendeleev as an example. Although Mendeleev was not a modeler, he relied heavily on abstraction. He synthesized the data from studies of many samples of the elements, then abstracted away almost all of the properties of the elements, leaving only atomic weight and a few others. This process of abstraction is what allowed him to determine the correct ordering of the elements on the Periodic Table. I have very deliberately chosen to call Mendeleev's theoretical style *abstract* direct representation.

Another common mistake is to associate any kind of approximation with modeling. Facing a theoretical representation which is extremely approximate, theorists might warn us to not take it too seriously because "it's just a model!" or "we were just modeling!" Translating these expressions into the terms of this paper, theorists are warning us that their representations (equations, graphs, etc.) are only accurate descriptions of model systems, not real systems. While this sort of warning is often made in connection with models displaying a low degree of fidelity to their intended target phenomena, there is no unique connection between approximation and modeling. In fact, both modelers and ADRs may rely heavily on approximation.

Darwin's construction of the atoll formation theory is an example of ADR, yet it involved certain kinds of approximations. For example, Darwin treats large regions of the underwater geology as uniform, even though he knew that there was local heterogeneity. So we should not take the existence of either abstraction or approximation to be a sign of modeling. Theorists of different varieties may be required to approximate and are always required to abstract.

Finally, some scientists and philosophers have taken idealization to be a defining characteristic of modeling. This issue is much more subtle then the others and there may well be a deep connection between idealization and modeling, despite the lack of consensus in the literature about exactly what idealization is. In my view, idealization involves more than approximation because it involves committing oneself to representational ideals which fall short of completeness. In the language of this paper, it involves lowering one's standard's of fidelity explicitly, perhaps because such an action can promote some other desirable theoretical virtue.

To delve deeply into this issue takes us away from the main focus of this paper, so I will simply make one suggestion: Theorists who practice ADR typically aim to give complete representations. Although they know they will fall short of such an ambitious goal, this is their goal nevertheless. But this is typically not the case with modelers. Because the world is one step removed from the theorist's representation and analysis, she has more latitude to explore non-actual possibilities. This involves more than simply approximating as a matter of course. It involves committing oneself to giving non-complete representations by lowering one's standard's of fidelity. Concluding that all idealization involves modeling is probably a mistake, but there is a deep connection between modeling and idealization.

6 Conclusion: Who is a Modeler?

When one focuses on the structure of mature theories, it is easy to assume that there is just one type or strategy of theorizing. While I have not said much about the structure of theories in this paper, I have argued that the first assumption is a mistake: modeling and ADR are distinct kinds of theorizing. A complete philosophical understanding of modern theoretical practice must recognize diversity in theory construction.

Modeling is distinguished from ADR by a theorist's construction and analysis of a model, which is used to analyze and represent a real world phenomenon indirectly if at all. When a modeler wants to describe a real phenomenon, she begins by choosing a model, not a real phenomenon to analyze. The description and analysis of this model can be done mathematically, pictorially, or even verbally. If the model is found to be appropriately similar to the real world phenomenon of interest, then the modeler's representation and analysis of the model is also an indirect representation of the real world phenomenon. However, sometimes a model is analyzed to answer very general questions that do not ask about any specific real world phenomenon or ask about a phenomenon that is known not to exist.

Although modeling and ADR may not uniquely divide up the domain of theoretical practice, they are two of the most important kinds of theorizing. Future accounts of theories and theorizing should carefully distinguish between them and further elucidate their properties. Only after we have a better understanding of the diversity inherent in theoretical practice will we be in a position to answer some of the most philosophically challenging issues about theorizing, including when modeling is an advisable strategy, what kinds of models to construct, and how to go about choosing the most fruitful models for a given scientific problem.

Michael Weisberg
Department of Philosophy
University of Pennsylvania
433 Logan Hall
Philadelphia, PA 19104
USA
weisberg@phil.upenn.edu

References

Brock, W. H. [1992]: *The Norton History of Chemistry*, New York: W.W. Norton.

Carey, F. A. and R. J. Sundberg [2000]: *Advanced Organic Chemistry: Structure and Mechanisms* (4th ed.), New York: Plenum.

Carroll, F. A. [1998]: *Perspectives on Structure and Mechanism in Organic Chemistry*, Pacific Grove: Brooks/Cole.

Cartwright, N. [1983]: *How the Laws of Physics Lie*, Oxford: Oxford University Press.

Cartwright, N. [1989]: *Nature's Capacities and Their Measurement*, Oxford: Oxford University Press.

da Costa, N. C. A. and S. French [2003]: *Science and Partial Truth*, New York: Oxford University Press.

Darwin, C. [1842/1984]: *The Structure and Distribution of Coral Reefs*, Tucson: University of Arizona Press.

Domotor, Z. [2001]: 'Mathematical models in philosophy of science,' in N. J. Semelser and P. B. Baltes (*eds*), 2001, *International Encyclopedia of the Social and Behavioral Sciences*, Amsterdam: Elsevier Science.

Downes, S. M. [1992]: 'The importance of models in theorizing: A deflationary semantic view,' *PSA 1992*.

Eddington, A. S. [1929]: *The Nature of the Physical World*, New York: Macmillan.

Fisher, R. A. [1930]: *The Genetical Theory of Natural Selection*, Oxford: The Clarendon Press.

Ghiselin, M. T. [1969]: *The Triumph of the Darwinian Method*, Berkeley: University of California Press.

Giere, R. N. [1988]: *Explaining Science: A Cognitive Approach*, Chicago: University of Chicago Press.

Godfrey-Smith, P. [forthcoming]: 'The Strategy of Model Based Science,' *Biology and Philosophy* , forthcoming.

Gordin, M. D. [2004]: *A Well-ordered Thing*, New York: Basic Books.

Hesse, M. [1966]: *Models and Analogies in Science*, Notre Dame: University of Notre Dame Press.

Hofbauer, J. and K. Sigmund [1998]: *Evolutionary Games and Population Dynamics*, Cambridge: Cambridge University Press.

Kingsland, S. E. [1995]: *Modeling nature: Episodes in the History of Population Ecology* (Second ed.), Chicago: University of Chicago Press.

Levine, I. N. [1991]: *Quantum Chemistry* (4th ed.), Englewood Cliffs, NJ: Prentice Hall.

Lloyd, E. A. [1994]: *The Structure and Confirmation of Evolutionary Theory* (Second ed.), Princeton: Princeton University Press.

Lowry, T. H. and K. S. Richardson [1997]: *Mechanism and Theory in Organic Chemistry* (third ed.), New York: Harper & Row.

May, R. M. [2001]: *Stability and Complexity in Model Ecosystems* (Landmarks in Biology ed.), Princeton: Princeton University Press.

Morgan, M. S. and M. Morrison [1999]: 'Models as mediating instruments,' in M. S. Morgan and M. Morrison (*eds*), *Models as Mediators*, pp. 10–37, Cambridge: Cambridge University Press.

Odenbaugh, J. [forthcoming]: 'Models,' in *Blackwell Companion to Philosophy of Biology.*

Orzack, S. H. and E. Sober [1993]: 'A critical assessment of levins's 'The strategy of model building in population biology (1966),'' *Quarterly Review of Biology* **68** (4), 533–546.

Pearson, E. S. [1927]: 'The application of the theory of differential equations to the solution of problems connected with the interdependence of species,' *Biometrika* **19**, 216–222.

Quine, W. [1969]: 'Natural kinds,' in *Ontological Relativity and Other Essays.* New York: Columbia University Press.

Roughgarden, J. [1979]: *Theory of Population Genetics and Evolutionary Ecology: An Introduction*, New York: Macmillan Publishing Co.

Roughgarden, J. [1997]: *Primer of Ecological Theory*, Upper Saddle River, NJ: Prentice Hall.

Scerri, E. [2004]: 'How ab initio is ab initio quantum chemistry?' *Foundations of Chemistry*, **6**, 93–116.

Scerri, E. R. [2001]: Prediction and the periodic table, *Studies in the History and Philosophy of Science* **32** (3), 407–452.

Scerri, E. R. [2006]: *The Periodic Table: Its Story and Significance*, New York: Oxford University Press.

Shapere, D. [1977]: 'Scientific theories and their domains,' in F. Suppe (*ed*), *The Structure of Scientific Theories* (second ed.), pp. 518–565, Chicago: University of Illinois Press.

Suárez, M. [2004]: 'An inferential conception of scientific representation,' *Philosophy of Science* **71**, 767–779.

Suppe, F. [1977]. 'The search for philosophic understanding of scientific theories,' In F. Suppe (*ed*), *The Structure of Scientific Theories* (2nd ed.), Chicago: University of Illinois Press.

Suppe, F. [1989]: *The semantic conception of theories and scientific realism*, Chicago: University of Illinois Press.

Suppes, P. [1960]: 'A comparison of the meaning and use of models in mathematics and the empirical sciences,' *Synthese* **12**, 287–300.

Suppes, P. [1962]: 'Models of Data', in E. Nagel, P. Suppes, and A. Tarski (*eds*), *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress*, Stanford: Stanford University Press, pp. 251–61.

van Fraassen, B. C. [1980]: *The Scientific Image*, New York: Oxford University Press.

van Fraassen, B. C. [1991]: *Quantum Mechanics*, Oxford: Oxford University Press.

Volterra, V. [1926a]: 'Fluctuations in the abundance of a species considered mathematically,' *Nature* **118**, pp. 558–60.

Volterra, V. [1926b]: 'Variazioni e fluttuazioni del numero d'individui in specie animali conviventi,' *Memorie Della R. Accademia Nazionale Dei Lincei* **II**, pp. 5–112.

Weisberg, M. [forthcoming]: 'Robustness Analysis,' *Philosophy of Science*, forthcoming.

Weisberg, M. [2003]: *When Less is More: Tradeoffs and Idealization in Model Building*, Ph. D. thesis, Stanford University.

Winther, R. [forthcoming]: 'Parts and Theories in Compositional Biology,' *Biology and Philosophy*, forthcoming.

Wimsatt, W. C. [1987]: 'False models as a means to truer theories,' in M. Nitecki and A. Hoffmann (*eds*), *Neutral Models in Biology*, New York: Oxford University Press, pp. 23–55.