Models for Modeling[*]

Michael Weisberg
University of Pennsylvania

Draft of January, 2008 (4e)
Under review. Please do not cite without permission.

1 Two Aquatic Puzzles

In the late 1940s, the San Francisco Bay Area was feeling the growth pains of rapid development. The need to find new land, to create additional traffic crossings, and to provide fresh water engendered a number of plans for modifying the Bay to accommodate the rapid growth. Perhaps the most notorious of these plans was proposed by John Reber, a former school teacher and theatrical producer.

Reber's plan was to build two salt-water barriers in the Bay. These barriers would consist of earth and rock dams, and their tops would be sealed so that they could hold new highways and rail tracks. Reber claimed that his plan would produce 20,000 acres of filled land, increase the deep-water harbor by more than 50

[*]I am greatful to Brett Calcott, Zoltan Domotor, Patrick Forber, Paul Humphreys, Ben Jeffares, Christian List, Elisabeth Lloyd, John Matthewson, Ryan Muldoon, Kim Sterelny, Scott Weinstein, and especially Peter Godfrey-Smith, Daniel Singer, and Deena Skolnick Weisberg, for their comments, challenges, and advice. Early drafts of this paper were presented at The Australian National University and I would like to thank the Research School of Social Sciences for its gracious hospitality. This work was partially supported by NSF grant SES-0620887.

miles, and conserve almost 2.5 million acre-feet of fresh water. The plan's critics worried that it would destroy commercial fisheries, render the South Bay a brackish cesspool, and create problems for the ports of Oakland, Stockton, and Sacramento (Jackson & Paterson, 1977).

Among the opponents of the plan was the Army Corps of Engineers, who were charged with studying the environmental impacts of the various salt water barrier plans. The Corps recognized the benefits that the Reber plan might bring to the area, but were quite certain that adding salt water barriers to the existing Bay ecosystem would have serious, unintended consequences. In the meantime, Reber garnered much public support and assured the citizens of the Bay Area that his plan could only bring benefits. The Corps recognized that a battle of words would not be helpful in advising regional authorities; they needed hard data. But how could such data be collected without actually building the salt water barriers and risking harm to the bay?

To better study the plan the Corps built a hydraulic scale model of the San Francisco Bay. This wasn't any ordinary scale model. It was the "San Francisco Bay in a Warehouse," (Huggins & Schultz, 1967) an immense structure housed in a Sausalito warehouse that started out at a size of about 1 acre (340 ft long in the north/south direction) and has grown to be about 1.5 acres today. Hydraulic pumps simulated the action of tidal and river flows in the Bay, modeling tides, currents, and the salinity barrier where fresh and salt water meet (Huggins & Schultz, 1973).

Armed with this high-fidelity model of the San Francisco Bay, the Army Corps of Engineers was now in a position to evaluate the Reber plan. Reber's proposal called for a 600 foot wide, 4 mile long earth and rock dam that stretched from San Quentin to Richmond, and a second barrier 2000 feet wide and 4 miles long, just south of the Bay Bridge, connecting San Francisco to Oakland. The Army Corps studied this proposal by building scaled-versions of Reber's barriers, adding these to the Bay model, and measuring the changes in current, salinity, and tidal cycles.

As the Corps predicted, the model showed that Reber's plan would have disastrous consequences for the Bay and its ecosystem. The proposed barriers did not allow for any fresh-water flushing to take place behind them. Far from creating the fresh water lakes envisioned by Reber, the barriers would actually create stagnant pools, without the circulation required to maintain healthy aquatic ecosystems. The possibility of introducing openings in the barriers to allow for water flow was also explored. This would not, of course, generate fresh water lakes, but it was thought that this might allow the dams to be built without creating stagnant water behind them. The model showed that building such porous barriers was also a very bad idea because doing so introduced extremely high velocity currents in the Bay. This would disrupt the normal ecosystem and make navigation in the central part of the Bay hazardous. The Corps was thus in a strong position to denounce Reber's plan on the basis of model-derived data (United States Army Corps of Engineers, San Francisco District, 1963).

A very different sort of model was built to address another aquatic puzzle. After World War I, there was an unusual shortage of aquatic life in the Adriatic sea. This seemed especially strange because fishing had slowed considerably during the war. Most Italians believed that this should have given the natural populations 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 an 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 the well known mathematician and physicist Vito Volterra, who approached the problem not by studying the fishery statistics directly and not by building a physical model, but by constructing a mathematical model composed of one population of predators and one population of prey (Volterra, 1926).

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). The result was what we now know as the Lotka-Volterra model of predation, which is described by the following two differential equations:

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

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

By analyzing the models described by these equations, Volterra solved the puzzle of the fishery shortages: His model predicted that intense levels of a general biocide, which kills both predators and prey at the same time, would be relatively favorable to the prey, whereas lesser degrees of biocide favor the predators. From this he reasoned that heavy fishing, a general biocide, favors the prey and light fishing favors the predator. Because WWI had slowed Adriatic fishing to a trickle, his model suggested that the shark population would be especially prosperous during this time of reduced biocide. This is not something that Volterra or anyone else would have expected a priori. However, armed with the dynamics of his mathematical model, Volterra found a solution to this perplexing problem.

These two episodes are paradigm cases of scientists solving problems by *modeling*, the indirect study of real-world systems by the construction and analysis of models. Contemporary literature in philosophy of science has begun to emphasize the *practice* of modeling, which differs in important respects from other forms of representation and analysis central to standard philosophical accounts (e.g. Godfrey-Smith, 2006; Wimsatt, 2007; M. Weisberg, 2007). This literature has stressed the constructed nature of models (Giere, 1988), their autonomy (Morgan & Morrison, 1999), and the utility of their high degrees of idealization (Levins, 1966; Wimsatt, 1987; Batterman, 2001; Hartmann, 1998; Strevens, 2006; M. Weisberg, forthcoming). What this new literature about modeling lacks, however, is a comprehensive account of the models that figure in to the practice of modeling. Instead, most authors have borrowed accounts of models from the *semantic view of theories* literature, a family of views which argues that the best analysis of scientific theories should be conducted in terms of models (Suppes, 1960a; Suppe, 1977).

The semantic view literature provides many starting points and insights from which we can build an account of models appropriate to understand the practice of modeling. Yet the extant literature is not fully adequate for this purpose because it is focused on models that figure in the foundations of mature theories, not the process of model construction and analysis. In particular, and most importantly for the current purposes, this literature has not fully explored the role of theorists' intentions in all aspects of modeling, including the individuation of models, the coordination of models to real-world systems, and the evaluation of the goodness of fit between models and the world.

This paper thus offers a new account of both concrete and mathematical models, with special emphasis on the intentions of theorists, which are necessary for evaluating the model-world relationship during the practice of modeling. Although mathematical models form the basis of most of contemporary modeling, my discussion begins with more traditional, concrete models such as the San Francisco Bay model. By examining some of the key features of concrete models and their interpretations, we learn much that can be used in the more complex analysis of mathematical models.

2 Anatomy of a Concrete Model

Concrete models are real, physical objects that are intended to stand in a resemblance relationship to other systems in the world. The phenomena to which a theorist intends her model to apply are called the *intended targets* of the model. The Bay model is a paradigm case of a constructed, concrete model, whose target is the San

Francisco Bay. Other historically important examples include ancient Greek models of the planets, Maxwell's mechanical model's of the ether, Watson and Crick's model of the structure of DNA, and scale models of airplane wings and engines. Some of these models, such as the Bay and DNA models, stood in successful resemblance relations. Others, such as the Greek models of the planets and Maxwell's ether models failed to resemble their intended targets because those target systems did not exist.

In addition to models which are literally constructed, scientists can also work with naturally occurring concrete models: structures and phenomena that already exist in nature and resemble other phenomena of interest. Perhaps the most widely-used natural models are model organisms (Griesemer & Wade, 1988; Griesemer, 2003; Weber, 2005; Winther, 2006). Fruit flies, for example, are often called the "test tubes of molecular biology" because of their ubiquity and utility in genetics. For mammalian studies, especially those involving medical research, mice, rats, dogs, and non-human primates can all be studied in place of studying humans. Of course, fruit fly molecular biology is not the same as the biology of all other animals, and mice, while similar to human beings in some ways, are obviously different in many others. No isomorphism exists between these natural models and target systems, yet for particular purposes, these natural models are similar enough to intended targets that studying them is a useful proxy.

A more complex, but substantially similar, case involves the use of natural experiments in population dynamics, geology and climatology (Richardson, 2006). In this case, there isn't a particular object or organism that stands in for another organism or class of organisms. Rather, a dynamic phenomenon taking place in time becomes a model for targets that are inaccessible temporally or spatially. High-pressure water quickly diffusing through rocks in one place might serve as a useful model for low pressure, long-diffusing water in another. While the relationship between these models and theorists' targets are very similar to natural models, far more emphasis is put on behavior than on structural similarity in these cases.

Finally, a variant of the natural experiments studied by geologists are natural population experiments studied by anthropologists. For example, Jared Diamond (1999) argues that geographical factors and the availability of crops and animals for domestication are two of the major factors which determine how effectively one population can dominate another when they come in contact. Diamond's goal in developing this theory is to explain the success that Europeans had in conquering the Americas and Africa. These interactions are so complex and took place on such a large scale that the most effective strategy for determining the operating mechanisms was to study small-scale, regional conflicts. The Maori people from northern New Zealand, for example, nearly exterminated the Moriori from the Chatham Islands when the former invaded. The Moriori had reverted to a hunter-gather lifestyle because Polynesian tropical crops could not grow in the Chatham's cold climate, while the Maori had a relatively advanced agricultural society, which allowed for military specialization. The Maori and Moriori case illustrates some of the causal factors at play in the larger-scale questions discussed in Diamond's book. Like the Army Corps studying the Reber plan's consequences with the Bay model, this natural experiment allowed more direct comprehension of the operant causes without trying to extract them directly from the intended target.

To give a more thorough account of concrete models, we must first ask exactly what kinds of things count as concrete models and how these models are to be individuated. In some ways, the answer to this first question is rather easy. There are few limits on the kinds of things that can serve as concrete models. Even the simplest objects or phenomena stand in many kinds of resemblance relations to other things. These resemblance relations can form the basis of the use of these objects as models, where one thing stands for at least some of the properties of the other. Simple and complex machines, paper and plastic shapes, organisms, and highly complex scale models such as the Bay model all count as concrete models and all stand in myriad resemblance relations to real-world target systems.

How do these concrete objects become scientific models, as opposed to merely objects? They weren't intrinsically scientific models and no magic makes them so. Concrete objects become scientific models because scientists *intend* them to serve as models. The intention that a concrete and steel structure which can be flooded with water should be a model of the San Francisco Bay is what makes that structure a model and what orients it towards its intended target.

Because the spatio-temporal boundaries of concrete models are determined by theorists, not intrinsically, individuating a single model from classes of models and parts of models is not complex. Their individuation simply follows the intentions of their users. The same object may be a whole model for one theorist and a part of a larger model for another. Nevertheless, we should take care to distinguish between the models themselves and our representations of them, what I will call

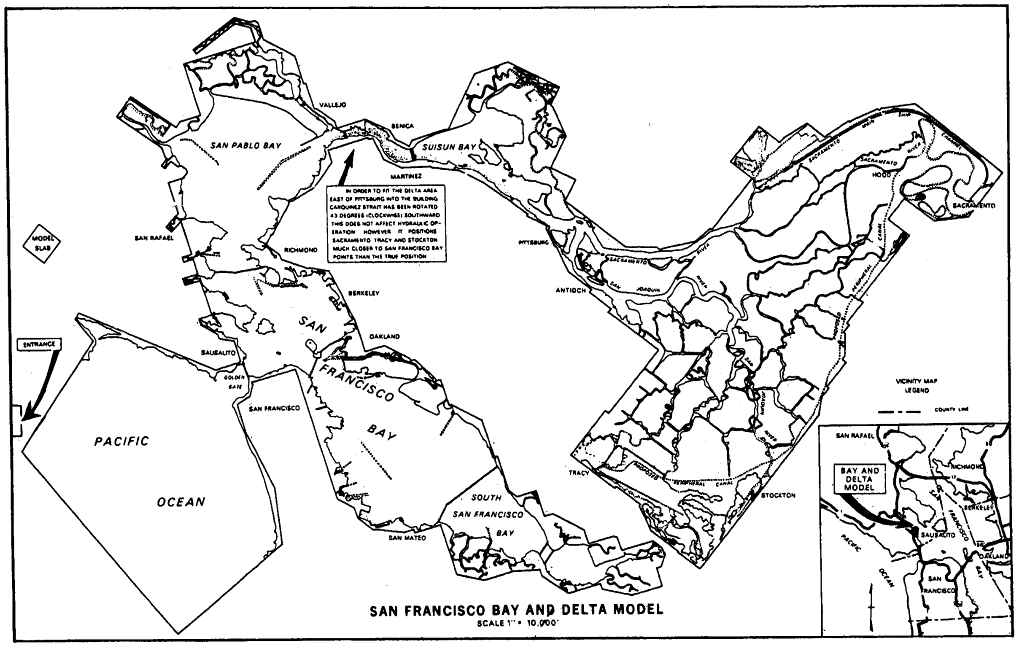

Figure 1: Technical drawing of the San Francisco Bay model showing the model's scale (1:10,000) and orientation. This drawing shows that the portion of the model representing the Suisun Bay and San Jaquin Delta was rotated 43 degrees so that it could fit in the warehouse.

model descriptions. When we talk about models, write about them, or show a picture or diagram, we are employing a model description. In Figure 1, for example, we see a technical diagram for the Bay model, which is a model description of the Bay. Figure 2 is another kind of model description, this time a photograph.

In the case of the diagram of the Bay model and the Bay model itself, the relationship between model and model description is uncomplicated. The diagram is intended to be an accurate description of some of the properties of the Bay model, specifically the scale and orientation of its features. This relationship can be described with standard semantic notions such as truth and reference. If the diagram

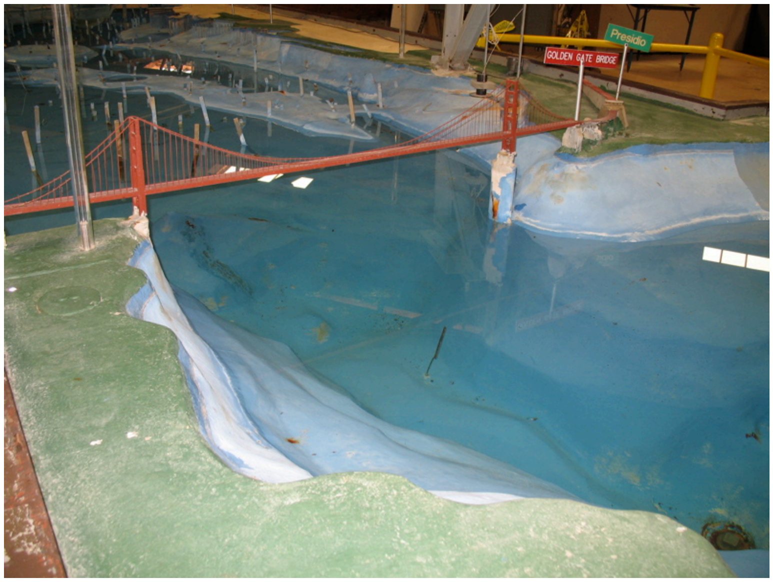

Figure 2: Photograph of a segment of the SF Bay model showing its representation of the Golden Gate Bridge.

is accurate, then it will truthfully describe the properties of the model. Of course, the model is almost certain to have properties not mentioned in the diagram, for example small imperfections in the concrete or the steel superstructure. Standard community conventions about how to read technical diagrams tell theorists that this is simply not specified one way or another. In this respect, the diagram is more abstract then the model, leaving the imperfections in the concrete and other details about the model itself vague.

It is important for theorists to know what is left vague or abstracted away on the one hand, and what is being asserted to be absent. For example, if the diagram of the Bay model does not include a drawing of the sump pump at the mouth of the Bay model, it is not clear whether the diagram is being non-committal about the existence of the pump or asserting that it does not exist. In order to determine this sort of thing, or anything else for which there are not standard reading conventions, model descriptions are often accompanied by a *commentary*. This commentary clarifies how the description ought to be read.

There isn't a one-to-one correspondence between models and model descriptions. A single model can be described in many different ways, such as with blueprints, pictures, equations, or computer programs. In general, a single description can also pick out more than one type of model. While a model description of perfect fidelity describes a single model type perfectly, any amount of vagueness in description allows a single description to pick out multiple models. In general, the vaguer, less precise, or less specific a model description becomes, the greater number of models it represents.

Although it is natural to think of model descriptions as being set down before concrete models are constructed or found, this is not strictly necessary. In some cases, the model is constructed before or without a description. In others, the description comes first. And perhaps most commonly, the two are produced in tandem. When Watson and Crick built their model of DNA, they constructed the model first. In fact, the key to solving the structure of DNA involved seeing physical characteristics of the model and using these to think about the ways that the backbone and nucleic acids of DNA could arrange themselves. Only after the physical model was constructed was a mathematical description of the model written down, in that case to check the validity of the structure from X-ray crystallography data.

In other cases, the description of the model preceded its building. In the Bay model, for example, the Army Corps of Engineers constructed the Bay model by first making detailed technical drawings. The model was then constructed according to the specifications of these drawings. When the Reber plan and other salt water barriers were studied, modifications were made to the model. In these cases, the Corps worked from technical drawings about the temporary modifications they would make in order to evaluate the soundness of these salt water barriers. So there is no general order for the construction of models and their descriptions; differing circumstances dictate creating one or the other first. And in some cases, model descriptions are not generated at all.

3 Resemblance and the Role of the Construal

We now turn to the relationship between models and their targets. When talking about concrete models, the representational relationship between models and targets must be one of similarity or resemblance, not truth or approximate truth. Even if relations like truth and approximate truth hold between models and targets in some way or other, they are far removed from the way theorists assess their models. By themselves, models do not assert anything of scientific interest about their targets. It is only the theorist, reasoning about the resemblances between the model and the world who can make true and false assertions. Of course, philosophers have long known that resemblance and similarity relationships are riddled with puzzles and it will be necessary to make some further refinements in order to

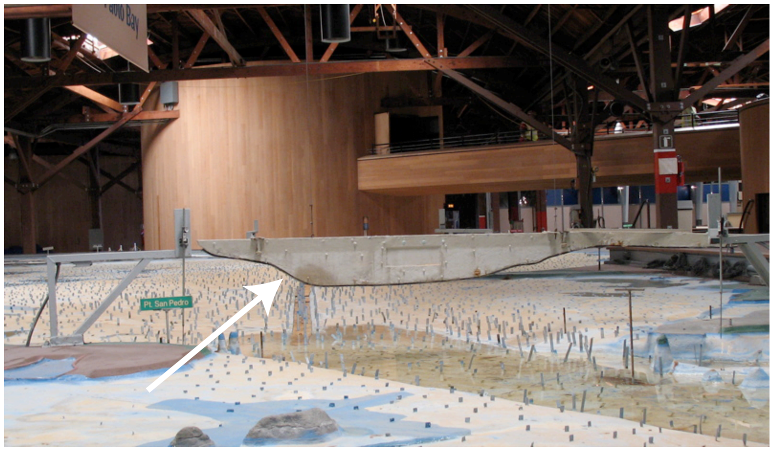

Figure 3: Model gate (indicated by arrow) used to simulate the dams proposed by Reber.

accept that this relationship could be a scientifically important one.

Before attempting to fill in some of the details, let's return to why the Army Corps built the Bay model and some of the uses to which they put it. One motivation for building the model was to find out what would happen if Reber's plan to fill in large portions of the Bay were carried out. For example, Reber claimed that his proposed landfills would create large reservoirs of fresh water. After ensuring the hydrodynamics of the model faithfully resembled the dynamics of the Bay, the Corps created gates which could be lowered in to the model to simulate the construction of Reber's dams, which we can see in Figure 3. After lowering the gates, the Corps observed the instability of fresh water reservoirs in the model. From this, they reasoned that the same proposed reservoirs would not be stable in the Bay itself.

The Army Corps reasoned along the following lines: The model Bay is very much like the real Bay in certain respects, especially those having to do with hydrodynamics. Since our model has these properties, it is reasonable to study the Bay model as a substitute or stand-in for the actual Bay. This kind of reasoning requires that the Bay model was a simulacrum of the real Bay; it resembled it or, to use Giere's expression, was similar to the Bay in certain theoretically important respects and degrees (Giere, 1988). This similarity relationship is really a resemblance relationship.

Concrete modeling requires both *structural* and *dynamical* resemblance relations. Structural resemblance means that the model and target either share or possess similar properties and relations. The physical differences between the Maori's and Moriari's islands are structurally similar to the differences between Eurasia and the Americas. The Bay model's topology resembles the San Francisco Bay's topology. Bond angles and connectivities in Watson and Crick's DNA model are similar to the bond angles and atomic connections in real DNA, and so forth. Dynamical resemblance is resemblance between behaviors of the model and behaviors of the target. More formally, it is a resemblance between the succession of model states and target states. Laboratory fruit-fly genetics, for example, has dynamical resemblance to the genetics of natural populations. The tides and currents of the Bay model resemble the tides and currents of the San Francisco Bay.

Resemblance relationships have a checkered history in philosophy. While it is tempting to use them to explain many semantic phenomena, attempts at formalizing them have been widely thought to be unsuccessful and this has made the philosophical community wary of relying on them.[1] Nelson Goodman voiced some of the most influential critiques of resemblance relations. He believed that appeals to resemblance and similarity had a deleterious effect on philosophical and scientific discussions because they ended up merely labeling something unknown, rather than giving a characterization of the phenomenon in question (Goodman, 1972).

Goodman raised several important arguments against resemblance relations, which must be confronted if a philosophical account of the concrete model/world relationship is built upon them. Perhaps the most important of Goodman's objections is what Quine dubbed "The Problem of Imperfect Community" (Quine, 1969). For any three objects, there will be some respect in which two of the objects resemble each other more than the third. To take a trivial example, imagine a green square and red square and a red circle. Which of these things is not like the other? There is no obvious answer to this question because either pair could be said to resemble each other more than the other pair.

This problem scales up beyond the trivial to real scientific examples. Consider the molecules ethyl alcohol, dimethyl ether, and diethyl ether. Ethyl alcohol has the structure: CH_3CH_2–OH, dimethyl ether: CH_3–O–CH_3, and diethyl ether: CH_3CH_2–O–CH_2CH_3. Which pair of molecules is more similar to each other than they are to the third? Dimethyl ether and ethyl ether are both ethers, hence they can engage in similar chemical reactions and dissolve similar substances. On the

[1] On the other hand, recent work on metaphor (e.g. Camp, 2003) relies on these relations, and gives a sophisticated new account of them.

other hand, ethyl alcohol and ethyl ether both have ethyl groups (chains containing two carbon atoms). Ethyl alcohol and dimethyl ether are structural isomers, meaning that they have the same atoms, just arranged in different orders; they also have the same molecular mass. However, ethyl alcohol is completely soluble in water, whereas both ethers are only partially soluble in water. Ethyl alcohol boils at 78.4°C, while the two ethers boil at a much lower temperature (34.6°C and -23°C respectively). So in neither the trivial case nor in a scientifically realistic case can we say that there is a context-free similarity metric that can be applied. This is the problem of imperfect community.

Of course, chemists make similarity judgments all the time and the problem of imperfect community seems to have gone unnoticed in chemical practice. In this particular case, most chemists would judge the ethers to be the more similar pair of molecules. This is because in organic chemistry, the salient features of molecules are almost always functional groups and conformation, not the atoms they are made up of or the exact length of the carbon chains. Of course, there are contexts where these other features matter more than functional groups, and in those cases chemists adjust their similarity judgments accordingly.

Goodman recognized that this was a potential response, but he didn't believe that this was a very philosophically satisfying one. "One might argue that what counts is not degree of similarity but rather similarity in a certain respect. In what respect, then, must inscriptions be alike to be replicas of one another?" (1972, 438) He continues, later in the article,

> More to the point would be counting not all shared properties but

> rather only *important* properties — or better, considering not the count but the overall importance of the shared properties. Then a and b are more alike then c and d if the cumulative importance of the properties shared by a and b is greater than that of the properties shared by c and d. But importance is a highly volatile matter, varying with every shift of context and interest, and quite incapable of supporting the fixed distinctions that philosophers so often seek to rest upon it. (444)

In everyday appeals to similarity, Goodman may well be correct. Shifting contexts will shift similarity metrics, and appeals to similarity as a brute relation may fail. However, Goodman's response seems inappropriate in dealing with the molecular similarity I described above. Scientific discourse neither requires absolute measures and weightings of similarity nor does it even require that all domains use the same criteria for similarity. Rather, similarity judgments are dependent on myriad factors including background theories, practices in a community, and the intentions of individual scientists. Goodman is correct in thinking that no fully general, purely logical analysis of similarity can be given. But this would not be a very useful account of similarity for scientific practice even if it could be generated.

Consider, now, how this response plays out in the case of the San Francisco Bay model. I would argue that the problem of imperfect community certainly arises, but it doesn't matter very much. In some ways, the model is more similar to a swimming pool than it is to the Bay. It has deep and shallow parts, it is made of concrete and reinforced with steel, it is filled with water which is circulated by pumps, and so forth. Similarly, the model is more similar in some respects to a pinball machine table than it is to the Bay. It sits a few feet off the floor, it is relatively flat, and it has barriers and channels. While these Goodmanized resemblance relations, which make the model more like a swimming pool or a pinball machine are real, they are simply ignored in scientific discourse. This does not reflect a cavalier attitude on the part of scientists. Rather, it simply shows that scientists working with physical models do not concern themselves with wholesale similarity or resemblance. They *choose* which respects are the scientifically relevant ones. This brings us to a very important part of my account of concrete models: the representational intentions of modelers, which I call their *construals*. These construals, along with background theories and community conventions, determine which similarity relations are the relevant ones.

Construals are composed of three parts: the *intended scope* of the model, the *assignment*, and *fidelity criteria*. The intended scope of the model consists of the target systems a theorist intends the model to resemble, as well as a specification of the aspects of the target system that the model is intended to resemble. The assignment coordinates particular aspects of the model to particular aspects of the target system. In many concrete models, such as the Bay model, the assignment is trivial because the model is a scaled physical representation of a target. However, in some cases, such as using a real spring to model a covalent bond, more specific guidance needs to be given about which aspects of the concrete model are essential and meant to represent some aspect of the target. Finally, fidelity criteria set the evaluative standards a theorist will use in assessing her model. These criteria determine both the level of fidelity required by the structure and dynamics of the model, as well as the level of fidelity required for the model's output. Construals often accompany models explicitly as part of a commentary on the model, but they are importantly independent of models. Different theorists may construe the same model differently.

Before turning to mathematical models, let us take stock of the account of concrete models developed so far, as this will guide the development of an account of mathematical models. Concrete models are physical objects that stand in two important relationships.They stand in standard semantic relationships to their descriptions. While notions like truth and satisfaction can be used to described the relationship, its many-many nature make the term 'associate' more appropriate. To target systems in the world, concrete models stand in construal-mediated similarity relationships. Because of the problem of imperfect community, it is impossible to assess this relationship without understanding the theorist's construal, or intentions about how the model is to be coordinated to the world and with what standard of fidelity it is intended to be coordinated. This leaves us with the picture of model world relations we find in Figure 4. Models are associated with model descriptions and can be compared to real-world systems via construal-mediated similarity relations.

To further develop this account, much more needs to be said about the nature of target systems, including their degree of abstraction and how theorists handle phenomena with lesser degrees of localization than the San Francisco Bay. Since many of these issues arise for mathematical models, I now turn to an account of mathematical models and develop it on the basis of what we have learned so far

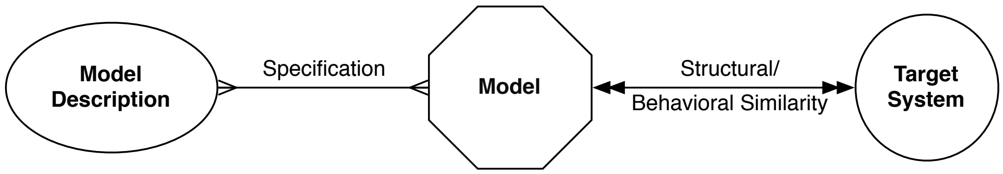

Figure 4: A concrete model and its relations to a model description and a real-world target. The connection to the target system is optional, as some models are studied for their intrinsic interest or without knowing whether or not they are similar to a real world target.

about concrete models.

4 Mathematical Models

The Lotka-Volterra model is a paradigm mathematical model and I will use it to illustrate many aspects of such models. Recall that the model is described by the following two differential equations:

$$\frac{dV}{dt} = rV - (aV)P \tag{3}$$

$$\frac{dP}{dt} = b(aV)P - mP \tag{4}$$

In these equations, V is the size of the prey population and P is the size of the predator population. The variable r stands for the intrinsic growth rate of the prey population and m stands for the intrinsic death rate of the predators. The other parameters describe the prey capture rate (a) and the rate at which predators can convert prey into more predator births (b) (Roughgarden, 1979, 432).

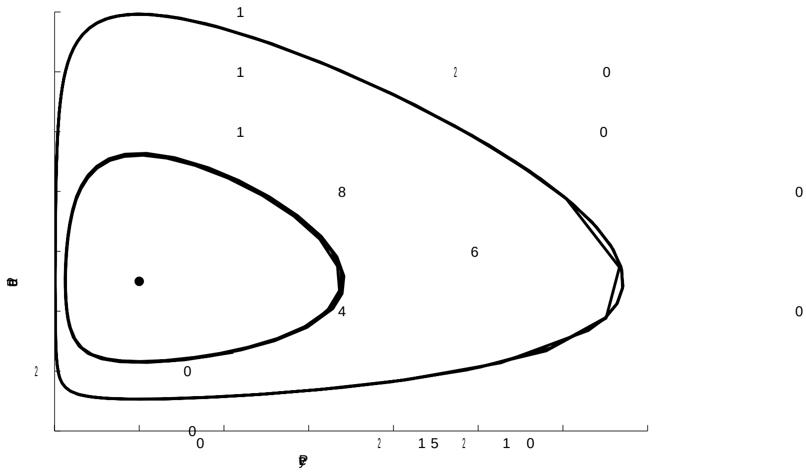

Figure 5: Graph of two dimensions of the Lotka-Volterra model's state space. This graph is often called the phase portrait of the model.

Volterra recognized that solving the equations for their equilibrium would give him the average population density for each of the two species. These average population densities are described by the following equations:

$$\hat{V} = \frac{m}{ab} \tag{5}$$

$$\hat{P} = \frac{r}{a} \tag{6}$$

From these equations, Volterra learned that, in his model predator-prey system, the average abundance of predators primarily depends on prey births, and the average abundance of prey depends on predator deaths. This means that a general biocide, something that effectively increases the predator death rate and decreases the prey birth rate, would actually hurt the predators more than the prey.

Recognizing that heavy fishing was a biocide, Volterra arrived at his answer to the anomalous fishery statistics: Before the war, heavy fishing was shifting the predator-prey balance towards the prey. The low-intensity fishing during the war shifted the balance to favor the predators, but this would abate as heavy fishing resumed after the war. (Roughgarden, 1979; M. Weisberg, 2006; M. Weisberg & Reisman, forthcoming)

Volterra solved the puzzle in the Adriatic by doing mathematics. He wrote down equations that described a mathematical object, a set of points in a three-dimensional *state space* corresponding to the variables P, V, and t. Yet he solved a puzzle about real-world population dynamics. His mathematical model told him something about the Adriatic, just as the Bay model told the Army Corps about the San Francisco Bay. It is relatively straightforward to understand how the constructed, concrete Bay model represents the San Francisco Bay. But it is a much more complex issue, and one that is at the center of many philosophical discussions of scientific models, how a mathematical model can represent a real world system. How is it that the Lotka-Volterra model, or the one-locus model of population genetics, or the harmonic osscilator model of molecular vibration, can model real world systems? The answer to this question depends on what kind of thing a mathematical model is.

5 Two Accounts of Mathematical Models

The Lotka-Volterra model is obviously not concrete and physical like the Bay model or Watson and Crick's model of DNA. Exactly what kind of thing a mathematical model is, however, has been the subject of considerable debate in the literature about models. In a recent review of this debate, Margaret Morrison and Mary Morgan (1999) divide accounts of mathematical models into two traditions. Proponents of *concrete* accounts of mathematical models take mathematical models to be something like imaginary structures that would be concrete if they were real (Hesse, 1966; Black, 1962; Campbell, 1957; Giere, 1988; Godfrey-Smith, 2006). What Morrison and Morgan call the *abstract* tradition includes accounts of models as set theoretical structures (Suppes, 1960a, 1960b), as well as those that take mathematical models to be trajectories through state space (van Fraassen, 1980). I will call this the *mathematical* tradition, to avoid confusion later on. Let's consider these accounts in turn.

5.1 Imaginary But Potentially Concrete

Accounts from the concrete tradition regard mathematical models as imaginary systems that would be concrete if they were real. A biological model of population dynamics, on this view, despite being described using mathematics, is actually an imaginary population of organisms, much like a population of the real world. So on this account, the Lotka-Volterra model of predation consists of an imaginary population of predatory animals and prey animals. These imaginary populations have the properties explicitly attributed to them in the act of modeling — such as growth and death rates, numerical responses, and functional responses. All of their other properties, however, are left vague. This is very similar to the way we construct fictional worlds from a novel or other written text. The text contains only some of the details and the rest must be filled in by us in order to make a coherent story. J. R. R. Tolkien doesn't tell us whether Frodo is left or right handed, but he must be one or the other or ambidextrous. So in order to draw inferences from fiction, readers may have to fill in details, even if these details do not really matter to the story or the author (Lewis, 1978; Ryan, 1980; D. S. Weisberg, 2008). Similarly, in order to draw inferences from the model, the theorist mentally fills in additional properties. But often, and this is part of the point of mathematical modeling, the theorist can simply leave these properties vague.

A recent formulation of this position comes from the work of Godfrey-Smith (2006). He explains that part of the motivation for this interpretation of model ontology is that theorists think of themselves as working with concrete yet imaginary systems when they write down the equations describing their models. He writes:

> I take at face value the fact that modelers often take themselves to be describing imaginary biological populations, imaginary neural networks, or imaginary economies. An imaginary population is something that, if it was real, would be a concrete flesh-and-blood population, not a mathematical object.

Proponents of this perspective point to several advantages of such an account.

The first advantage is that individuating models and specifying the relationship between models and model descriptions are straightforward. Each imaginary system is a model and such systems can be represented in many different ways using words, equations, picture, or graphs. Model descriptions will always underdetermine models conceived of in this way, but this poses no problem and may even be an advantage because imprecise model descriptions can be used to generate families of models with greater degrees of generality (M. Weisberg, 2004).

Another advantage of this kind of account is that the similarity relation between a model and the world is intuitive, just as in the concrete model–world case. A model is similar to a target phenomenon in the world just in case it resembles that target. It is not easy to give a formal analysis of this similarity relationship, but the basic idea behind it is the same as in the case of concrete models. On this view, although imaginary, mathematical models are physical systems that have structural and behavioral similarity relations to real-world targets. More elaborate accounts of the relation will point to the role of theorists' interpretations of different parts of the model's structure via their construals.

Finally, this account has the advantage of taking seriously the way that theorists refine model descriptions on the basis of the mental picture they have of model systems. Godfrey-Smith points to the writings of theorists who describe themselves as first thinking about the model, as if they have some kind of mental picture of it, and then proceeding to write down their model description (equations) on the basis of this mental picture. This is one of the most important insights of the concrete account of mathematical models.

An especially clear example of this can be found in John Maynard Smith's evolutionary genetics textbook. In the course of describing a model of the accuracy of RNA replication, he reasons as follows:

> Imagine a population of replicating RNA molecules. There is some unique sequence, S, that produces copies at a rate R: all other sequences produce copies at a lower rate, r.

In these first steps, he asks us to think about a collection of RNA molecules undergoing the process of replication. Presumably, we can only imagine them because we have had some prior experience with RNA and we can also assume that whatever is standardly true (whether or not we know about it) of RNA is also true of this population. He then asks us to consider a restriction to our initial imagined population: The replication rate is not consistent through the population. It is instead sequence-dependent and one sequence has a greater rate of replication than all of the others.

Maynard-Smith goes on to describe the model in much greater detail:

> A sequence produces an exact copy of itself with probability Q. If x_0 and x_1 are the numbers of copies of S and non-S respectively, then ignoring deaths,
>
> $dx_0/dt = RQx_0,$
>
> $dx_1/dt = rx_1 + RQx_0$
>
> In writing down these equations, I have assumed that when an error

occurs in the replication of a non-S sequence, it gives rise to another non-S sequence will give rise to an S sequence.

In this next step, Maynard Smith further constrains the model by giving more information about the nature of the replication of the RNA molecules. In particular, he specifies the probability of exact replication and, as a result, also specifies the probability of non-exact replication. To complete the model, he goes on to derive an equation that describes the preservation of optimal molecules in the population as a function of differing degrees of accuracy in replication, but that is not our primary concern. Instead, I want to use Maynard Smith's comments to further illustrate the concrete approach to mathematical models.

Godfrey-Smith and other proponents of the concrete approach would argue that this is an especially explicit illustration of a typical way theorists go about modeling. First, Maynard-Smith imagined the model that he was taking about, in this case a population of self-replicating RNA molecules. He then went on to mentally fill in specific properties of the model and, at the same time, wrote down equations which recorded these specifications. As he thought more about the model and analyzed it in detail, he was able to refine it and make it more specific, recording these refinements in the equations that describe the model. All of the modifications to the model were modifications to an imaginary population of RNA molecules.

The concrete account of mathematical models thus has good deal of appeal. Models can be individuated straightforwardly. It gives us an obvious place to begin an analysis of the model–world relation by analogy to physical concrete models and it helps make sense of a very common mode of discourse among modelers.

Each of these advantages, however, raises its own philosophical puzzles. In the next sections, I will discuss several of them, arguing that at least two of them present nearly insurmountable difficulties for the concrete view. Before doing so, I want to consider the standard alternative account.

5.2 Mathematical Objects

The standard alternative account to the kind defended by Godfrey-Smith is that mathematical models are mathematical objects. This is the dominant position in the literature on models, modeling, and the semantic view of theories. Several versions of this view have been discussed in the literature. In one, mathematical models are seen as exactly the same kind of thing as logician's models. This view, which is the original version of the semantic view, analyzed mathematical models by writing down the set-theoretic predicates corresponding to the model (Suppes, 1960a). In more recent literature, models are usually treated as trajectories through state space. This view was first articulated by Bethe (1961) and van Fraassen (1980), and has subsequently been refined by Lloyd (1994). It has become more or less the default version of the semantic view of theories, and is taken by many philosophers of science to be the correct account of models for modeling.

The state-space version of the mathematical objects view begins with the concepts *state* and *state space*. A system's state is a description of all of its properties. For a mathematical model understood in the state-space tradition, the state specifies all and only the properties that the model actually possess. In dynamical processes, states are typically indexed temporally, such that a state description tells us the state of a system at a particular time t. System states can be represented mathematically as a vector, which corresponds to values for a set of variables corresponding to the determinable properties of the system. A state space, then, is a multidimensional space where each dimension corresponds to a particular state and points in that space are specified with state vectors. So for any state of the system, there is a point in the space which corresponds to that state. *State transitions* can be represented by curves or trajectories through the space.

A common way to fill in the account is to say that mathematical models consist of the complete set of trajectories through the state space corresponding to a fully precise specification of a model description's parameters. Thus, all allowable states and evolutions of these states corresponding to a parameter set constitute the model. The set of trajectories through state space which constitute a mathematical model are all the trajectories associated with a fully determinant set of parameters in the description of the model. Each trajectory corresponds to a different set of initial conditions or of values of the independent variables in the model description. Sometimes, this is explained by saying that the model is the state space itself, meaning the structure of the state space. A more accurate gloss is that the model is a *trajectory space*. When the model's state space has a temporal dimension, then the path through state space will represent the temporal evolution of a system's state. Such a model is called a *dynamical model*, and this has been the paradigm case through most of the literature.

Sometimes scientists use the term 'model' to refer to the set of paths through state space associated with an *uninstantiated* model description, an equation with no values set for its parameters. This also parallels discussions of set theoretical predicates associated with model types. For reasons of simplicity, I will avoid this usage and call the full set of trajectories associated with such an uninstantiated equation a *family* or *class* of models. Individual models correspond to equations with exact values set for the parameters, but without values set for the variables. Nothing very important turns on this, however, and one could translate from my way of talking to the other.

Since this account treats mathematical models as abstract, mathematical objects, the model–target system relationship is more complex than on the concrete view. Despite this, the model–world relationship of the mathematical account has been worked out in greater formal detail. In subsequent sections I will defend my own version of this account, but for now, let's consider the traditional way that the this issue has been dealt with in the semantic view and beyond.

In van Fraassen's treatment, a model successfully represents its target when at least one of its trajectories is isomorphic to a mathematical representation of the target. In other words, there must exist a mapping of the trajectory to a representation of the target that preserves the structure and relations of the target. This needs some fleshing out because a set of trajectories through a state space is not isomorphic to, say, a population of organisms undergoing natural selection.

For van Fraassen, the isomorphism is supposed to hold between the model (or what I would call some subset of trajectories associated with the model) and what van Fraassen calls the *appearance* of the system, the subset of the system's state corresponding to observable properties. For example, if I am interested in modeling the periodic motion of a simple pendulum, then the appearance of the pendulum system is the position and momentum of the pendulum during the period of time of interest.

The appearance of the system is the set of actual values of the system through time, not measured values. In moving from a set of measured data points to the appearance, a scientist needs to make an inductive inference of some kind. This inductive inference fills in the as yet unobserved values and removes noise, yielding a mathematical relationship between the observed data. In simple cases, this inference can be simple, linear curve fitting, but more complex possibilities are available as well. The important thing is that real properties of the target, represented mathematically, must be isomorphic to the model.

Many philosophers have noted the stringency of requiring isomorphism. Lloyd, for example, argues that many models are intended to be idealized and that isomorphism is overly stringent. She writes, "In practice, the relationship between theoretical and empirical model is typically weaker than isomorphism, usually a homomorphism, or sometimes an even weaker type of morphism." (1994, 168n.2) Other philosophers have gone further, requiring even weaker relations. Da Costa and French, for example, have argued that the appropriate model-world relation is one of *partial isomorphism*. This three-part relation allows aspects of the model's structure to either be isomorphic to, not isomorphic to, or indeterminate with respect to a mathematical representation of data collected for a real-world system. A model of a gas and a data model abstracted from a real gas, for example, might exhibit such a partial isomorphism between the relations among model molecules on the one hand and data collected about real molecules on the other.

Another option, discussed by Lloyd but not fully developed in the literature, would be to focus on metric relations between the model and a mathematical representation of the target, rather than model-theoretic relations. In such a measure, one would define a relevant metric for each important state variable in the state space. Geometric measures could then be assessed such as, in the simplest case, the distance between model states and the states of the real system. More complex quantitative as well as qualitative relationships could also be assessed.

One can thus see that there are many different relations that have been proposed for the model–world relation in the mathematical account of mathematical models. The literature has tended towards isomorphism accounts and model-theoretic correlates. However, the consensus of the literature has been changing, and much of the recent work on models discusses alternatives.

5.3 Problems for the Accounts

Both the concrete and mathematical accounts of mathematical models have many things to be said in their favor. In this section, I discuss some of their disadvantages and pose these disadvantages as challenges to be overcome in my own account.

5.3.1 Problems for the Concrete Account

The major disadvantages of the concrete account concern the metaphysics of models and the nature of the model-world relation. The metaphysical worry is that imaginary objects seem like strange kinds of things. If we make them less strange and say that they are merely psychological, then one must worry about whether there can be cross-scientist agreement as to their properties. But if they are not merely psychological and if they exist in some non-psychological way, what is this kind of existence? This sort of worry should be familiar from philosophical discussions of abstract objects as well as possibilia.

Godfrey-Smith has developed a simple and reasonable response to this problem. Although he does not attempt to give an account of the metaphysics of such imaginary objects, he argues that they are no more or less worrisome that the imaginary objects in ordinary fiction. Mathematical models are like Tolkien's Middle Earth or Herbert's Arakis. Metaphysicians and philosophers of language will ultimately need to provide an account of the metaphysics of these worlds, and this account may be fictionalist or deflationary. However, it is perfectly obvious that we can reason about these worlds, talk about them, analyze counterfactuals about them, and so forth. Exactly the same can be said of mathematical models understood as imaginary, but concrete systems. The metaphysics remains unknown, but we can still reason with them and about them.

While Godfrey-Smith's response seems plausible, there are lingering worries. The most important of these is what I will call the problem of variation. Insofar as models are imaginary objects, akin to the Orcs of Middle Earth, there may be considerable differences in the way these are conceived of by different scientists. In the fiction case, this variation obviously happens (at least until a movie is made) and poses no problem. If I think that Orcs have human-like feet and you think their feet look a bit more like bear paws, this doesn't pose a problem unless the shape of Orc's feet becomes part of the story. Of course, if they did become part of the story, Tolkien would almost certainly have given us the necessary detail to understand how the story unfolded. Insofar as Tolkien and the story were silent on the issue, it remains an interesting thing to think about, the sort of thing people debate at fan conventions, but nothing critical turns on it.

While the problem of variation poses little problem for fiction, this is not the case with mathematical models. In scientific applications, it would be a serious problem if different scientists, especially ones with the same target system in mind, actually were thinking about slightly different models. There has to be at least a high degree of consensus about the basics in order for modeling and model-based representation to work.

A second problem with the concrete account is more fundamental: How can working scientists make accurate determinations about the absolute or relative fit of a model to a real-world target? While the notion of similarity at the heart of concrete accounts is an intuitive one, intuitive judgments about structural and behavioral resemblance give us little quantitative information of the sort that can be measured in the laboratory. This may not always be required, because some models are developed to investigate very general kinds of phenomena, and are not expected to have any more than qualitative fit to particular targets. But in many other cases, scientists demand that their models have a precise and accurate fit to their intended targets. Doing so assures them that the model can be used to make accurate predictions about the future and that the mathematical structure of the model corresponds to causal structures in the world that give rise to the phenomenon of interest.

The problems discussed so far concerning the concrete account can be handled by the mathematical account. First consider the problem of variation. On the concrete account, this is generated by the fact that different scientists might have different imaginary populations in mind. But on the abstract account, mathematical models are mathematical objects and ought to be universal. How to spell this out, of course, depends on the metaphysics of mathematical objects, but whatever account of the ontology of these objects is correct, that account will insist that trajectories in the Lotka-Volterra state space are consistent across scientists. There is thus no problem of variation on the mathematical account.

Determining model–world fit is also not a problem on the mathematical account. True, saying exactly what the fit ought to consist in has been controversial, but when this relation is specified, then it becomes a mathematical fact whether or not a target and a model have the right kind of relationship. Indeed, one of the best arguments in favor of the mathematical account is that the account mirrors the kinds of determinations actually made by scientists dealing with data and models. However, there are also many challenges to the mathematical account of mathematical models. Before developing a new version of the mathematical account, I turn to some of the problems faced by the versions on offer.

5.3.2 Problems for the Mathematical Account

The first problem for the mathematical account has to do with model individuation. If models are simply mathematical objects, then when two distinct models use the same mathematics, we will not be able to individuate them as separate objects. This situation occurs frequently. Take, for example, the harmonic osscilator model. The same exact mathematics can both be used to describe an idealized spring and a chemical bond. Proponents of the mathematical account of mathematical models would thus conclude that there is a single model being applied to these two cases. A single differential equation (model description) describes the same set of trajectories in a state space (the model) in both cases. However, common scientific usage would have us think that these models are similar, but not exactly the same. Although related, the models are about different phenomena and are interpreted in different ways.

The second problem with the mathematical account is considerably more serious than the first and involves the way that causal information is encoded in models. Many traditional accounts of models in the semantic view only demanded that models were *empirically adequate*, isomorphic (or something weaker) to a mathematical representation of the empirical substructure of a real world system. Many scientists, however, are realists and demand that unobservable state variables and causal structures be accurately represented by their models. It is one thing to have a model with which one can make accurate predictions, it is another to have a model that makes accurate predictions for the right reasons. Such a model would represent real causal structure of the target phenomenon. Can this be done with purely mathematical objects?

The answer would seem to be no. Mathematical objects can have structural and relational properties, but not causal ones. They naturally show correlation, but not causal dependence. It is not obvious how they can distinguish between a properly-formulated forward-looking causal path and a backwards causal path, or even a common cause. So if models are mathematical objects, how can they successfully represent causal structure?

A final problem is accounting for the fact that theorists talk about their models in concrete terms. When discussing a model of predation, for example, a theorist will often describe two populations of organisms that have properties like birth rates and capture rates. These sound like biological properties of concrete objects, not mathematical properties of abstract objects. So a mathematical account of mathematical models is going to need to make sense of this concrete way theorists have for talking about these abstract objects.

These problems can all be handled straightforwardly by the concrete account. There is little problem of individuation with that account because harmonic osscilator spring models are completely distinct from harmonic osscilator bond models. The former contains imaginary springs, the latter imaginary bonds. Since the concrete account treats mathematical models as concrete objects, these objects can have causal structure in a straightforward manner. Since they would be concrete if real, they have all the properties of concrete objects including causal ones. Finally, the concrete account easily handles theorists' talk about models as concrete objects because, on this account, they are imaginary objects that would be concrete if they were real.

The concrete account easily handles some of the most important problems for the mathematical account, while the mathematical account easily handles the problems for the concrete account. This suggests that a hybrid account might be optimum if one could be constructed. In the remainder of this paper, I present such a hybrid account that is at base a mathematical account, but which finds room for many of the important properties of the concrete account.

6 Mathematics and Folk Ontology

Like the defenders of the mathematical account, I believe that mathematical models are ultimately mathematical objects. But I differ from them in two important respects. I emphasize that the mathematical structure must be fully interpreted in order to be compared to a target phenomenon in the world. As in the concrete case, this happens via the theorist's construal, which contains the assignment, intended scope, intended target, and fidelity criteria. The other area in which I disagree concerns what I will call the *folk ontology* of model. I believe that theorists' mental pictures of their model play a significant role in modeling, especially at the early stages of theorizing. This brings my view closer, in some respects, to the concrete account of mathematical models.

6.1 Models, Model Descriptions, and Construals

On my account, mathematical models are mathematical objects, but which kind of mathematical object depends on the model. The dynamical models of much of physics, chemistry, and population biology are usually sets of trajectories through a state space. However, many of the computational models of cognitive science are finite state machines, some models in social theory are directed acyclic graphs, while some other models consist of probability distributions over possible system states, or ensembles. While some of the philosophical literature about mathematical models has tried to argue that there is a single kind of mathematical structure that always corresponds to a model, I see no reason to make this restriction. Many kinds of abstract structures can be deployed in a representation and the determination of whether or not a structure is a model has to do with the representational intentions of the modeler. Nevertheless, a very large number of the mathematical models of interest are dynamical, and all of these models can be described as (minimally) consisting of states and transitions. I will often use this idiom and focus on models that can be so described, but this should not be taken to be a hidden restriction.

Regardless of which mathematical structures can serve as models, mathematical models raise a number of ontological questions not normally considered for concrete models. In particular, the ontological status of mathematical objects — everything from numbers to high-dimensional state spaces — is an outstanding philosophical puzzle. Recent defenses of Platonism, empiricism, and various forms of deflationism and fictionalism make it clear that the ultimate nature of mathematical objects is still unresolved[2]. Consequently, there is little one can say with certainty about the kind of things that mathematical objects are. Because of this, I will bracket the ontological question about mathematical models for this paper.

[2]For a review of this literature, along with a somewhat skeptical conclusion about it, see Balaguer, 1998.

One of the clearest advantages of the concrete account is the close connection between it and the way theorists talk about and seem to conceive of their models. While my account is a mathematical one, I think it is important to preserve this insight, lest mathematical models be thought of as bare structural entities. One way to do this is to place special emphasis on theorists' interpretive intentions, locating intuitive ways of talking and thinking about models here. I call these intentions the *construal* of the model.

Construals are 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 to evaluate the relationship between parts of the model and parts of real-world phenomena. 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 to be studied and the explicit coordination of the model with parts of the real-world phenomenon. This explicit coordination is important for two reasons. First, 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 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 pendulums. Chemists use harmonic oscillators to model vibrations in bonds. So 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. The assignment is a formal record of this type of coordination.

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. Consider the Lotka-Volterra model of predation. This model's main dependent state variables are population densities, but these states and the transitions between them are represented with continuous mathematics. This means that the model can describe transitions between states where the state variable is an irrational number. Volterra certainly did not intend the fact that his model could describe the dynamics of population densities with negative and irrational numbers of predator and prey to correspond to any real or possible population of fish with negative or irrational densities. Thus, in his construal of the model, Volterra only assigned rational values for the state variables (and probably only certain ranges of those numbers) to population densities 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 a target phenomenon intended to be represented by the model. While the assignment coordinates parts of the model and parts of the target, the intended scope is about the target alone. One can think of it as a precise specification of the target.[3]

Intended 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

[3] A similar position is taken by Suppe (1977).

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. 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 intended scope and hence a narrow target, 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 — the predictions it makes about the values of dependent variables given some independent variables — must be to the output of the real world phenomenon. They are often specified as error tolerances. For example, a dynamical fidelity criterion for a predator–prey model might state that the population density 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

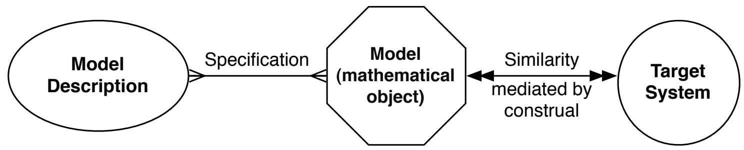

Figure 6: Simple version of the relationship between mathematical models, model descriptions, and targets.

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 Figure 6, I have represented the view described so far. Mathematical models are mathematical objects described by equations, graphs, or other descriptions. Often, but not always, these objects are states and transitions in a state space. The model–world relationship is mediated through the modeler's construal which specifies the meaning of the components of the model, the targets it is intended to be applied to (if any), and the required degree of match between the model and target. This will be elaborated in greater detail in §7, but there is one more major component to my account of mathematical models: *folk ontology*.

6.2 Folk Ontology

One of the important insights of the concrete account of mathematical models, especially as developed by Godfrey-Smith, is that scientists' imagination of concrete systems guides the model descriptions they write down. Godfrey-Smith believes that these imagined systems actually are models, whereas I do not. However, the insight that theorists' imagination of concrete phenomena guides them in writing down and analyzing equations and even in formulating their models is an important one. This needs to be captured in a complete account of mathematical models.

Recognizing that the metaphysics of concrete but imaginary systems is complex and by no means settled, Godfrey-Smith suggests how we might currently construe these systems.

> …to use a phrase suggested by Deena Skolnick [Weisberg], the treatment of model systems as comprising imagined concrete things is the "folk ontology" of at least many scientific modelers. It is the ontology embodied in many scientists' unreflective habits of talking about the objects of their study — talk about what a certain kind of population will do, about whether a certain kind of market will clear. (Godfrey-Smith, 2006)

I suggest that we take this suggestion literally and think about theorists' folk ontology as an additional component in our account of models along with the model itself and its construal. So what exactly is the folk ontology of a model and how does it operate?

Let us look at how the Lotka-Volterra model is presented in Maynard-Smith's important monograph on ecological theory. Several sections after he introduces Volterra's model of predator-prey relations, he asks us to suppose that "some number [V] of the prey can find some cover or refuge which makes them inaccessible to the predator." (Maynard Smith, 1974, 25) This gives us a model which is described by the following differential equations:

$$\frac{dV}{dt} = rV - aP(V - V) \tag{7}$$

$$\frac{dP}{dt} = baP(V - V) - mP \tag{8}$$

He goes on to explain how this new model makes two interesting predictions. First, if the total number of prey in cover is a constant fraction of the total, this does not alter the nature of the oscillation and unstable equilibrium. However, "if the number of prey in cover is constant … the effects of cover are stabilizing, since it changes a conservative into a convergent oscillation." In other words, a constant number of prey hiding from predators stabilizes the oscillation of the model.

The details of this model and its comparison to the original Lotka-Volterra model are interesting, but what I want to focus on is the use Maynard-Smith made of concrete imagery. He began by telling us to imagine a Lotka-Volterra predator-prey system. He then gave us some information about how to modify our imagination — some fraction of the prey population was allowed to find cover, avoiding the predators. No doubt any theorist or other reader would find herself imagining some population of prey heading for cover. Huge variation exists, of course. I first imagined a brushtail possum running up a tree for cover. Others may have though about a lizard hiding under a rock from a Kookaburra. Still others may have thought of clownfish, which hide in anemones. On the concrete view, these mental pictures (or possibly some abstraction of them) is actually the model. But on the view I am advocating, these mental pictures are *aids to thinking about the model*, but are not part of the model itself.

I follow Deena Skolnick Weisberg in calling these mental pictures the folk ontology of models. Just as folk psychology helps people make predictions about the behavior of others, and folk thermodynamics helps people figure out what will be too hot to touch, folk ontology aids theorists in developing mathematical models and the equations that describe them. And just as folk psychology will undoubtedly vary from person to person, but be in near enough agreement to make predictions, so folk ontology can vary among theorists.

Godfrey-Smith and Weisberg are thus correct that many theorists describe imagined concrete systems when they are talking about models, but I advocate interpreting this talk as a commitment to folk ontology, not to models being the way that theorists imagined. Given this view, it is natural to ask whether folk ontology is an essential part of the practice of modeling or is it something that can be completely dispensed with. I believe that the former is correct; folk ontology is essential to modeling in at least three contexts.

The first context in which theorists' folk ontology comes in the development phase of a mathematical model. Take the initial formulation of Volterra's predator-prey model as an example. We don't have the kind of access to Volterra's mental representations that we would like, but it is probably fair to say that he began by imagining a population of predators and a population of prey and attributed to them certain properties. As he wanted to perform a mathematical analysis of this population, he set this idea to paper, writing down equations specifying the model that he had imagined. We do not have a record of this, 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 it. In either case, the mental picture he had — his own folk ontology of the model — guided him in formulating the initial model description and making sure that these equations picked out what he had in mind. So one role of a theorist's folk ontology is to guide the development and refinement of a mathematical model.

The second important context in which folk ontology is important is in thinking about very complex mathematical models. Consider some of the complex mathematical models employed in chemistry. Even highly idealized models of the reactions of simple molecules consist of potential energy surfaces in state spaces of high dimensionality. Of course, no chemist can hold this picture in her mind and hence cannot directly reason about the model. All she can do is manipulate it on the computer. However, she does have access to a mental picture which, more or less, corresponds to the assumptions and idealizations of the model. This need for such mental pictures is all the more dramatic in statistical thermodynamics. Statistical mechanical models of gases are actually ensemble models. They are "a hypothetical collection of an infinite number of noninteracting systems, each of which is in the same microstate (thermodynamic state) as the system of interest" (Levine, 2002, 749). Such an ensemble is not something that can be thought about directly; highly approximate mental pictures are the only ways to think about these systems.

Finally, the folk ontology of mathematical models plays an important role in coordinating models encoded in different representational systems. These models might be very different mathematically, but they nevertheless share many features or assumptions in common. For example, predator-prey models are most commonly developed in an aggregate way, where the main quantities tracked are populations of organisms. However, the contemporary ecological literature has increasingly turned to individual-based approaches, where each organism is represented explicitly as an individual (Grimm & Railsback, 2005; M. Weisberg & Reisman, forthcoming).

Very similar models can be developed in the individual-based and aggregate frameworks in the sense that these models have the same kinds of feedback loops. Yet such models are mathematically distinct because, for example, the state spaces of individual-based models will often have hundreds more dimensions than the state spaces of aggregate models, dimensionality scaling with the number of organisms. On the mathematical account of mathematical models, then, these will be distinct kinds of models. Yet on the concrete account, they can potentially be the very same model, described using different mathematical language. Ecological models are intrinsically individualistic on the concrete view since even imaginary populations are composed of individuals.

Neither of these pure perspectives is satisfactory. Aggregate and individual-based models of the same phenomenon clearly have significant differences in their structural properties. Yet to say that they have no more than a superficial relationship seems too strong. Here is a place that the folk ontology of our models can help.

Folk ontology lets us tie the very complex behavior of the individual-based model back to the simple aggregate model and can function in a similar way whenever theorists need to compare topically similar, but structurally different models.

For all of these reasons, theorists' folk ontologies about their models seem to be a crucial part of scientific practice. Without these mental pictures, it would be difficult to develop mathematical models in the first place, to hold mathematically complex models in mind, and to coordinate similar models embedded in different representational systems. A full account of models must include the role of these mental pictures, but they are not, I believe, most appropriately thought of as the models themselves.

7 The Model–World Relationship

In an earlier section, I discussed the role of the construal in determining the degree of similarity between a mathematical model and the world. However, the discussion did not take a firm stand on the nature of this similarity relationship. For concrete models, similarity involves structural and behavioral resemblance, but for mathematical models, "similarity" is just a place-holder for a different set of mathematical or logical relationships. This section discusses these precise relations of similarity, and ultimately argues that no single relation can fill this role in all cases.

The nature of the model–world relationship is probably the most active area of investigation in the literature about scientific models. Philosophers working on this topic have proposed many accounts but reached little consensus, even for simple cases like the relationship between the harmonic osscilator model and a pendulum or spring. The situation is even more complicated when we consider the phenomena that much of modeling is aimed at understanding: highly complex systems of the sort studied by ecologists and chemists. In these cases, even extremely accurate models are far from complete representations of these phenomena and contain many idealizations.

Some of the more recent developments in the semantic view of theories address approximation, abstraction, and idealization explicitly. Proponents of partial truth accounts of model-based representation argue that good models are typically only partially isomorphic to targets. In other words, some substructures of the target are isomorphic to the model or some substructure of a model, typically via models of data. Partial truth analyses address some of these vexed problems, but do not account for how target phenomena are manipulated so that they can be compared to models. My account of the model–world relation pays special attention to the mismatch of complexity between model and phenomenon. This part of the account deals with what I will call the *preparation* of the phenomenon for representation with a model.

7.1 Preparation

Phenomena in the world undergo preparation before they can be compared to and analyzed with models. Preparation involves three steps: scope specification, abstraction, and parameterization. The latter two steps can be conducted with high fidelity, but more often than not they involve the introduction of approximation.

Breaking this process into three steps is not meant to imply that they always follow one after another temporally. Rather, I want to emphasize that they represent three conceptually distinct stages that a phenomenon must pass through before it is ready to be compared to a model. These steps may all be carried out together and some may only be performed implicitly.

7.1.1 From World to Phenomenon

The first step in preparing a phenomenon to be compared to a model involves the determination of which spatio-temporal regions of the world the scientist wants to study. This process of setting up the spatio-temporal boundary is simply the process of individuating phenomena from the buzzing, blooming confusion of the world *in toto*. What gets included in the boundary is some main object, property, or dynamical process, along with anything exogenous to this object, property, or process that has a causal influence on it. I will call the primary object, property, or process of interest the *phenomenon of interest*. Often, theorists are interested in classes or types of target phenomena. This partitioning of the universe into phenomena may itself raise philosophical problems about object and event individuation, but I will not discuss this topic here and will assume through this article that theorists can carve off parts of the world as individual phenomena.

7.1.2 From Phenomenon of Interest to Target System

With very few exceptions, the initial restriction of phenomena and phenomenon types leaves the theorist with a system which is still too complex to be described, let alone compared to a model. Thus after individuating a phenomenon, a theorist must make decisions about which parts of a phenomenon she will consider and which ones she will not. This is the the process of *abstraction*, a process of further narrowing the phenomenon such that only certain aspects of the phenomenon are considered. The outcome of this process of abstraction is the generation of what I call *target systems*. Target systems are spatio-temporally restricted phenomena abstracted in such a way to make them manageable. Obviously, there are many possible abstractions for any given phenomenon. This means that the process of abstraction will generate a set of potential targets from any given phenomenon.

Some comparison between mathematical models and target systems is already possible at this point. In particular, judgments about structural similarity or resemblance that Godfrey-Smith and Giere think are important are already possible when theorists have an abstracted representation of a target phenomenon. This, of course, requires that models be taken to be imaginary, concrete systems or that the theorist's folk ontology is compared to the target. However, if theorists want to make more exacting comparisons of model and target system, they need to complete the final step of target preparation, where an abstracted phenomenon is prepared to be formally comparable to a mathematical model.

7.1.3 Mathematical Representation of Target

The final step of preparation involves taking a a target system and representing it mathematically in such a way that it can be compared to a mathematical model, a process I call *parameterization*. If a theorist wants to compare the target to a state-space model, she has to also represent the target system in a state space. In such a case, each property making up the state of the target system is assigned to state-variables in a mathematical representation. Once we know the general form of the state for the target system, then we can construct a state space for the target system. This will then let us abstractly represent the properties of the target as points in this state space. It will also let us represent the evolution of the target system as transitions in the state space.

From a purely mathematical point of view, the state space of the target is also a model. Because of this, Suppes (1960b) coined the term 'model of data' for the target's state space. As this name implies, the target's state space can be populated with empirically measured data. Depending on the intended comparison to the model, statistical inferences can be made in order to "clean up" the data and make inferences about missing values. With her model of data in hand, a theorist can begin the process of comparing the model to its target.

7.2 Initial Comparison: Dynamical Sufficiency

Having completed what I call 'preparation,' the theorist can explicitly compare the model to a target system. This happens in several conceptually distinct stages, where different aspects of a the target state-space is compared to the model-state space. There are at least three stages: the assessment of *dynamical sufficiency*, *dynamical adequacy*, and finally *representational adequacy*.

The first kind of comparison of the model to the target is the assessment of the dynamical sufficiency of the model. The purpose of this assessment is to determine if the model is of sufficient dimensionality and possesses the appropriate dimensions to be able to make predictions about the state transitions in the target system. If the target system contains seven properties as well as time and spatial dimensions, the model must have the representational resources to be able to predict the changes in these quantities, assuming the correct boundary conditions and transition rules are set (Godfrey-Smith & Lewontin, 1993).

The easiest way to ensure the dynamical sufficiency of the model is to have the same dimensions as the mathematical representation of the target. This ensures dynamical sufficiency. However, this is not strictly necessary because there are sets of sufficient variables that, when tracked, are sufficient for describing the behavior of much more complex biological and physical systems (Levins, 1966). Indeed, one of the arts of good model building is finding these sets of variables, obviating the need to model a target with something equally complex.

7.3 Dynamical Adequacy

The second kind of assessment between the model and target system examines what I call dynamical adequacy. This is a measure of the fit between the sequence of states of the model and the states of the target, which is often referred to as the *accuracy* of the model. Commonly, these states are temporally sequential states. In such cases, dynamical adequacy is a measure of the relationship between one of the trajectories of a model through its state space and the trajectory of the mathematical representation of the target system through the same state space. The relevant kind of relation and the degree to which this relation must hold are determined by the dynamical fidelity criteria, part of the theorist's construal.

In situations where dynamical sufficiency is ensured by the model having the same dimensionality as the mathematical representation of the target system, isomorphism is a possible relation between one of the trajectories of the model and the representation of the target. However, in the many cases where the model is of lower dimensionality than the target's representation, it is unclear how isomorphism could ever be a possible relation between model and target. In such cases, the best one could hope for is homomorphism. If one insisted on a stricter relation, then one would simultaneously need to require that dynamical sufficiency is achieved by matching the dimensionality of the representation of the target.

In addition to model-theoretic measures of dynamical adequacy, there are many potential metric measures of dynamical adequacy. The simplest such measures tell us the distance, in units relevant to the system, of the model trajectory from the real trajectory in state space. Statistical measures of goodness of fit are the most obvious such relationships. But more complex metric measures can be developed as well, such as ones that bias the importance of the fit in some dimensions over others. These distance measures can be chosen over the entire state space, or, more usually, to some important subspace.

Finally, there are also many kinds of qualitative measures of the goodness of models. These measures concern themselves with gross features of the model and target such as the direction and rate of change, the correct sequence of important events, and the preservation of spatial and causal relations. Some of these measures take the form of informal observations, but considerable effort has been made to develop mathematical tools for the study of these qualitative relationships. Such tools include loop analysis (Puccia & Levins, 1985; Justus, 2005, 2006) and model–target comparison of specific features of the phase space, such as points of stability, equilibrium, limit cycle, bifurcation, and so forth.

These qualitative measures of fit between model and target systems are especially important in modeling, but have received comparatively little philosophical attention. Their importance lies in their promise in explaining how highly idealized models can be compared with real-world systems. There is no sense in which a harmonic osscilator, even one with dampening, is isomorphic to a real spring. Any real spring has far more degrees of freedom and will exhibit anharmonicity and rotation. However, there is clearly a scientifically important relationship between the simple harmonic oscillator and a real spring: Over certain regions, the spring exhibits periodic oscillations that can be described with standard periodic functions like sine and cosine, which are the terms in the solution to the differential equations describing a harmonic osscilator. While most of the models literature has been silent on this type of fit, a great deal of contemporary modeling requires it, and philosophers of science would do well to investigate it further.[4]

Unlike in many previous accounts of model–world relations, I do not argue that one of these relations *is* the model–world relationship. Scientific practice is simply too diverse in this area and I believe that one can find examples of any of these

[4]There are, of course important exceptions to this. One way of understanding Robert Batterman's research program (e.g. 2002, 2001) is as an attempt to understand how grossly inaccurate models nevertheless describe central features of target systems. Vadim Batitsky and Zoltan Domotor (2007) have also discussed this issue in connection with assessing models of chaotic systems.

measures being used as a criterion of dynamical adequacy. Individual scientists and scientific communities determine which of these standards to adopt as well as the degree of fit required for the model to be considered adequate. Depending on the complexity of the system, the data available, the explanatory or predictive goals of the projects, and a number of other factors, theorists can adopt more or less restrictive dynamical fidelity criteria. While the choice of standards is not dictated to scientists a priori, there are more or less rational fidelity criteria to choose depending on the circumstance and scientific goals of the project.

7.4 Representational Adequacy

If one is a realist, then dynamical adequacy is not the only aspect of the model–target system similarity relationship of interest. Dynamical adequacy tells us if the model can make the correct predictions about how the system behaves. However, modelers are often also interested in finding out whether the model reflects the causal structure of the target system that is responsible for producing the behavior. In other words, modelers may want to know if the model is dynamically adequate for the right causal reasons. A model which is *representationally adequate* correctly captures the causal structure responsible for the behavior of the target system. Typically, this causal structure is captured by the transition rules associated with the model.

Determinations of representational adequacy are more complex than determinations of dynamical adequacy because the interpretation of a model as reflecting causal structure in the target is more heavily dependent on the construal, specifically on the assignment. There are formal ways to determine whether the structure of the model and the structure of the representation of the target are isomorphic or share common metric properties, but in order to compare the causal properties of the model to the target, we need to rely on the construal. Specifically, we need to know how the mathematical structure of the model is being interpreted, whether one of the directions of the model represents time, if a certain set of points corresponds to an equilibrium state, and so forth. All of these transcend the mathematical structure alone.

Because of its dependence on the construal, it is even harder to give a recipe for determinations of representational adequacy than it is to give one for determinations of dynamical adequacy. In most cases, however, the determination begins by considering whether the model is intended to even have a causal interpretation. The Lotka-Volterra predator-prey model, for example, is typically intended by ecologists to be interpreted causally, although possibly with low standards of fidelity. The coupling between the two differential equations is intended to represent a causal coupling between two populations of organisms. Conversely, the one-locus model of natural selection is not intended to be interpreted causally. The difference equations this model are supposed to tell us what will happen to the gene frequencies in subsequent generations. The model was derived from a causal scenario—Mendel's laws acting on an infinite population—but as typically interpreted, the recursions are not supposed to reflect causal information.

It is worth emphasizing that this initial judgment cannot be made solely by looking at the mathematics. There is no mathematical reason that a differential equation can represent causal structure, but a difference equation cannot. Rather, it is the use to which these equations are being put that determines whether they are representing causal relations are not. Hence it is imperative to keep in mind the construal, specifically the assignment, when determining whether one has a causal model.

Once a theorist knows that she is working with a causal model, further determinations of representational adequacy can take place. For example, consider the typical case where the coupling between a set of differential equations describing the model is supposed to represent a causal coupling in the world. In such a case, the theorist endeavors to determine empirically whether or not the target possesses such a casual relationship. Similarly, if the equations describing the model have an additive structure, where various variables are linearly combined to determine the value of another, she can determine whether this reflects a real causal relationship in the target, possibly through interventionist experiments.

How is the causal structure of the real system supposed to be determined such that it can be compared to the model? This is of course an excruciatingly difficult question, one that has occupied philosophers of science and metaphysicians for a very long time. There is no hope settling this issue here, but a few observations are worth making.

A growing consensus in the philosophy of science literature about causation is that the kind of counterfactual relations involved in interventionist theory of causation are at least a core element of causal relations (Woodward, 2003). Scientists explicitly intervene on and control variables so that they can study the effect that some variables have on others. Something similar is done with models themselves: Parameters can be changed, ranges of boundary values can be tested, and new couplings or functions can be introduced to models, all with the goal of investigating how this affects the model's behavior.

So one way that modelers investigate the causal properties of the real world system and try to ensure these are represented by the model is to plan parallel experiments between target system and model. As much as is feasible, they manipulate properties of the target system and see the systematic effects of these manipulations. They can then conduct parallel manipulations on the parts of the model intended to represent those parts of the target system. Converging results between model and target system are evidence that the model is capturing the real causal structure of the system.

Thus a full determination of representational adequacy of a model can rarely be accomplished with a single model described by a fully instantiated model description. Rather, one needs to examine different initial conditions (within a single model) and different parameter sets (from other models in the family) in order to determine the systematic effects of these changes on the model's behavior. Determining the representational fidelity of a model involves active investigation of the properties of models drawn from a family.

Finally, I should note that as with dynamical fidelity, a modeler's representational fidelity criteria set the standards for how good of a match there must be between the model's causal structure and the target system. Sometimes, modelers may not care at all about capturing causal structure; perhaps only the output mat-

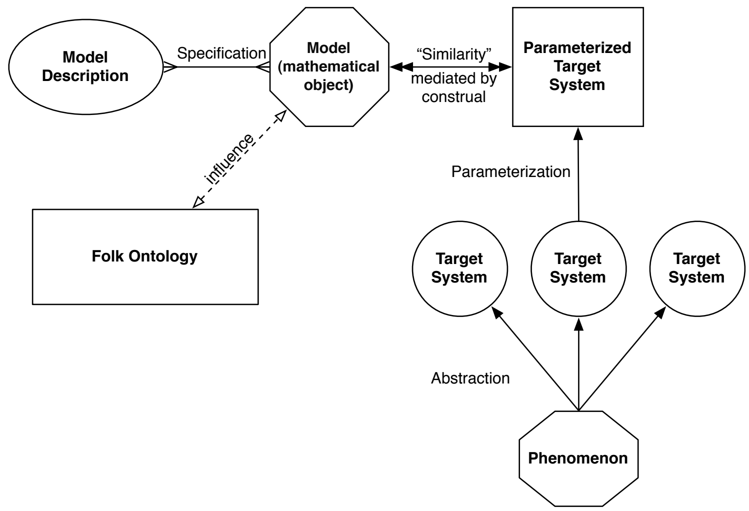

Figure 7: Mathematical models, their auxiliary properties, and their relationship to real-world phenomena.

ters. Sometimes, they may be content to have the the primary causal factors of the target captured in the model. And in other cases, theorists may be after a very precise and accurate that can not only make extremely high fidelity predictions, but that also captures, with high fidelity, the causal structure of the target system.

8 Models for Modeling

We come to a complete picture of mathematical models and their relationship to real-world phenomena, represented pictorially in Figure 7. Mathematical models are mathematical objects, described by model descriptions. A number of different kinds of mathematical objects can serve as mathematical models, but often they are sets of points and trajectories in state spaces. The model is interpreted via the theorist's construal, which determines what each dimension of the state space means. When a model is intended to be compared with a real world phenomena, the nature of this phenomenon is specified in the intended scope, another part of the theorist's construal. Such a phenomenon must then be prepared to be compared to mathematical models. This happens in several stages. First, a spatio-temporal region of the world is carved off according to the specification of the scope. Secondly, much of this spatio-temporal region is abstracted away, leaving a target system, which is those aspects of target object that will be considered by the theorist. Finally, the theorist creates a mathematical representation of the target system, which can be compared to the mathematical model. This comparison is assisted by the theorist's folk ontology and is evaluated according to standards set by the theorist's fidelity criteria.

With a complete account of scientific models sketched, I want to highlight some of the features of this account that play the most important roles in the practice of modeling. Exploring these connections further is a significant undertaking, but I will outline three major themes.

First, the account of mathematical models I have laid out draws a strong parallel between mathematical models and concrete, physical models. Both concrete and mathematical models stand in many-many relationships with their model descriptions, require interpretation to be fully realized, and can potentially stand in many different kinds of model–world relations with target systems. Most of the modern literature about scientific models pays little attention to concrete models, as their use has become considerably less important with the rise of computational methods. Nevertheless, I believe that the strong parallels between these accounts will motivate further research into the representational capacities and fruitful uses of concrete models. We should not confine philosophical study to mathematical models alone.

More central to the concerns of extant accounts of modeling practice are the many roles for theorists' intentions that I have discussed in this article. Theorists' intentions play a role in determining what counts as a model, how the model is individuated, which aspects of the world are to be parts of the target, which bits of the model represent which bits of the world, and what standards of fidelity are used to evaluate the model. All of these roles are worthy of their own study, and a complete account of modeling practice requires understanding all of them. If we are to understand the ways in which modeling can be conducted rationally, then much more needs to be known about the kinds of decisions that can be made about these factors and what effect these decisions have on the representational power of models.

Finally, I have strongly emphasized what I called the folk ontology of mathematical models. These are the beliefs and mental images that individual and communities of scientists associate with the abstract mathematical objects that strictly constitute mathematical models. I have argued that they play several indispensable roles in modeling. If this is correct, then traditional attitudes about the unimportance of psychology and pragmatics (e.g., those in Hempel, 1965) can have no place in detailed accounts of models and modeling. Even if we choose, as I have, more formal analyses of the structure of models and model-world relationships, much of what is important about the practice of modeling will be missed if they are the exclusive focus of our philosophical attention.

References

Balaguer, M. (1998). *Platonism and anti-platonism in mathematics.* New York: Oxford University Press.

Batitsky, V., & Domotor, Z. (2007). When good theories make bad predictions. *Synthese*, *157*, 79-103.

Batterman, R. W. (2001). *The devil in the details: Asymptotic reasoning in explanation, reduction, and emergence.* New York: Oxford University Press.

Batterman, R. W. (2002). Asymptotics and the role of minimal models. *British Journal for the Philosophy of Science*, *53*, 21–38.

Beth, E. (1961). Semantics of physical theories. In H. Freudenthal (Ed.), *The concept and the role of the model in mathematics and natural and social sciences* (pp. 48–51). Dordrecht, Holland: Reidel.

Black, M. (1962). *Models and metaphors.* Ithaca: Cornell University Press.

Camp, E. A. (2003). *Saying and seeing-as: The linguistic uses and cognitive effects of metaphor.* Unpublished doctoral dissertation, University of California, Berkeley.

Campbell, N. R. (1957). *Foundations of science.* New York: Dover.

da Costa, N. C. A., & French, S. (2003). *Science and partial truth.* Oxford: Oxford University Press.

Diamond, J. (1999). *Guns, germs, and steel: The fate of human societies*. New York: W. W. Norton.

Giere, R. N. (1988). *Explaining science: A cognitive approach*. Chicago: University of Chicago Press.

Godfrey-Smith, P. (2006). The strategy of model based science. *Biology and Philosophy*, *21*, 725–740.

Godfrey-Smith, P., & Lewontin, R. (1993). The dimensions of selection. *Philosophy of Science*, *60*(3), 373–395.

Goodman, N. (1972). Seven strictures on similarity. In *Problems and projects*. Indianapolis: Bobbs-Merril.

Griesemer, J. R. (2003). Three-dimensional models in philosophical perspective. In S. de Chadarevian & N. Hopwood (Eds.), *Displaying the third dimension: Models in the sciences, technology and medicine*. Stanford: Stanford University Press.

Griesemer, J. R., & Wade, M. J. (1988). Laboratory models, causal explanation, and group selection. *Biology and Philosophy*, *3*, 67–96.

Grimm, V., & Railsback, S. F. (2005). *Individual-based modeling and ecology*. Princeton: Princeton University Press.

Hartmann, S. (1998). Idealization in quantum field theory. In N. Shanks (Ed.), *Idealization in contemporary physics, poznan studies in the philosophy of science and the humanities 64* (pp. 99–122). Amsterdam: Rodopi.

Hempel, C. G. (1965). Aspects of scientific explanation. In *Aspects of scientific explanation and other essays* (pp. 331–496). New York: The Free Press.

Hesse, M. B. (1966). *Models and analogies in science*. South Bend: University of Notre Dame Press.

Huggins, E. M., & Schultz, E. A. (1967). San francisco bay in a warehouse. *Journal of the IEST*, *10*(5), 9–16.

Huggins, E. M., & Schultz, E. A. (1973). The san francisco bay and delta model. *California Engineer*, *51*(3), 11–23.

Jackson, W. T., & Paterson, A. M. (1977). *The sacramento-san joaquin delta : the evolution and implementation of water policy : an historical perspective*. Davis: California Water Resources Center, University of California.

Justus, J. (2005). Qualitative scientific modeling and loop analysis. *Philosophy of Science*, *72*, 1272-1286.

Justus, J. (2006). Loop analysis and qualitative modeling: limitations and merits. *Biology & Philosophy*, *21*, 647-666.

Levine, I. N. (2002). *Physical chemistry* (5th ed.). Boston: McGraw-Hill.

Levins, R. (1966). The strategy of model building in population biology. In E. Sober (Ed.), *Conceptual issues in evolutionary biology* (First ed., pp. 18–27). Cambridge, MA: MIT Press.

Lewis, D. (1978). Truth in fiction. *American Philosophical Quarterly*, *15*, 37–46.

Lloyd, E. A. (1994). *The structure and confirmation of evolutionary theory* (second ed.). Princeton: Princeton University Press.

Maynard Smith, J. (1974). *Models in ecology*. Cambridge: Cambridge University Press.

Morgan, M. S., & Morrison, M. (1999). Models as mediating instruments. In

M. Morgan & M. Morrison (Eds.), *Models as mediators* (pp. 10–37). Cambridge: Cambridge University Press.

Puccia, C. J., & Levins, R. (1985). *Qualitative modeling of complex systems : an introduction to loop analysis and time averaging.* Cambridge, Mass.: Harvard University Press.

Quine, W. (1969). Natural kinds. In *Ontological relativity and other essays.* New York: Columbia University Press.

Richardson, R. C. (2006). Chance and the patterns of drift: A natural experiment. *Philosophy of Science*, *73*, 642-654.

Roughgarden, J. (1979). *Theory of population genetics and evolutionary ecology: An introduction*. New York: Macmillan Publishing Co.

Ryan, M. L. (1980). Fiction, non-factuals, and the principle of minimal depature. *Poetics*, *9*, 403–422.

Strevens, M. (2006). *Why explanations lie: Idealization in explanation.*

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.

Suppes, P. (1960a). A comparison of the meaning and use of models in mathematics and the empirical sciences. *Synthese*, *12*, 287–300.

Suppes, P. (1960b). Models of data. In E. Nagel & P. Suppes (Eds.), *Logic, methodology and the philosophy of science: Proceedings of the 1960 international congress* (p. 251-61). Stanford: Stanford University Press.

United States Army Corps of Engineers, San Francisco District. (1963). *Technical* *report on barriers : A part of the comprehensive survey of san francisco bay and tributaries, california*. San Francisco: Army Corps of Engineers.

van Fraassen, B. C. (1980). *The scientific image*. Oxford: Oxford University Press.

Volterra, V. (1926). Fluctuations in the abundance of a species considered mathematically. *Nature*, *118*(558–60).

Weber, M. (2005). *Philosophy of experimental biology*. Cambridge: Cambridge University Press.

Weisberg, D. S. (2008). *The creation and comprehension of fictional worlds*. Unpublished doctoral dissertation, Yale University.

Weisberg, M. (2004). Qualitative theory and chemical explanation. *Philosophy of Science*, *71*, 1071–1081.

Weisberg, M. (2006). Robustness analysis. *Philosophy of Science*, *73*, 730–742.

Weisberg, M. (2007). Who is a modeler? *British Journal for Philosophy of Science*, *58*, 207-233.

Weisberg, M. (forthcoming). Three kinds of idealization. *Journal of Philosophy*.

Weisberg, M., & Reisman, K. (forthcoming). The robust volterra principle. *Philosophy of Science*.

Wimsatt, W. C. (1987). False models as a means to truer theories. In M. Nitecki & A. Hoffmann (Eds.), *Neutral models in biology* (pp. 23–55). Oxford: Oxford University Press.

Wimsatt, W. C. (2007). *Re-engineering philosophy for limited beings*. Cambridge, MA: Harvard University Press.

Winther, R. G. (2006). Parts and theories in compositional biology. *Biology &*

Philosophy, *21*, 471-499.

Woodward, J. (2003). *Making things happen*. Oxford: Oxford University Press.