

Black Holes, Black-Scholes, and Prairie Voles:
An Essay Review of *Simulation and Similarity*,
by Michael Weisberg*

Cailin O'Connor

Department of Logic and Philosophy of Science
University of California, Irvine

James Owen Weatherall

Department of Logic and Philosophy of Science
University of California, Irvine

March 8, 2016

Michael Weisberg's *Simulation and Similarity* reflects the state of the art in the philosophical literature on scientific models. It presents an account of modeling that aims to accommodate essentially all examples of scientific models discussed in the extant philosophical literature, and many more besides. The view that results is pluralist, in the sense that Weisberg tends to divide and conquer. He recognizes three broad classes of models, and within each class, he identifies multiple sub-classes. Cutting across these classes, he draws distinctions between “targeted” and “untargeted” modeling practices, and analyzes these separately. Cutting in yet another direction, he identifies three varieties of idealization, which he argues function in different ways, but all of which involve some kind of distortion of the model with respect to its target. Many of these distinctions are intended to broaden the tent, to draw attention to features of the practice that others have neglected: the philosophical literature, he argues, has tended to focus on a few kinds of model, rather than addressing the breadth of modeling practice.

All of this is to the good: drawing attention to the richness and diversity of models in the wild is a valuable contribution in itself, and faced with this diversity, his pluralism serves him well. There are also striking successes in his attempts to regiment the subject. Indeed, the framework he has devised has already given form to the modeling literature in the wake of the book. He has set the terms for future work on this subject. Perhaps it goes without saying that this is a major accomplishment on two counts: the book succeeds in its goal of

*We are grateful to Steve Downes, Sarita Rosenstock, Kyle Stanford, Michael Weisberg, and John Norton for detailed comments on previous drafts of this review. This review was written while the authors were Visiting Fellows at the Center for Philosophy of Science at the University of Pittsburgh. We are grateful to the Center for its support, and to the other Fellows—Francesca Biagoli, Agnes Bolinska, Michel Janssen, Nancy Nersessian, Mike Stuart, and Matthias Unterhuber—for a lively and helpful discussion of the review.

providing tools for thinking about a very broad class of models, and those tools have already been recognized as fruitful by others working in the field. There are few stronger forms of praise.

That said, for reasons we describe in detail below, the very virtues we have just noted are at the heart of our ultimate dissatisfaction—with the book, yes, but perhaps more with the widespread and influential program in philosophy of science that this book contributes to.

In short, Weisberg is most convincing when he argues that various models function in ways that are importantly different from one another—although we tend to think there remain further depths to plumb along these lines. These arguments provide a helpful corrective to a literature that often focuses on isolated examples. The book is less convincing when it pivots to argue that the particular distinctions Weisberg draws exhaust the dimensions along which models may fruitfully be seen to differ. The problem with this is not that a broad range of models, and indeed perhaps all models anyone has ever case-studied, cannot be shoe-horned into a tripartite taxonomy with sufficiently many sub-parts. Rather, the problem is that the more one appreciates the richness of modeling practices in science—a richness Weisberg has done more than anyone to highlight—the less compelling it is to think that the philosophically and scientifically important features of models are the ones they have in common. The term “modeling”, much like the term “science”, picks out a set of practices that do not constitute any sort of natural category. For this reason, studying models in science at the level of generality and abstraction attempted here is not just herculean, but quixotic.¹

We will proceed in two stages. First we will describe the structure and arguments of the book in more detail. Along the way, we will emphasize the places where we think Weisberg is most successful. Second we will return to the issues just raised and explore them in more detail by drawing on even more examples of models and modeling practices, from physics, economics, finance, and evolutionary biology, that we think cause problems for the unified view of modeling that Weisberg provides.

Simulation and Similarity begins by presenting and elaborating the taxonomy mentioned above. Weisberg divides models into three major categories—mathematical, computational, and physical. *Physical* models consist of real world, physical materials. Weisberg uses the San Francisco Bay-Delta model—an enormous scale model of the bay complete with hydraulic tides—as a paradigm case of this type of modeling. *Mathematical* models consist in mathematical structures, such as systems of equations. The Lotka-Volterra model from biology, which describes dynamical relationships between predator and prey populations, is the principal example for this category. Finally, *computational* models iterate some computational process, usually tracking rules of change for some system. As the key example of this sort of model, Weisberg uses Schelling’s famous segregation model, where actors on a

¹Others have also worried about univocal accounts of modeling in science. (See Downes (2011) for an overview.) Our worry is perhaps more basic, since we do not see enough of a family resemblance to justify understanding “models” as a fruitful unit of analysis at all.

lattice decide whether or not to move based on the racial makeup of their direct vicinity. (We return to this taxonomy in the next section.)

In chapter 3, Weisberg goes into further depth on the anatomy of models. After fleshing out the three part taxonomy in more detail, he discusses the role of *model descriptions* in modeling. For Weisberg, a description does not define models (as it does for Giere (2010), for example). Instead, descriptions and models stand in many-to-many relationships where a single model can be described in multiple ways, and a single description can translate into many models, especially if it is vague. There are some puzzles, here, related to whether it is ever possible to pin down a model with a description on Weisberg’s account; one also wonders whether there are facts about abstract (mathematical) models that are not specified by, say, a system of equations, and if so, what the character of those facts is meant to be. It might be better for Weisberg to merely allow that the model-description relationship is usually many-to-many, while also allowing that in some cases, a description does uniquely pick out a model. But this is a small complaint.

In this chapter, he also describes *construals*—the interpretation part of the model—in more detail. Construals, according to Weisberg, consist in four things: an *assignment*, specifying what in the model maps to what in the target system; *scope*, determining what in the target system is being represented by the model at all; and two types of *fidelity criteria*, which capture the degree of model-world match that will be acceptable to the modeler. *Dynamical fidelity criteria* are supposed to track (roughly) how accurate the output of models must be, with respect to real world observables, for the model to be considered successful. *Representational fidelity criteria*, on the other hand, track how well the structure of the model maps the structure of the world.

Here Weisberg emphasizes the role of the modeler in giving models semantic content, via the construal. This is an important point that, although related to arguments made by others, such as Giere (2004) and Van Fraassen (2010), who emphasize the pragmatics of representation, is still often overlooked—for instance, in the philosophy of physics literature, where some authors try to read modal information off of a collection of models of a theory without regard for pragmatic issues. Models do not stand in rigid relations with situations or objects in the physical world that they “represent”. Rather, as Weisberg rightly argues, various physical and mathematical objects—structures—have the capacity to represent many other things, in many ways. How they do that representational work is a matter of interpretation, design, and use, and is ultimately flexible.

Weisberg moves on in chapter 4 to describe the fictions account of modeling—that mathematical models are in fact concrete and imaginary and so can directly represent systems that are concrete and real. He rejects it as a useful way to think about what models are and how they represent. We find his critiques convincing, but will not go into detail here, as these largely negative arguments are not central to Weisberg’s unified account. Weisberg does describe a positive role for fictions in modeling, which is as the “folk ontology” of many modelers.

In the next three chapters of the book, Weisberg describes, in much greater detail, the process of building and using models. He also dramatically expands his bestiary of examples, offering consistently insightful analyses of many different models, from many fields. He begins with target-directed modeling, which is meant to capture the “simplest of case modeling” where modelers attempt to understand a single, specific target system. This chapter walks

through modeling practice: how such models are developed, how they are analyzed, how they are fit to their target systems. Of course, these processes vary wildly across modelers and modeling communities, and so much of the discussion draws on specific examples, with modest generalizations of the form, “Some modelers do X some of the time”. But there is no more to be said, and so it is commendable that Weisberg does not try to do more.

At the end of this chapter, he turns to the question of how modelers can compare mathematical models to target systems if the former consist of mathematical structures and the latter physical objects. His answer is that mathematical models are directly compared to mathematical representations of the target system, not the system itself. Concrete models, on the other hand, can be directly compared. In this way he takes an opposite approach from the fictions view, which solves the problem by arguing that mathematical models are in fact concrete, by instead presenting the target system as in fact mathematical. This view reappears throughout the book, but it seems to us merely defer problems. How are mathematical representations of the target system to be compared with the target system? If an answer is forthcoming, then why bother with the middle man? (This is not to say we endorse the fictions account.)

Before moving on to discuss modeling without specific targets, Weisberg turns to the topic of idealization. Idealized models intentionally distort aspects of the systems they represent. Weisberg recognizes that idealization is many-faceted, and so, once again, provides a three part taxonomy of idealization practices. *Galilean Idealization* is motivated by a desire for tractability, with an implicit goal of eventually reducing the level of idealization as much as possible. *Minimal idealization*, on the other hand, reduces complexity of systems in order to generate scientific explanation, especially of causal relationships in the target system. *Multiple-Models Idealization* includes the production of multiple models with independent assumptions to represent the same target phenomenon, for example to generate explanation, or to predict highly complex phenomena. Here Weisberg uses a taxonomy in a flexible, and ultimately successful, way. Instead of treating this as the only correct division, he clearly intends it to be one useful way to differentiate and understand modeling practice.

One especially nice feature of Weisberg’s account is that it emphasizes the ways in which the idealizing choices made by modelers are driven by their particular goals. While there are many goals a modeler might have, Weisberg broadly articulates some important ones, such as providing a complete representation or maximally simplifying a model’s structure, and then shows how these goals relate to the strategies for idealization he describes. Unsurprisingly, this chapter has been particularly influential—in part because it builds on earlier work in Weisberg (2007) that already had a following.

This analysis sets the stage for modeling that is not directed towards a specific target. Weisberg attempts to tackle the heterogeneity of modeling practice here by (again) dividing these models into three types—*generalized modeling*, *hypothetical modeling*, and *targetless modeling*. Generalized modeling is arguably the most similar to target directed modeling, as it focuses on target systems, but general rather than specific ones. Hypothetical models represent possible, but not actual, target systems, such as infinitely growing populations and perpetual motion. Targetless modeling is supposed to refer to exploration of constructed systems that do not have any intended target. It is unclear how this sort of practice can count as modeling on Weisberg’s view given that there are no intended interpretations of the structures at hand. This small complaint aside, his analysis here is successful in the same

ways as the previous two chapters—it draws on detailed, specific examples of modeling and uses the categories he outlines flexibly to understand these.

From here, Weisberg turns to his account of model-world relationships. Like previous authors such as Giere (2010) and Cartwright (1983), he takes this relationship to be one of similarity, rather than isomorphism or homomorphism. So far, so good. To capture the similarity between models and their target systems, the account offers a formalism, called *weighted feature matching*. We will not reproduce it fully here, but the idea is that one can generate a list of relevant features of the model and the world—‘oscillation’, ‘is a Lyapunov function’, ‘equilibrium’, etc. This list can then be used to generate sub-lists of features shared by the model and world, those only in the model, and those only in the world. These sub-lists are weighted and combined to generate a number between 0 and 1 that measures the model-world similarity. Weisberg gives desiderata that such an account should meet—allowing for qualitative and not just quantitative similarity, for example—and argues that his account is up to the task. We found this account of similarity less appealing than the other views defended in the book, but we will defer further commentary to the next section.

The final chapter addresses robustness analysis—the practice whereby modelers assess the explanatory and predictive reliability of their models by testing the robustness of these models to slight changes. Modelers can use this process to determine which of their results arise from relevant causal features of their models, and which from irrelevant features. Weisberg distinguishes between *parameter robustness*, which is robustness over alterations of parameter values in a model; *structural robustness*, which is robustness over structural changes to features of the model; and *representational robustness* where modelers use different representational assumptions to model the same phenomenon.

Weisberg’s story here is very nice, and effectively counters unintuitive skeptical claims from the literature about the usefulness of robustness analyses, such as those made by Orzack and Sober (1993). Weisberg ultimately argues that robustness analysis can provide what he calls *low level confirmation* involving the production of conditional statements of the following form: “Ceteris Paribus, if agents’ decisions about where to live are guided by the Schelling utility function and movement rules, then segregation is inevitable” (169). The idea is that if a modeler manages to show a robust dependency relationship between a causal factor (Schelling’s utility function) and a result (segregation), they can later use this dependency to assert the result (ceteris paribus) should the causal factor in fact hold of the world.

With this overview of the book in place, it should be clear how much Weisberg accomplishes, and how strong many of these chapter-by-chapter contributions are. But we are also now in a position to step back and ask whether the book does what it sets out to: Has it succeeded in giving us a satisfactory account of scientific models and modeling, taken as a whole? We think not. In fact, after reflecting on the considerations raised in the book, we do not think an account anything like what is envisioned here could be successful. In this part of the review, we will defend both of these reactions: first, we will argue that the unified account of modeling given in the book does not do what it sets out to; second, we will argue that no such account could succeed in doing so.

To make our case, we need to say something about what a satisfactory account would accomplish. Weisberg himself offers a compelling answer to this question. Towards the end of chapter 2, in the course of a defense of his particular taxonomy of models, he distinguishes several levels of analysis of scientific models. One level is a matter of descriptive sociology: what sorts of models do scientists recognize or distinguish? Another concerns metaphysics: what sorts of things are models, fundamentally, and how many kinds are there? The final level is what Weisberg calls the *epistemic* one: what sorts of taxonomic categories, methods of analysis, and so forth do we need to understand modeling practice? Weisberg is principally interested in the final one—and so are we. So it is with this question in mind that we evaluate the account as a whole: To what extent does the book provide the conceptual tools needed for fruitful analysis of scientific models, taken all together?

The first thing to say is that much of Weisberg's book actually *does* provide tools of just the right sort. In particular, chapters 5-7 and 9, where Weisberg discusses targeted and untargeted modeling, idealization, and robustness analysis, operate successfully at the epistemic level of analysis. Weisberg brings in many examples, he discusses them at a useful level of detail, he draws important, easy-to-overlook distinctions, makes insightful comparisons, and leaves the reader with a clearer understanding of how various aspects of modeling practices work. Nothing we say in what follows should impugn these chapters; indeed, they should set the example for the future of this literature. At the same time, we do not detect anything in them that rises to the level of a univocal account of scientific models.

Even so, as we noted at the beginning of the review, the book certainly *does* attempt to do the more ambitious thing. As we see it, the grand account of models appears in chapters 2 and 3, with a coda in chapter 8 on the “representation relation” between models and the world. And so, to evaluate the account as a whole, on the epistemic level, we need to focus on what the book says about taxonomy, anatomy, and representation. We will begin with the taxonomy, and then expand the discussion to consider the other parts of the account.

Weisberg's taxonomy has some important virtues. It certainly allows us to draw useful distinctions for some purposes. Indeed, it is helpful for the three models that he focuses on in his discussion. Moreover, we do not dispute that a broad range of models—perhaps all models—can be fit into the three categories he offers us. But such virtues are not sufficient for this taxonomy to be the basis for an adequate account, in the grand sense, at the epistemic level of analysis. For this, it must be established that these are the right categories, or the right sorts of distinctions, for fruitful analysis of any scientific model at all. Weisberg recognizes this: he offers an extended argument in chapter 2 for why the three categories he describes *are* the right ones for understanding any model at all.

It is here that the biggest problems lie. As we will presently argue, there are myriad examples of models that fall within his categories that are as importantly different from one another as models that fall into different categories. And conversely, there are models that we are forced to put into different categories even though they are importantly similar—and importantly dissimilar from other models in their respective categories. Finally, there are models that most naturally fall in *two* categories, or on their boundary, in such a way that one is forced to make choices about how to classify or individuate models that obscure how they work. These problem cases cast doubt on whether the categories in the book do what

they need to at the epistemic level of analysis.²

The best way to proceed from here is simply to discuss examples, and to show how these examples stretch the usefulness of the three part taxonomy to the breaking point. We will begin with an example from evolutionary game theory, which is used in economics, evolutionary biology, and ecology. One of the principal tools of evolutionary game theory is the so-called *replicator dynamics*. These dynamics are primarily used to model the evolution of strategic behavior in populations of humans and other organisms. The basic rule of these dynamics says that behaviors in a population that are more successful than the population average will become proportionally more common, while behaviors that are less successful will die out.

There are two ways of characterizing the replicator dynamics. One is the *continuous time* dynamics, which is formulated as systems of differential equations that determine a smooth trajectory of population change through a state space. On Weisberg's account, this is a paradigm case of a mathematical model. A second form is the *discrete time* replicator dynamics. Here one considers populations changing in discrete time steps, where the dynamics determine how the state of the population at one time step will change at the next time step. This form of the dynamics is best construed as an algorithm or a procedure, and thus counts as a computational model for Weisberg.

There are deep similarities, however, between the two dynamics just described. The continuous time version of the dynamics can be derived from the discrete time version via a limiting process as the time step approaches zero. They also tend to give (provably) similar results for the evolution of a population. Both produce either trajectories (discrete or smooth) through a state space, or long-term steady states which vary with initial state, depending on the analysis performed. For these reasons, the two sets of dynamics are often used to approximate one another. Many modelers use these two dynamics interchangeably. For many problems, either will do, and little hangs on the choice, besides the practical issues related to which will be easier to analyze. Furthermore, it is clear that the representational content of the two sorts of models in these cases is essentially identical.

What is important about the example is that the taxonomy seems to force, or at least emphasize, distinctions that, although they track real facts about different versions of this model, actually obscure the most important aspects about how the models are *used*. Models with discrete and continuous time replicator dynamics often play essentially identical roles in inquiry, meaning that for philosophers interested in an epistemic level account of modeling, these models should not be differentiated.³

²This is not to say that we think a different taxonomy—say, one that draws yet more distinctions—is the right one. We are skeptical of the idea of a “right” taxonomy at all. Conversely, we do not think that generalizations should never be drawn, and that one should always focus on individual cases. What we do think is that generalizations need to be sensitive to the epistemic aims of the generalizer, and that what counts as “saliently similar” will often vary with context and interest.

³In a sense, Weisberg anticipates this sort of concern. On the one hand, he briefly suggests—but does not elaborate on—the possibility of “hybrid” models that somehow bridge the different categories. It is hard to see how this can be accommodated, however, given how he defines the models in each category. Perhaps more promising, he also allows that procedures may be understood as a variety of mathematical structure, so that computational models may be understood as a special sort of mathematical model. If one were to adopt this view, one could argue that there really is just one type of model when one considers the discrete time and continuous time replicator dynamics, perhaps with different descriptions. But Weisberg does not

This first example illustrates how the taxonomy forces distinctions between models that, for epistemic purposes, are not distinct. We will now turn to examples of models within one category that are as importantly different from each other, from the point of view of modeling practice and epistemic purpose, as they are from models in other categories. In particular, the models that fall under the category of “mathematical” are strikingly diverse.

The Lotka Volterra equations, remember, serve as Weisberg’s key example of a mathematical model. Like the replicator dynamics, these equations define trajectories through a state-space, parametrized by time, that are meant to represent how real-world populations should be expected to evolve over time given some initial state. Weisberg explicitly denies that all mathematical models involve trajectories through a state-space, but, for his examples, he does focus on mathematical models of this form.

Consider, instead, the models one encounters in the theory of general relativity. These are known as *relativistic space-times*. They consist in a smooth, four dimensional manifold whose points represent events in space and time. Further structure encodes geometrical relationships between these events, such as elapsed duration or spatial distance as determined by some observer. Although these models satisfy a differential equation known as Einstein’s equation, it is of a different character from the ones discussed above. Importantly, in general one cannot think of this differential equation as characterizing trajectories of (complete) instantaneous states through some state space. Instead, one thinks of relativistic space-times as providing an intrinsically four-dimensional characterization of a possible universe.

Another example is given by the Black-Scholes formula, which one encounters in mathematical finance. The Black-Scholes formula relates a theoretically “correct” price for a certain kind of option contract, known as a European call option, to a number of other parameters characterizing the option and prevailing market conditions, such as the market price of the underlying asset. The Black-Scholes formula can be derived from a stochastic partial differential equation. But even so, one is not usually interested in the dynamical properties of solutions—that is, how price changes over time—so much as in the relationships between the various parameters at a *fixed* time. The reason is that these relationships allow traders to extract otherwise unobservable information from a market at a time by studying readily observable information.

Yet another example is the (normal form) *game*, in a game theoretic sense. This model represents a particular strategic interaction between actors and consists in a list of players, information, strategies, and payoffs, plus some further structure determining, for example, how players’ payoffs are related to the strategies they take. Games are the sorts of things evolutionary dynamics may be applied to, but without these dynamics they are significantly different from the evolutionary models described above. There is no change in states at all. Typical outputs are combinations of strategies that are, in a precise sense, stable, known as *equilibria*. The notions of equilibrium and stability here are not dynamical, though, because the model does not have the capacity to represent change over time.

All of these examples are mathematical models. But what do they have in common with one another, or with the Lotka-Volterra model? They all use mathematics in some

adopt this view. The reason is that although a procedure may be understood as a kind of mathematical structure, described by discrete mathematics, he believes that because the step-by-step process is what does the representational work, these sorts of models are sufficiently different from other mathematical models to merit their own category.

essential way, but in each case the mathematical tools are different, the roles those tools play in the reasoning are different, the ways in which they represent the world are different, the methods of analysis are different, the inputs, outputs, and intermediate manipulations are different, and so on. It is difficult to see what is gained by grouping them all together. More, the ways in which they differ from one another are precisely the sorts of ways in which a computational model might differ from the Lotka-Volterra model, or any of these other mathematical models just described. To flag computational models as having a special status in the account, without similarly distinguishing dynamical systems models from what might be called “geometrical models”, or distinguishing geometrical models from what might be called “static-constraint models,” and so on, is arbitrary.

We should emphasize that these concerns are not a quibble about whether computational models are really a sub-type of mathematical models. Similar concerns arise with physical models, vis a vis the other two categories. Consider, for instance, a three dimensional model of DNA or a projection of stars onto the inner dome of a planetarium. These are, presumably, physical models: they are physical objects, whose physical properties have the capacity to stand in certain representational relationships. The physical properties doing the representational work are geometrical, and what is represented are geometrical relationships of relative size, distance, etc. But these are just the sorts of relationships represented by a relativistic space-time. Indeed, the geometrical character of the relationships encoded in a relativistic space-time lend themselves to characterization using pictures, known as space-time diagrams. One can even consider three dimensional figures of this sort. Conversely, the double helix of the DNA model may be thought of as a manifold with further structure, representing the distance between nucleotides. So is a double helix just an alternate description of a mathematical model? Or is the mathematics just a way of characterizing certain physical models? Is there a compelling reason to distinguish a space-time diagram from the corresponding Lorentzian manifold, and say one is a physical model and the other a mathematical model? More importantly, how do these questions help us reason about the uses of models in science?

Consider another class of physical models that Weisberg describes: model organisms. Model organisms are organisms such as mice, voles, yeasts, rhesus monkeys, and so on, which are usually carefully bred from pure lines and which can then be the basis for controlled experimentation. For instance, one might intervene by changing individual genes to see the effects on development, or one might change developmental conditions to see how it affects adult behaviors, or one might expose the organisms to hormones or suspected carcinogens or potential medications to study the effects.

There are certainly ways in which model organisms are similar to a scale model like the Bay-Delta model: in a sense, one is taking advantage of natural processes that would be far too complicated to implement any other way in order to learn about the likely processes of other systems. But the sorts of processes are very different from one another. Indeed, in many cases the interventions one is interested in studying in model organisms concern chemical pathways in the body. We often think of these pathways as consisting in a sequence of states, with transformations from one state to the next occurring in time steps—in other words, as procedures. Of course, these are procedures implemented in vivo, rather than on a computer. But it is still the procedure—and our attempts to manipulate it—that does the important work, at least in some cases. So are these computational models? If not, then

is the Schelling model, when actually implemented with checkers on a checker board, still a computational model, or does it become a physical model?

So physical models may bear important epistemic similarities to both mathematical *and* computational models. It also bears emphasizing that the various dynamical physical models just mentioned—the Bay-Delta model and model organisms—are importantly different from the essentially *static* geometrical models we find in the planetarium and DNA cases. Different, arguably, to the same degree that the various mathematical models are different from one another.

The battery of examples just presented may seem like overkill for making our first point, concerning whether Weisberg’s taxonomy succeeds at the epistemic level. But we raise them because it is by reflecting on examples such as these, as well as the many examples in the book, that we are led to our second point: ultimately, “scientific models” is simply not a fruitful unit of analysis, at the epistemic level or any other. To work at this level of abstraction forces one to group together models so dissimilar in terms of their structure, their function, their interpretation, their role in practice, and so on, that one is left either making claims that cannot really apply to everything in the category, or else with generalities that reveal very little.

This point is clearest when we consider what Weisberg *does* manage to say when he steps back from the various categories discussed, and asks what all of the models considered have in common. The answer given is that a model consists in a *structure* with an *interpretation*. Now, this basic picture has a venerable history. It is precisely the definition of a model that one inherits from first order model theory. There, however, the terms are controlled: by “structure,” one means an L -structure for some first order language L , which is a collection of sets satisfying certain language-dependent properties; an interpretation is a map that associates these sets with the terms, predicates, relations, etc. of the language, in such a way that axioms in the language can be evaluated as propositions about the sets.

In the more general setting considered here, however, both “structure” and “interpretation” have been stripped of these precise meanings and nothing new is offered in their place.⁴ And what *could* be offered? What notion of structure is sufficient for it to be the case that mathematical objects, physical objects, and procedures are all examples of structures? We do not mean to claim that a model is *not* a structure with an interpretation. Rather, the point is that a sofa can also be a structure with an interpretation, as can the word “love”, and the Battle of Hastings. One might as well have an account according to which a model is a “thing”, and leave it at that.

Admittedly, there is something unfair about this way of putting the point. Weisberg does have good reason to distinguish between a model’s “structure” and the interpretation of that structure. Emphasizing these two parts allows him to make important points concerning the relationship between the physical and mathematical objects that we use in modeling and the purposes to which we put those objects. Indeed, as noted above, one of the most convincing and important parts of the book is Weisberg’s emphasis on the role of modeler intent in giving semantic content to models. Distinguishing structure from interpretation, and arguing that typically models have both parts, helps him accomplish this.

⁴For each broad class of model he offers an analysis of their type of structure, but he does not say what structure in general is supposed to be.

A response of this sort pulls the discussion firmly back to the epistemic level, though it makes the point about structure and interpretation look a bit more modest than what the book appears to claim. Perhaps the most charitable reading of the book is one on which we always assume this sort of modesty; if so, the critical comments here should be construed as a warning against a misreading of the book as trying to accomplish more than can be hoped for. That said, it is not clear that the totalizing ambitions can be read away in every case. The most important instance of this—aside from the chapter devoted to developing the general account in the first place—occurs in chapter 8, where the weighted similarity matching account of representation appears. Here the book squarely tackles a big question: what is the representational relationship between models and the world?

Of course, this is a problem that has been framed and studied by some of the greatest philosophers of the late twentieth century. The account given here is open to some isolated criticism—though one can criticize competing proposals in similar ways, as Weisberg does convincingly. But we cannot help but feel that harping on the details, here, is beside the point. The fundamental problem—just as much for the other eminent philosophers who have written on the subject as for Weisberg himself—is that the analysis *begins* with the assumption that there is a single relationship that bears between models and the world.⁵ But why should we suppose this? The more one digests examples of modeling practices across fields, the less plausible it seems to think that the same basic relationship holds between a mouse exposed to radon gas and humans suffering from cancer, as between a relativistic space-time and the universe over the course of its entire history, or as between a bargaining game and negotiations over Iran’s nuclear program, or as between the Black-Scholes formula and traders’ expectations about market volatility. It is a socio-linguistic fact that scientists tend to use the word “model” often. But one cannot infer from this that there is a natural activity or category of practice that the term tracks.

This is not to say that studying how the many things scientists call models work, including how they represent the world, is not important for philosophy of science. But care must be given to what we can hope to learn. It seems to us that any successful analysis must focus on sets of models and modeling practice that hang together in ways relevant for the analysis at hand. Weisberg does a lot of this sort of analysis very well in the book. But he is also tempted by the final generalizing step, from saliently similar to models once and for all. And it is this that is to be resisted.

References

- Nancy Cartwright. *How the laws of physics lie*. Cambridge University Press, 1983.
- Stephen Downes. Scientific models. *Philosophy Compass*, 6(11):757–764, November 2011.
- Ronald N Giere. How models are used to represent reality. *Philosophy of science*, 71(5): 742–752, 2004.

⁵Of course, it also supposes that this relationship is one of representation—an assumption that has been a subject of recent debate (Downes, 2011).

Ronald N Giere. *Explaining science: A cognitive approach*. University of Chicago Press, 2010.

Steven Hecht Orzack and Elliott Sober. A critical assessment of levins's the strategy of model building in population biology (1966). *Quarterly Review of Biology*, pages 533–546, 1993.

Bas C Van Fraassen. Scientific representation: Paradoxes of perspective. *Analysis*, 70(3): 511–514, 2010.

Michael Weisberg. Three kinds of idealization. *The Journal of Philosophy*, pages 639–659, 2007.