1-2018

Fictionalism, Semantics, and Ontology

Gordon Purves
Sacred Heart University, purvesg@sacredheart.edu

Recommended Citation

Purves, G. (2018). Fictionalism, semantics, and ontology. *Perspectives on Science, 26*(1), 52-75. doi:10.1162/POSC_a_00267

Fictionalism, Semantics, and Ontology

Gordon Michael Purves
Sacred Heart University

This article expands upon the argument of a previous work which defended a variational account of scientific fictions. Specifically, I show that this understanding of scientific fictions can provide guidance for realist interpretations of scientific theories and models. Depending on a model's variational properties, different ontological commitments are appropriate, providing a principled way for a realist to moderate her views according to the structural properties of a given model. This reasoning is then applied the Lee-Yang theory and Kubo-Martin-Schwinger statistics, two foundational models in quantum statistical mechanics. The Lee-Yang theory is analyzed in a way that permits a robust realist interpretation, whereas KMS statistics is shown to involve a use of fictions that shields the theory from confirmation and makes it inappropriate for strongly realist interpretation, without contradicting broadly realist commitments.

1. Introduction

In a previous article (Purves 2013), I argued that some recent philosophical work on the use of fictions in science (for instance, Winsberg 1999, 2003; Teller 2009; Morgan and Morrison 1999; Giere 2010) is, while illuminating about some aspects of current scientific practice, unduly limited to cases of well-established fictions. In other words, this earlier work contended, a philosophical account of scientific fictions can do more than merely describe scientific practices, but can aid in the resolution of disputes about the proper interpretation of scientific theories and the epistemic status of some scientific models. In the present paper, I expand upon the variational account of fictions and idealizations proposed in that earlier work in order to offer an analysis of theory confirmation in a context of scientific fictions, and to discuss the implications this account of confirmation has for the interpretation of scientific theories. More specifically,

Perspectives on Science 2018, vol. 26, no. 1

doi:10.1162/POSC_a_00267

I will use this account of fictions to specify what may (or may not) be an appropriate commitment for realist interpretations of scientific models. As a result, certain "problem cases" for the realist may be disposed of as being inappropriate for a realist interpretation, or as being appropriate for at most a modest, partial, or structural realist interpretation, without challenging the realist program more generally.

To demonstrate the use of this interpretive framework, I will then apply these ideas to certain models of phase change in quantum statistical mechanics (QSM). Without question, QSM has enjoyed substantial empirical success. Nonetheless, several models in QSM may prove difficult for a realist to interpret. I will show that, granting the present account of fictions, the empirical success is not a sufficient condition for claiming that a realist must be able to provide an ontological story about the model's success. Indeed, I will show in the specific case of the Kubo-Martin-Schwinger model of the ferromagnetic-paramagnetic phase transition that the structure of the model suggests that a strict realist (and even, perhaps, a structural realist) interpretation is empirically inappropriate for the model in a way that need not make it a counterinstance for the realist (or structural realist) agenda more broadly.

2. Fictions and Monotonicity

The idea that science frequently and productively involves the use of various entities and methods termed 'fictions' is not a new one, dating back at least to suggestions of the early positivists and, most explicitly, Hans Vahinger's *The Philosophy of 'As If'*. Debate on the existence and nature of scientific fictions has recently seen a resurgence among researchers studying models and simulations. My earlier work (Purves 2013) argued in favor of an epistemically normative approach to the study of fictions to complement earlier work that focused on describing how scientists use those entities that the community of researchers agrees are fictional. Instead, I present an analysis of fictions as a particular type of false posits that can be used in the construction of a model, to be juxtaposed against idealization and approximation.[1]

In this earlier work, I proposed that we can understand fictions against the contrast class of idealizations, which may be thought of as another type of strictly false posit in a model. More specifically, I proposed that these

1. There is nothing in my account that necessarily limits fictionality to simply a single posit or element of a model, and it seems plausible to say that sets of posits and whole models may be fictions as well. Nonetheless, for simplicity I will ignore this from here on, and simply state at the outset that everything I have to say about individual posits applies equally well to sets of posits and to whole models.

two types of false posits may be differentiated by their variational properties.[2] The central idea here is that we don't just look at a single use of a single model. Rather, we should look at how a model's accuracy changes as a posit is modified in different ways. (For my purposes, a posit may be thought of as any individually variable part of a model, that is, any part of the model that can be replaced so as to produce another intelligible model. As a result, any conjunction of posits will necessarily be a posit.)

Following Laymon (1987, 1989), I characterize idealizations as those posits in a specific model, which in principle may be made more realistic[3] or replaced with more realistic alternatives that would make the model more accurate. More specifically, they are posits for which all determinately more realistic alternatives (granting that it may not be clear in some cases, either because of computational intractability or expressive opacity) produce models that are more accurate. In this way, we can make sense of an idealization of something, or an idealized representation of a part of a physical system, insofar as there is (presumably) some "most realistic" alternative to the idealization that the idealization is really an idealization of. One would commonly say, then, that the idealization is a "true enough" version of the fully un-idealized posit for some given purpose.

Fictions, on the other hand, cannot be justified as being a "true enough" representation of some part of a system. Fictions are elements of a model that work in spite of their falsehoods. Thus, fictions are those posits that are false in such a way that there is some variation or alternative that could replace it in the model that is, taken independently of the context of the particular model in question, a more realistic representation of the part of the physical system that is the prima facie representational target of both posits, but which nonetheless causes the model as a whole to become less accurate. Since fictional and idealized models should obviously be interpreted differently,

2. The relationship between the terms "fiction" and "idealization" is far from uniform. Some authors consider fictions to be a type of idealization others view idealizations as a type of fiction still others use the terms interchangeably. Nothing in particular rests on this semantic point, and so for simplicity I will treat fictions and idealizations as two different types of false posits. Similar semantic differences arise with regards to whether an entire model is a fiction or an idealization, or just a part. Since I am interested in the parts of models, my convention is to treat fictions as parts of models, though again the difference is merely semantic and nothing rests on it other than simplicity of expression.

3. It is, of course, no trivial task to specify fully what is meant by a more or less realistic posit. Part of the reason, I think, is because there is not a general sense of realism that can be applied in all situations. Rather than being diverted into a much longer discussion of various theories of approximate truth, all that need be granted for the present discussion is that there are some cases where one posit is more realistic than another, even if there are some posits that cannot be properly compared. As a result, there may be some situations where the interpretive framework present here does not fully or clearly apply.

we can see that—at least in those models for which this account can unambiguously be applied—the variational properties of a model can provide some illumination about which interpretations of a model are acceptable and which are inappropriate for the model.

An important consequence of this account is that a posit is only determinately fictional or nonfictional relative to the context of a specific model, i.e., relative to its conjunction with a specific set of other elements of a model, other posits, assumptions, theories, etc. So, a posit (or a set of posits, or a model) is fictional just in case there is a more realistic alternative that can replace it that yields a model with worse predictions, explanations, or descriptions. A posit is an idealization just in case all more realistic alternatives produce more accurate predictions. This contextuality leaves open the possibility that a posit may be fictional in one model, but that the very same posit could be nonfictional in another model, even another model that is describing many of the same physical processes, should the other assumptions being made and modelling elements adopted allow for the consistent improvability of the former fiction. We can also now see the sense in which a posit's fictionality is variational and structural: fictionality arises from how the accuracy of our predictions and descriptions changes as a posit varies across a range of possibilities, not from any intrinsic property of a particular model or element of a model, nor from any property of the community of researchers using a model. (That having been said, how a community of researches treats a model, modelling method, or model element will often be a reliable indicator of these structural properties.) Fictionality is thus a property born by a posit due to the relationship between the predictions and descriptions that can be made by a model incorporating it and those of the models that are identical (as much as possible), except that the posit in question is replaced with others that are more realistic.

It is also worth noting that this definition of a fiction is, in a certain sense, a success condition: fictions are those posits or elements of a model that are able to provide more accurate predictions than some more realistic alternatives. In most cases, then, one would need a detailed epistemology of the specific case study to determine which elements of a particular model are fictional. As a result, we will not always have access to the empirical facts necessary to directly test a posit's fictionality, especially in situations where the model is being used to study phenomena to which we have no empirical access, as is often the case. To address this fully would take us far from the focus of the present project, so I will only briefly gesture towards Winsberg's (1999, 2003, 2010) excellent work on calibration and the epistemology of simulations. The idea, in brief, is that we may relatively safely infer that a posit will produce accurate predictions and descriptions when used in a model describing systems to which we do not have empirical access if it does so in other similar models describing phenomena we can actually test (either experimentally or through the use analytic calculations). Thus, we may infer from the fact that an element of a model is successfully fictional in one context (i.e., that it produces good predictions in spite of the fact that worse predictions result from replacing it with a more realistic alternative) to a similar success in another context. Of course, such an inference is far from certain or easy—as noted above, it is entirely conceivable that something may be a fiction in one model but a nonfiction in another. This type of inference will then require an argument that parity holds between the two uses of the posit.

3. An Epistemology of Falsehoods

The fact that science uses fictions and idealizations (and the fact that they are pervasive in almost all branches of science) raises some immediate problems for a theory of confirmation. As Cartwright famously argued (Cartwright 1983, pp. 100–27), the fact that idealizations and fictions[4] are generally necessary to be able to derive any descriptions of concrete phenomena from general theories means that those theories are shielded from any confirmation or refutation of those phenomena. To see why, consider a simplified version of the classic hypothetical-deductive model of confirmation (Hempel 1945), ignoring for the moment the obvious problems with that model and focusing instead on the broad picture of confirmation. According to this model, a hypothesis, H, is confirmed by observation, O, just in case H (presumably conjoined with a description of the state of the system under consideration) entails O and O is actually observed to occur. However, Cartwright argues, if the only way to derive O from H is to conjoin H with idealizations and fictions, F, then there is no clear way in which the observation of O is a success for H. After all, if O were observed to be false, all that would logically entail is that the conjunction $(H \,\&\, F)$ is false—which we already knew to be the case. Given, then, the pervasiveness of fictions and idealizations in model building, and the fact that they are often necessary for any derivations to be computationally tractable at all, it seems to follow that the use of these false posits shield fundamental theories from confirmation or refutation. Thus, "the fundamental laws of physics do not describe true facts about reality. Rendered as descriptions of facts, they are false; amended to

4. In this early work, Cartwright does not use the language of fictions and is more focused on explanation than confirmation. However, it seems clear from her later work and from how the discussion of this problem has progressed over the ensuing decades that the discussion that follows is a fairly direct consequence of the earlier work.

be true, they lose their fundamental, explanatory force," (Cartwright 1983, p. 54). Put another way, Cartwright is arguing against a realist interpretation of fundamental laws and theories, in favor of a focus on the phenomenological laws and models that result from incorporating idealizations and fictions.

A great deal of detail and sophistication has been elided from this account, as have the vast number of replies that have been proposed. My point here is neither to provide a complete history of this discussion, nor to provide a detailed response to Cartwright. My focus, rather, will be to explore whether and when fictions and idealizations, understood variationally as presented above, shield realistic interpretations of fundamental laws and theories from confirmation, and when a realist may reasonably take positive evidence to be confirmatory. As a preliminary note, however, it should be stressed that the argument that follows is not intended as a defense of scientific realism. Instead, my goal is to show how the realist may be principled in their response to apparent counterinstances either in treating a model instrumentally, anti-realistically, phenomenalistically, or in nonetheless claiming that the evidence indirectly supports a general law or theory. The argument that follows, then, will use the present account of fictions to show how one may be a principled modest realist about fictionalized and idealized models.

Presumably, in situations where we may safely assume we are not making use of any fictional or highly idealized posits, successful predictions confirm (to some degree) the general accuracy of the model's representation of the phenomenon of interest (in whatever way one interprets "accuracy"). In doing so, the confirmation may be said to "trickle down" to confirm each of the posits in the non-idealized, non-fictional model, and from there to the theories from which they were derived or which otherwise motivated them.

The problem, as Cartwright and many others have pointed out, is that very few models are constructed in this way. Frequently, models are partially composed of fictional or idealized posits, which we obviously will not want to be confirmable by a successful prediction since we know them to be false. The account that must be offered in such situations depends crucially on what kind of falsehood is at play. Laymon's (1987, 1989) account of theory confirmation provides a blueprint for how to deal with the confirmation of (at least some) models that incorporate idealizations. The basic idea is that the non-idealized elements of a model are confirmed or disconfirmed by comparing empirical evidence to the trend of the predictions as the idealizations in the model are made more realistic by being de-idealized. If the predictions would become more accurate or more accurately detailed as idealizations are made more realistic, and if we have good reasons to suppose that this trend would continue beyond the limit of what calculations are tractable, then we may reasonably take the non-idealized elements to be confirmed to a certain degree.

A successful series of predictions thus speaks in favor of the model's overall representational accuracy—i.e., its empirical adequacy. It shows that, while the model is not an exact representation of the system under investigation, it is close and would improve if the idealized parts were de-idealized, and thus the non-idealized parts of the model may be lent confirmation by its empirical success. The confirmation, however, should not extend to all of the different elements of the model in the same way, since we don't want to say that a successful prediction indicates that an idealized posit is an accurate representation of its *prima facie* target. The clearly non-idealized (and non-fictional) elements are straightforwardly confirmed, and to a certain degree so are those theories (if any) from which they may be derived. For the false posits to be idealizations, we must have some good reasons to think that the model would become more accurate if the idealizations were removed. As a result, the general theories from which the model elements are derived are not blocked from confirmation completely, as Cartwright argued. The confirmation suggests that the theories would provide accurate predictions if applied exactly, were such applications computationally feasible.

The idealized posits themselves, on the other hand, are "confirmed" to be falsehoods that are representationally useful. This confirmation is, of course, not directly evidence of the truth of an idealization, but is evidence that the idealization will suffice in models similar to the one being tested, that is, that other similar models that incorporate this idealization may be "true enough" for analogous purposes, and thus would not shield the rest of the model from confirmation or refutation. In some situations, though, we may conclude slightly more. If we have good reason to suppose that the idealization may be incrementally made more realistic, and in doing so we will continue to get incrementally better predictions, we may, with Laymon, conclude that the posit is properly interpreted as truth-like. More importantly, we can say that the version of this posit that is at the end of this process is (to a certain extent) confirmed, even if that particular model is intractable, and that we may have confidence that the idealized posit is relevantly similar to its de-idealized cousin. In this way, confirmation operates in much the same way as the calibration of computer simulations that Winsberg discusses (1999). Since the idea here is much the same as that which Winsberg discusses at length, straightforwardly applied to Laymon's way of thinking, I will not address it any further and will, instead, move on to the confirmation of fictional models.

The introduction of fictions further complicates a characterization of the success of a prediction, and the determination of the appropriate epistemic attitudes that a confirmation may call for. Laymon's response obviously will not work here, since by definition fictions are those posits which are not improvable in the way necessary for idealized confirmation. It seems just as clear, however, that we do not want all models that incorporate what Laymon terms "non-monotonically improvable" posits to be treated on the same level. Consider, by way of example, the method of image charges in classical electrodynamics.

A simple example of the use of this method involves a single charged particle, say a positively charged ion, next to a large conductive plate. The ion will induce a charge in the conductive plate by pulling the electrons in the plate towards it (and to a much lesser extent pushing the positively charged nuclei away), and leaving the region near it on the plate with a net negative charge. The exact details of this phenomenon are incredibly complex, since in principle the electric field from the ion extends infinitely, and thus the electric field will exert a force on every charged particle in the plate. The model quickly and predictably becomes intractable when one attempts to incorporate the effects on and from each of these charged particles.

The method of image charges can step in, however, and accurately model the behavior of the electric field around the plate. This comes in two stages. First, it models the large conductive plate as though it were infinite and the charged particle as though it were a point charge. Both of these adjustments are well understood as idealizations—infinite plates are reasonably accurate models of finite plates, but are computationally much more tractable, and gradually easing off from the infinite limit by looking at very large finite plates (i.e., much larger than the one in the system under consideration), though less tractable, would produce a more accurate approximation, and similarly for giving a point charge's extension. Thus, two idealizations replace two posits of our model. Second, this idealized model is fictionalized. According to an analytic uniqueness theorem, the induced electric field of a point charge near an infinite plate is identical to that produced by a fictional particle on the opposite side of the plate. This particle, the image charge or mirror charge, is an exact duplicate of the ion, except that it has the opposite charge and its position mirrors that of the original charged particle on the opposite side of the plate.

What is interesting about this model, for present purposes, is that there is a rigorous mathematical proof showing that the image charge generates the exact same electric charge as the infinite conductive plate would have, were it computationally possible to directly calculate it, and thus the interaction between the charged particle and the plate is identical in many respects to that between the charged particle and the image charge. Moreover, were we to try to make the image charge incrementally more realistic (say, by splitting the negative charge amongst several smaller particles, and spreading them slightly across the surface, or bringing the charge closer to the face of the plate, etc.), we would immediately end up with much worse predictions, and thus the image charge fits the variational definition of a fiction outlined above. However, given that the uniqueness theorem guarantees that this model would behave identically to a highly accurate idealized model, it seems odd to say, following Cartwright, that the practical necessity of the use of a fiction here is a problem for the theory of classical electrodynamics, or that the empirical success of using it to model real physical systems is somehow not a success for that theory.

Indeed, the empirical success of this type of model seems to demonstrate the power of classical electrodynamics to predict and explain electrical phenomena. The fact that fictions and idealizations were necessary to produce a tractable phenomenological model ought not to shield the non-fictional elements from confirmation or disconfirmation, at least in this case. This is because the success of the use of fictions can readily be explained by appeal to the truth of the underlying theory, through the use of the uniqueness theorem. I will refer to this kind of fiction as an isolated fiction, since the explanation for why we should not be surprised by its success demonstrates that the fictionality is isolated from the rest of the model, which may thus be confirmed to a certain degree just as with models involving idealizations.

By the same lights, however, when an explanation of a falsehood's empirical success by appeal to the truth of the underlying theory is impossible, Cartwright's argument should convince us that a realist interpretation of the general laws and theories is problematic. When a theory can only produce accurate models through the use of fictions, and when those models' empirical success cannot be explained by appeal to the truth of the theory, then the empirical success must be taken to be indicative of a problem with the underlying theory, and we can conclude that the theory can at most be interpreted as an empirically adequate fictional representation. In Cartwright's language, the empirical success only lends credence to the phenomenological laws that the model generated, rather than to the abstract general laws that were used in the formation of the model.

This is not to say simply that we lack an explanation of the empirical success in terms of the truth of the theory, but that no explanation is possible. In other words, whenever a theory's models only work when they incorporate fictions, and when the theory in principle cannot explain why this fiction would work in (at least approximately) the same way as a non-fictional alternative, the theory cannot be judged to be true in realist terms. Of course, the claim that no explanation is possible if the theory is literally true requires a significant and nontrivial argument, as well as an account of what it means to be "literally true." As we will see below, this argument will normally come by way of proving, directly or by taking a limit, or by a parity argument or some other indirect means, that, regardless of how a non-fictional alternative posit would act in the model, it cannot act in the same way as the working fiction. This, I think, can be plausibly taken to be a property that those theories we take to inform our ontology should not have, and as such, I will refer to these fictions as disconfirming fictions, their empirical success being indicative of the falsehood of any realist interpretation of the underlying theory.

It is important to be clear exactly what is being disconfirmed when a model incorporates a disconfirming fiction. Anti-realist, instrumentalist, or phenomenalist interpretations are still available, of course, since an element of the model can only be a fiction if it is to some extent empirically adequate. What is being disconfirmed, rather, is the underlying theory's ability to be a true explanation of the phenomena being modeled, as the realist would understand it. This is not to say that we should throw out the theory altogether, any more than we should throw out an outdated but still useful theory like classical mechanics. After all, the theory was still able to be used to generate an empirically accurate fictionalized model. Rather, as we will see below when discussing a disconfirming fiction in the field of quantum statistical mechanics, the appropriate epistemic attitude in these cases is to acknowledge an open scientific question about why, in this case, a problematic theory is able to provide accurate descriptions or predictions. Depending on how this question can be answered, the false explanation that the disconfirming fiction offered may fit into a broader account of the phenomenon of interest and the role of the problematized theory in explanations of it.

The types of fictions considered thus far have all required a clear and full understanding of the model under consideration, but this will obviously not always be the case. On the one hand, it may not be obvious whether a falsehood is a fiction or an idealization. Similarly, and more interesting for present purposes, a theory may be generally intractable and so it may be uncertain whether there is a ready explanation for a falsehoods' success that appeals to the truth of the theory, or there may be no explanation for the falsehood's success at all. In the former situation, we would not be in a position to judge how the model's descriptions and predictions would fare were the fiction replaced with a more realistic posit. This could arise if we have access to enough of the variational properties of a posit to identify it as a fiction (i.e., we at least suspect that there is at least one more realistic alternative that provides worse predictions), but not enough to determine how the results would compare to those of a model that instead used an ideally realistic posit in the fiction's place.[5] The success of the predictions that are generated by these fictions cannot be separated from the fictions themselves, just as with disconfirming fictions, but not because of any determinate variational property. Rather, we cannot separate success from fictionality because of our limited understanding of the fictions' variational behavior. Since we cannot isolate these fictions, any ontological conclusions would be highly tentative, and certainly not evidentially supported. However, unlike with disconfirming fictions, we are not in a position to judge that an anti-realist interpretation is thus best, and should instead take the empirical success to indicate that the model (at least in this context) is at least empirically adequate.

Further investigation, either through the construction of better tools granting new empirical access to the relevant phenomena, the discovery of new analytic solutions to the theory's equations, or the development of better computation and simulation techniques, may yet resolve the uncertainty, but our present knowledge is insufficient to do so. The best epistemic attitude, then, is a degree of epistemic caution. Were they optimistic, a realist may tentatively investigate how the world would be if the fiction turned out to be isolated (or isolatable). If, on the other hand, they were more pessimistic, a realist may look outside the theory for an explanation for the phenomena and, possibly, an external account of why the model works, to the extent that it does, though such a search is not demanded by the model's success as is the case with disconfirming fictions. Since these falsehoods fail to recommend a particular epistemic attitude, I term them indeterminate fictions.

I have here spelled out five different interpretations that may be given for the empirical success of a model, depending on whether it includes any false posits, and what kinds they are. A model that includes no idealizations or fictions will simply be confirmed by its accurate predictions, and that confirmation will trickle down to each posit and any theories that generated it. An empirically successful model that incorporates clear idealizations permits a realist interpretation. The underlying theories and modeling assumptions are confirmed, and if the idealization may be reasonably supposed to continue to yield more accurate predictions as it is de-idealized, then so is a fully de-idealized version. A model that includes an isolated fiction may be fully confirmed, allowing a realist interpretation,

5. I am using language here of a set of posits that vary with respect to their realism continuously, though in many cases such language will make no sense. The existence of a continuum from the more realistic to the less is not necessary for the present argument, though I am here mostly concerned with asymptotic behavior, for which at least poset ordering is necessary.

though the fiction itself can at most be empirically adequate. The whole of a model that uses a disconfirming fiction, on the hand, can at most be empirically adequate, and we should be hesitant to attempt any ontological interpretation of the underlying theory. Finally, the empirical success of a model that uses an indeterminate fiction indicates that the model is at least empirically adequate, while remaining silent about whether a realist interpretation is allowable, though permitting tentative explorations of both possibilities.

This division between types of fiction may be used to analyze some scientific model-building practices which utilize what are usually thought of as highly idealized assumptions, assumptions which are particularly far from the truth. I will restrict myself for the remainder of the present work to the use of the thermodynamic limit in derivations of phase transitions, with a particular focus on the ferromagnetic-paramagnetic transition. Specifically, I will argue in the next section that, in the case of the Kubo-Martin-Schwinger theory of phase transitions, the thermodynamic limit acts as a disconfirming fiction, requiring a radical reimagining of realist interpretations of quantum statistical accounts of these phenomena. This will be true even though, in classic statistical mechanics, the thermodynamic limit is able to act as a straightforward idealization, with all of the interpretive freedom that entails.

4. Quantum Statistics and the Thermodynamic Limit

Quantum statistical mechanics (QSM) is the field that attempts to rectify the apparent disconnect between how our microscopic and macroscopic theories describe the world. Specifically, QSM is the study of emergent quantum phenomena, including solid state physics, quantum thermodynamics, and changes of state and phase. As such, just like its classical counterpart, it is an attempt to use the theory of the very small to describe certain observable macroscopic phenomena. In this section, I will briefly examine the use of a central model-building assumption known as the thermodynamic assumption or the thermodynamic limit in a few QSM models. The thermodynamic limit describes the practice of treating systems that are in fact composed of a very large but finite number ($\sim 10^{23}$) of distinct particles as though they were composed of an infinite number of particles while maintaining a finite energy density in the composite system. In many theoretical contexts, such as in classical statistical explanations of phenomena like the smoothness of pressure on a boundary surface, the limit is a fairly straightforward example of an idealization in the sense described above: it simplifies calculations that are otherwise intractable but which we have good reason to suppose would monotonically improve if it were computationally possible to ease away from the limit. However, that the thermodynamic limit constitutes an idealization in these contexts is, without some further argument, not a sufficient reason to suppose that it will work as an idealization, or even an isolated fiction, in others. Conversely, of course, that the limit is strictly false in some model is not a sufficient reason to reject it or conclude that it is of merely instrumental value.[6] In this section, I will demonstrate that the thermodynamic limit is a disconfirming fiction in at least one model in quantum statistical mechanics, and that thus QSM ought not to be given a realist interpretation – or, at least, its explanation of the phenomenon in question, the existence of multiple equilibrium states during phase transition, ought not to be given a realist interpretation; and we have some reason to be skeptical of realist interpretations of other QSM models.

Much of the motivation for the thermodynamic assumption in this context comes from the difficulty of modeling phenomena such as phase transitions as quantum systems. Consider, by way of example, the transition from the ferromagnetic state to the paramagnetic state of iron (for detailed explanations of this phenomenon, see Ruelle 2004, pp. 1–22; Sewell 2002, pp. 15–8; for a more philosophical discussion, see Reutsche 2006). Below 771°C, iron (like other ferromagnetic materials below their own Curie temperatures) behaves in exactly the way we are all familiar, with a magnetic polarization. At higher temperatures, however, this polarization is washed out by thermodynamic effects. The shift from the lower temperature region to the higher temperature region corresponds to the shift from the ferromagnetic state to the paramagnetic state. This phenomenon has been well known and mathematically described for over a hundred years; nevertheless, as it turns out, providing a robust theoretical explanation for it has proven more difficult.

On a microscopic (and somewhat simplistic) level, the two phases can be fairly intuitively understood. Each atom in a piece of iron has a non-zero net magnetism (it has an unfilled electron shell, so the dipole magnetic moments of the up spin electrons do not exactly cancel out that of the down spin electrons). Each atom exerts a force on its neighbors pushing it to anti-align, just as two magnets next to each other will push their negative ends to be next to their neighbors' positive ends. At fairly low temperatures and relatively high densities (such as room temperature and pressure), this tendency will be dominated by a phenomenon called

6. Robinson (1994, pp. 84–9), for example, has argued in the context of the second and more interesting case study considered here that the thermodynamic assumption is false in ordinary quantum mechanics, and along with the fact that ordinary quantum mechanics is very highly confirmed he concludes that the assumption is illegitimate. My point here is that, though I agree with most of Robinson's conclusions, some further argument is necessary to demonstrate that the falsehood in fact closes off the possibility of answering foundational or ontological questions.

the exchange interaction. The basic idea of this interaction is that, when the orbitals of valence electrons overlap in a molecule, the lowest energy state of the two-electron system occurs when they are “farthest apart.”[7] Because of the magnetic push, this means that the system is in the lowest energy state when the dipoles are aligned, with positive ends next to positive ends and negative ends next to negative ends. The result is the ferromagnetic state: most of the valence electrons, and thus the atoms they are a part of, have aligned dipole moments, and so the piece of iron as a whole has a nontrivial net magnetic dipole moment and the whole piece of iron is magnetic. If the temperature is significantly increased, however, thermal collisions between the atoms will increase, allowing the valence electrons to rise to higher energy states. Consequently, the macroscopically more normal tendency of dipoles to anti-align (as they do in non-ferromagnetic substances) with near neighbors will again become dominant, and so the overall net dipole moment will move to very near zero, and the iron enters the paramagnetic state.

While this description fairly accurately captures the differences between the two equilibrium states and their respective microscopic causes, it fails to explain the sharpness of the threshold between them, or the details of a transition between a ferromagnetic state and a paramagnetic state, specifically the point where the two phases meet and overlap. The thermodynamic limit arises in attempts to deal with both of these problems: the sharpness of the threshold and macroscopic details of the point of contact between the two equilibrium states. I will only briefly address the statistical theory that deals with the former, the Lee-Yang theory, in order to introduce some key concepts, and then will focus on Kubo-Martin-Schwinger statistics, which addresses the fact that, during a phase change, two distinct equilibrium states can coexist.

Thermodynamics dictates that a pointlike phase change can only occur when there is a discontinuity in the free energy of the system, or one of its derivatives (for a good introduction to the Lee-Yang theory, see Blythe and Evans 2003). As such, a statistical explanation of the threshold phase change must begin by giving a reductive redefinition of the free energy of a system in terms of the properties of its microscopic constituents. The most natural way to do this is by making the free energy a function of the partition function. For instance, the Helmholtz free energy is given by:

$$A = -k_B T \ln Z$$

7. This is, of course, at best metaphorical, since the exchange interaction is a purely quantum mechanical phenomenon. More literally, the expectation value of the positions of the electrons are farther apart.

where k_B is Boltzmann's constant, T is the temperature, and the partition function, Z, is a sum over the microstates with energy E_r:

$$Z = \sum \exp \frac{-E_r}{k_B T}.$$

Since Z is a finite sum of analytic functions, and the free energy is an analytic function of Z, the free energy will also be analytic so neither it nor its derivatives can have any discontinuities. As such, if this were the end of the story, no phase transition would be theoretically possible because, again, phase changes can only occur when there is a discontinuity in the free energy or one of its derivatives. If we take the thermodynamic limit, however, then the sum becomes infinite and the partition function need not be analytic, since an infinite sum of analytic functions is not necessarily analytic. The Lee-Yang theory demonstrates that, in this limit the partition function, Z, can and does have discontinuities, and thus permits phase transitions.

The Lee-Yang theory demonstrates that threshold phase transitions can exist in the thermodynamic limit, and can differentiate between threshold transitions and non-threshold transitions, but it does not directly account for the fact that for a system to transition from one distinct phase to another, the two unique phases must, and empirically do, coexist at the transition point, regardless of whether the transition is sharp or gradual.[8] To understand how the thermodynamic limit arises in response to this problem, we must first address the statistical mechanical characterization of an equilibrium state, what is called a Gibbs state. A Gibbs state is a probability distribution over the microstates of a system that remains unchanged under all possible dynamic evolutions of the constituent microsystems. If a Gibbs state exists for a finite system at a particular temperature, therefore, it is necessarily unique. The basic idea is easy to grasp, and a detailed proof would distract from the topic at hand and is addressed more than adequately elsewhere (for the classic form of this argument, see Ruelle 2004). If a system is in a Gibbs state, then this says that for all possible dynamical transformations, the system will remain in the same state. Since

8. It is important to note that this problem cannot simply be solved by recognizing that a real bar of iron at 771°C will have some portions that are slightly above that point—and thus will be paramagnetic—and some that will be below—and thus ferromagnetic. The reason this cannot be the answer is twofold. First, coexistence is necessary for there to be a transition from one phase to the other, so saying that, in the transition of a large macroscopic system there are some parts that undergo the transition before others doesn't explain how the transition occurs. Second, even if as an empirical matter we cannot have a macroscopic system with a completely homogenous temperature, as a theoretical matter there is a determinate temperature for each part of the system, and so some will be at the threshold where the two phases must be allowed.

the set of possible dynamics is the set of possible paths through the microscopic state space, this implies that every point in the state space accessible from our initial state has a probability distribution equivalent to (or vanishingly similar to) that of the Gibbs state it started in. Thus, what it is to be an equilibrium state of a closed finite system (i.e., an equilibrium state of a system with set parameters such as temperature, pressure, volume, etc.) requires that no other equilibrium state be coexistent. In other words, if a system is in one Gibbs state that means that no other Gibbs state is accessible to it, since, by definition, if it were then either it would have the same probability distribution (and thus would be the same Gibbs state), or it would have a different probability distribution (in which case it was not in a Gibbs state to begin with).

While this description suffices for the pure paramagnetic phase and the pure ferromagnetic phase, it also illustrates why the theory of Gibbs states cannot describe the phase transition between the two. Indeed, the theory of Gibbs states seems to say that no phase transitions are possible, since a phase transition just is a transition from one equilibrium state to another, which was just demonstrated to be impossible. Equivalently, the theory of Gibbs states cannot hold at the phase transition point (771°C), at which both paramagnetic and ferromagnetic equilibrium states are possible.[9] (If this seems strange, consider a glass of ice water – it is at a temperature of 0°C, the transition point between solid and liquid water, where both solid and liquid equilibrium states thus are simultaneously present.)

If we shift to an infinite system, however, it turns out that the plurality of equilibrium states needed to account for the phenomenon of phase transition becomes possible, though, since Gibbs states are only defined for finite systems, a new theory of equilibrium states will be needed. That need can be filled by the theory of Kubo-Martin-Schwinger (KMS) states, a more general analogue to the theory of Gibbs states. Without getting overly distracted by the technical details, KMS states are unitarily equivalent probability density matrices (as opposed to probability distributions over phase space), and equilibrium states are identified with extremal disjoint KMS states that remain unchanged up to unitary equivalence under all dynamical evolutions. For finite systems, KMS equilibrium states are equivalent to Gibbs states, and thus also fail to offer any explanation of the phenomenon of phase transition. If, however, we transition to

9. In point of fact, there is not an exact transition point where both equilibrium states coexist; rather there is a small range centered around 771°C within which both equilibrium states coexist. The point is simply that there is an observed phenomenon—the coexistence of two equilibrium states—that is in need of explanation where none can be provided by the theory of Gibbs states. The apparent conflict between this range and the Lee-Yang theory discussed above will be briefly explored at the end of this work.

the thermodynamic limit (possible for the KMS theory, though not for the Gibbs theory), then superpositions of extremal disjoint KMS states become possible, and the phenomenon of coexisting equilibrium states at a phase transition point seems to be saved. As a result, "this idealization [i.e., the thermodynamic limit] is necessary because only infinite systems exhibit sharp phase transitions" (Ruelle 2004, p. 1).

To illustrate why this is the case, consider a KMS description of the possible states of a ferromagnetic system. Let us begin by idealizing (in the strict sense defined earlier) in the following ways. First, we will consider only the interactions between the valance electrons, as, through the exchange interaction, they are the greatest determinant of the magnetic moment of the system. Second, we will suppose that the iron atoms are only at the nodes of a three dimensional lattice, ignoring the slight variations in density. Finally, assume that each of the atoms are describable by a representation of the canonical commutation relations (CCR) for spin-half particles with the added condition that the spins of different sites on the lattice intercommute. (This additional condition is not an idealization as I have defined it, since the failure of atoms to intercommute is exactly what drives the exchange interaction and, thus, ferromagnetic behavior. This is unproblematic because we are only using it to be able to construct a representation of a particular state of the system as a whole, and not its dynamics. In the language outlined above, this posit is an isolated fiction, and thus doesn't affect the interpretation of the rest of the model.)

Initially restricting ourselves just to finite systems, consider two different representations, $S^{(+)}$ and $S^{(-)}$, of this lattice of spin-half particles (Reutsche 2006). $S^{(+)}$ begins with every particle in the z-spin-up state and then switches a finite number of them into the z-spin-down state. If the system is right at the phase change temperature, then the "push" towards aligned dipoles by way of the exchange interaction will be exactly balanced by the amount of thermal energy available to raise the electrons to the more compact anti-aligned dipole arrangement. As such, since the system is finite, any possible configuration of spin-up and spin-down states of the constituents is possible and thus any possible net magnetization could emerge. Analogously, $S^{(-)}$ begins with every particle in the spin-down state and then switches a finite number to spin-up, and can similarly account for every possible micro- and macro-state. What is important here is that these are equivalent representations. They are both representations because the algebraic properties (captured by the CCR and the additional intercommutation condition) are readily met by any given basis of $S^{(+)}$ or $S^{(-)}$. More importantly, they are unitarily equivalent because all of the physical properties of a system in $S^{(+)}$ could be had by a system in $S^{(-)}$, which is trivially true because the two representations span the same states. Again, however, neither of these representations can capture the physics of the change of state that occurs at 771°C, since there will always be a determinate magnetization at any given temperature.

If we now take the two representations to the thermodynamic limit, we find that the representations become unitarily inequivalent. $S^{(+)}$ is now an infinite lattice of spin-half particles, all but a finite number of which are spin-up in the z-direction, and conversely for $S^{(-)}$. They are both still representations, for it is still a simple matter to construct operators for which the CCR and intercommutation condition are still obeyed, so the algebraic properties of the system are still captured. However, it is just as simple a matter to show that not all of the same physical properties are possible for them, most relevantly that every state in $S^{(+)}$ has a net magnetization in the positive z-direction and every state in $S^{(-)}$ has a net magnetization in the negative z-direction. There thus is not a unitary map between the two representations, and they are not unitarily equivalent.

The final relevant move of this quantum statistical description of Curie's law is to recognize that these two inequivalent representations correspond to two equilibrium KMS states. A probability density matrix may be extracted in the usual way from a description of any state under one or the other representations, and in the case of infinite systems they will be extremal and, as we've just seen, disjointed. $S^{(+)}$ and $S^{(-)}$ are thus two equilibrium states, namely the two possible z-polarizations in the ferromagnetic state. The other possible ferromagnetic equilibrium states could be represented by simply picking a different axis and following the argument to this point through. Similarly, paramagnetic states could be represented by, e.g., beginning with every other constituent spin-up and the rest spin down, and then varying a finite number of them. Finally, these equilibrium states can co-exist at the same temperature, so long as they are thermodynamically allowed, because superpositions of thermodynamically distinct equilibrium states are allowed in the thermodynamic limit. Shy of the thermodynamic limit, all allowable superpositions are between unitarily equivalent representations of the same equilibrium state. To put things more simply, only in an infinite system can the paramagnetic and ferromagnetic phases coexist at the transition point.

5. Quantum Statistical Fictions

So far, so good—we have a model that can adequately capture the empirical fact that, during a phase transition, two different equilibrium states co-exist. The open question is, given the appeal to the thermodynamic limit, what have we actually demonstrated? This question is especially poignant because this kind of appeal to the thermodynamic limit is commonplace in QSM, but the limit is clearly a falsehood—no physical systems in fact have an infinite number of atoms or molecules. Given the earlier analysis of the different kinds of falsehoods, ultimately what needs to be determined is whether, in this case (and cases similar to it in relevant ways) the thermodynamic limit is an idealization, an isolated fiction, a disconfirming fiction, or an indeterminate fiction.

If the thermodynamic limit were a straightforward idealization, then we could explain its success in modeling phase change by showing that easing away from the limit leads to slightly more accurate results (though, presumably, at the cost of tractability, so the demonstration may not come in the form of actually constructing models shy of the thermodynamic limit, but instead through arguments by analogy, parity, or demonstrating the approximation analytically). Derivations of the fundamental thermodynamic relation, $dU = TdS - PdV$, from statistical mechanical principles, work in exactly this way. Without going into the details, the relation may be derived exactly from a microscopic definition of entropy only if we allow the number of particles to increase to infinity, but a very similar, and slightly more accurate result follows from more realistic specification of the number of particles. In such cases, we are free to proffer a loosely realist interpretation of the model, concluding that the world is similar to the model, and that if we were to shift away from the thermodynamic limit, our model would become more and more similar to the world. In the case of the fundamental thermodynamic relation, this means that we are free to take its empirical success to lend credence to the statistical mechanical explanation of thermodynamic laws.

In the derivation of phase transitions in quantum statistical mechanics, however, the thermodynamic limit behaves importantly differently. If we attempt to reverse the limit, we do not have good reason to think that we will get similar but more accurate (if less tractable) results. Indeed, we don't get similar results at all, since it is only in the limit that two different equilibrium states can coexist, and thus it is only in the limit that phase transitions are possible. As a result, we cannot think of the use of the thermodynamic limit in KMS models of phase transitions as a conceptual approximation—it must be a kind of fiction: an isolated fiction, a disconfirming fiction, or an indeterminate fiction.

First, consider whether we should think of the thermodynamic limit as an isolated fiction in KMS models of phase transition. Recall that isolated fictions are picked out by the ability of the underlying theory to give us a good reason for supposing that a model that uses the fiction will provide results that are very similar (and thus similarly accurate) to what would be provided by a model that replaced that fiction with an accurate description of that part of the world. We saw this kind of fiction with the method of image charges: classical electromagnetism can prove that the image charge produces exactly the same voltage and magnetic and electric fields as the actual induced charge distribution (within the idealization of an infinite plate and a point charge). However, for the same reason that the thermodynamic limit is not acting as an idealization in KMS models of phase transitions, it cannot be an isolated fiction. No finite system—whether large or small, accurate or inaccurate—can have multiple coexisting equilibrium states; only in the limit is the phenomenon saved. We thus have very good theoretical reasons for thinking that replacing the thermodynamic limit with a realistic number of particles will not produce similar predictions and descriptions of ferromagnetic systems.

Finally, the fiction is not simply an indeterminate fiction. It is not the case that we lack a full understanding of its variational properties. We know full well that the model cannot account for observed phenomena without an appeal to the thermodynamic limit. As such, the fiction must be disconfirming: the fiction plays an essential role in the empirical success of the model. But, as I argued earlier, a true (or approximately true) model (as understood by the realist), ought not to have the property of only providing accurate descriptions when incorporating a falsehood. As such, the KMS description of ferromagnetic phase transitions (and by parity other phase transitions as well) does not permit any straightforwardly realist interpretations.

This is not necessarily to say that it can only be thought of in instrumentalist terms, as a prediction generator. It does imply that we ought to be very hesitant to draw any ontological or foundational conclusions from its empirical success, but this should not be taken as an edict to re-interpret that empirical success as a theoretical failure. A more nuanced epistemic attitude may be called for, where the nature of the empirical success is indicative of the falsehood (in realist terms) of the KMS statistical explanation of phase transition, but does not suggest that that theory is to be rejected in its entirety or thought of as no more than an approximate generator of phenomenological laws. The fiction is disconfirming, in a limited sense, but not necessarily refuting. Rather, it raises a new scientific question: why is the theory of KMS statistics able to provide an explanation of phase transitions even though it does so only by assuming the falsehood of the thermodynamic limit? This question leaves room for KMS statistics, or some conceptual cousin of it, to play a role in the ultimate true explanation of phase transition, and also provides a starting point for the investigation into that explanation. In addition to driving further research by the apparent paradox of its disconfirmatory empirical success, that success remains a virtue in its own right, and so it can be used as a placeholder in other investigations that must make use of an account of phases, but which themselves are not directly concerned with the nature of phase change. In other words, the necessity of the appeal to a disconfirming fiction makes the whole of the theory of KMS statistics fictional in the non-pejorative truth-conducive sense that I have been using it, and as such it may be fruitfully used in other models, so long as care is made to ensure that it plays the role of an isolated fiction in those models. As such, KMS statistics can play the role of a locally anti-realist theory of phase transition, a theory that does not match either the physical reality involved or the mathematical structure of that reality, but which nonetheless can still be used in good science if one is careful.

Before concluding, it will be worthwhile to return briefly to the use of the thermodynamic limit in the Lee-Yang theory. Recall that this theory explains the fact that phase transitions seem to occur at thresholds, and thus, according to thermodynamics, require a non-analyticity in the free energy or one of its derivatives. Just like KMS statistics, however, this non-analyticity can only arise in the thermodynamic limit. We thus have the exact same problem with the use of the thermodynamic limit as we had with KMS statistics: only infinite systems can exhibit the kinds of phase transitions that the Lee-Yang theory is trying to explain, and so the limit is again a disconfirming fiction. However, unlike with KMS statistics, we have another avenue available other than rejecting the ontology of the underlying statistical theory: we can reject the phenomenon to be explained. Unlike the coexistence of two different equilibrium states at the transition point, the empirical basis for the claim that the transition occurs at a point is far less certain. Indeed, there is some empirical support that, slightly above 771°C, there are some regions that demonstrate ferromagnetic properties mixed into the generally paramagnetic system, and conversely for slightly below 771°C. Of course, why the transition occurs in a region surrounding 771°C rather than some other region is still in need of explanation, but we should not be surprised if that explanation does not come from finding non-analyticities in the free energy. It is beyond the scope of the present paper to delve deeply into what kind of an explanation would work, other than to note that Callender (2001) has made some progress in a similar direction, arguing that phase transitions occur around points where there would be a discontinuity in the free energy or one of its derivatives if the system were infinite. More work is needed to show why this in fact explains the phase transition, but let it suffice as an example of the kind of project that could fill the gap that I have argued is left by the interpretive problems surrounding the Lee-Yang theory.

6. Fictional Epistemic Attitudes

I have argued that the use of the thermodynamic limit constitutes a disconfirming fiction in quantum statistical accounts of phase transitions, implying that it cannot be used in a true foundational explanation of that phenomenon and cannot directly contribute to our understanding of the connection between microphysics and macrophysics. I would like to conclude now by pointing to two lights at the end of this rather bleak tunnel (at least, bleak for the realist). Specifically, I will note that this anti-realism should be viewed in a highly limited way, which, all things considered, is not so strange a conclusion to reach when discussing the limits of quantum mechanics. Moreover, the nature of the foundational failure of the thermodynamic limit can, in fact, shed some light on what type of theoretical description may not be so ontologically barren.

Quantum mechanics, perhaps more than any other scientific theory in history, has been plagued by its refusal to bend to the intuitions we have gleaned from everyday experiences. Quantum systems frequently cannot be thought of in the same way we think about those objects with which we are more intimately familiar. It is in part for this very reason that there is an open project of attempting to interpret the quantum formalism. I have here argued that such an interpretive project is a lost cause within the subfield of quantum statistical mechanics,[10] that the nature of its success is such that it must sacrifice a claim to be a description of the inner workings of our universe, if we are to understand those claims in an ontologically robust way. As I hope this discussion has established, however, this does not necessarily mean that QSM cannot play a role in good scientific practice, merely that its use is highly limited and that care is needed whenever it is used.

No amount of argument, no matter how subtle, can deny that the success of QSM is striking. Indeed, that QSM is able to so substantially connect the microworld with the macroworld is surely part of what has motivated so many practitioners to conclude that at least some of the foundational problems that have plagued ordinary quantum mechanics may be solved by moving further down these paths. The ability of a quantum system to treat macroscopic systems as though they were infinite may suggest that it is best to find a different quantum theory in which macroscopic systems really are infinite. As luck would have it, just such a theory is readily available in the guise of quantum field theory. In quantum field theory, operators are defined at geometric points on a field, and thus there is a strict infinity of property-bearing points. The present analysis may be taken to suggest, then, that, though the success of QSM fails to lend credence to a quantum statistical foundation of emergent phenomena, it may

10. At least insofar as it uses the thermodynamic limit as it does in accounts of phase transition. Such appeals, however, are endemic to QSM, and very frequently fall afoul of the argument of this paper, *mutatis mutandis*.

point in the direction of a quantum field theoretic foundation of the very same phenomena. Of course, this is far from a trivial project, in part because the thermodynamic limit now must face the substantial task of renormalization with its well-known interpretive challenges, and it is thus far beyond the scope of this paper. I mention it merely to indicate a final, positive role that disconfirming fictions may have: their failure to be fully confirmed by their success can suggest that perhaps they may be conjoined with a different theoretical framework within which they are less fictional.

References

Blythe, R. A., and M. R. Evans. 2003. "The Lee-Yang Theory of Equilibrium and Nonequilibrium Phase Transitions." *Brazilian Journal of Physics* 33: 464–473.

Callender, C. 2001. "Taking Thermodynamics Too Seriously." *Studies in the History and Philosophy of Science Part B: Studies in the History and Philosophy of Modern Physics* 32: 539–553.

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

Giere, R. 2010. "An Agent-Based Conception of Models and Scientific Representation." *Synthese* 272: 269–281.

Hempel, C. 1945. "Studies in the Logic of Confirmation." *Mind* 54: 1–26.

Laymon, R. 1987. "Using Scott Domains to Explicate the Notions of Approximate and Idealized Data." *Philosophy of Science* 54: 194–221.

Laymon, R. 1989. "Cartwright and the Lying Laws of Physics." *The Journal of Philosophy* 86: 353–372.

Morgan, M. S., and M. Morrison. 1999. *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge: Cambridge University Press.

Purves, G. M. 2013. "Finding Truth in Fictions: Identifying Non-Fictions in Imaginary Cracks." *Synthese* 190: 235–251.

Reutsche, L. 2006. "Johnny's So Long at the Ferromagnet." *Philosophy of Science* 73: 473–486.

Robinson, D. 1994. "Can Superselection Rules Solve the Measurement Problem?" *The British Journal for the Philosophy of Science* 45: 79–93.

Ruelle, D. 2004. *Thermodynamic Formalism* (2nd ed.). Cambridge: Cambridge University Press.

Sewell, G. L. 2002. *Quantum Mechanics and Its Emergent Macrophysics*. Princeton: Princeton University Press.

Teller, P. 2009. "Fictions, Fictionalization and Truth in Science." Pp. 235–247 in *Fictions in Science: Philosophical Essays in on Modeling and Idealization*. Edited by M. Suarez and A. Fine. New York: Routledge.

Winsberg, E. 1999. "Sanctioning Models: An Epistemology of Simulation." *Science in Context* 12: 275–292.

Winsberg, E. 2003. "Simulated Experiments: Methodology for a Virtual World." *Philosophy of Science* 70: 105–125.

Winsberg, E. 2010. *Science in the Age of Computer Simulation*. Chicago: University of Chicago Press.