Philosophy Compass 4/5 (2009): 835–845, 10.1111/j.1747-9991.2009.00236.x

Computer Simulation and the Philosophy of Science

Eric Winsberg*
University of South Florida

Abstract
There are a variety of topics in the philosophy of science that need to be rethought, in varying degrees, after one pays careful attention to the ways in which computer simulations are used in the sciences. There are a number of conceptual issues internal to the practice of computer simulation that can benefit from the attention of philosophers. This essay surveys some of the recent literature on simulation from the perspective of the philosophy of science and argues that philosophers have a lot to learn by paying closer attention to the practice of simulation.

1. Introduction

Computer simulation was pioneered as a scientific tool in meteorology and nuclear physics in the period directly following World War II, and since then has become indispensible in a growing number of disciplines. The list of sciences that make extensive use of computer simulation has grown to include astrophysics, materials science, engineering, fluid mechanics, climate science, evolutionary biology, ecology, economics, decision theory, sociology, and many others. There are even a few disciplines, such as chaos theory and complexity theory, whose very existence has emerged alongside the development of the computational models they study. Until relatively recently, however, philosophers of science have paid little attention.[1]

It is not obvious, of course, that philosophers of science *should* have much to say about computer simulation. It is far from obvious, after all, that the use of computer simulations in the sciences raises any particularly novel philosophical problems, or that it prompts any novel philosophical insights about the nature of science. Quite to the contrary, there is a great deal of received wisdom built into the theory-centric philosophy of science of the 20th century that leads naturally to the following view: computer simulation is nothing but the application of well-established scientific theories, and it is therefore unlikely to be of interest to philosophers. For instance, while Frigg and Reiss (2009) acknowledge that computer simulations have contributed enormously to changes in sciences, they deny that they have produced any changes that are worthy of philosophical attention.

One of the main goals of this essay will be to show that this is a misconception. There are, I want to argue, a variety of topics in the philosophy of science that need to be rethought, in varying degrees, after one pays careful attention to the ways in which computer simulations are used in the sciences. There are a number of conceptual issues internal to the practice of computer simulation that can benefit from the attention of the philosophers.

But before we begin, we should say a little about what, exactly, computer simulation is. There are three useful ways of thinking about what the term means. (i) In its narrowest sense, a 'computer simulation' is an algorithm, run on a computer, which uses step-by-step methods to explore the approximate behavior of a mathematical model, usually because the model contains equations that cannot be solved analytically. (ii) More broadly, we can think of computer simulation as a comprehensive method for studying systems that are best modeled with analytically unsolvable equations. In this broader sense of the term, it refers to the entire process of choosing a model, finding a way of implementing that model in a form that can be run on a computer, studying the output of the resulting algorithm, and using this entire process to make inferences, and in turn trying to sanction those inferences, about the target system that one tries to model. Suppose, for example, our target system is the flow of water in a basin. 'Computer simulation' in the second sense would refer to the entire process of choosing a core model from fluid dynamics, transforming the continuous model equations of that model into a set of discrete equations, finding a computationally tractable and reliable method for choosing an algorithm, studying the output of that algorithm to learn what we are interested in learning about the flow, and finally arguing that we are justified in believing what we think we have learned about the system. (iii) Finally, one can consider seriously the fact that the term computer simulation is a compound noun, and define the idea of simulation independently of computation. On this approach, a simulation is any system that is believed or hoped to have dynamical behavior that is similar enough to some other system such that the former can be studied to learn about the latter. Any object that we study because we think or hope it is dynamically similar enough for us to learn about basins of fluid by studying it is a simulation of a basin of fluid. If the former system happens to be a digitally programmed computer, then it is a computer simulation.[2]

In what follows, I will almost always have the second definition in mind. The exception will be in Section 5, where it will be useful to think of computer simulation as a species of simulation, generally.

2. A Novel Epistemology?

I want to begin with what I take to be the central philosophical issue regarding computer simulations: their epistemology. As computer simulation methods have made their way into novel disciplines, the issue of their trustworthiness for generating new knowledge has often loomed large, especially when they have competed for attention with experiments or analytically tractable modeling methods. The relevant question is always whether or not the results of a particular computer simulation are accurate enough for their intended purpose.[3] If a simulation of the global atmosphere is being used to predict future climate, does it predict the variables we are interested into a degree of accuracy that is sufficient, for example, to meet the policy-making needs to which it is being put? If a simulation of the gasses in a star is being used to understand the star's convective structure, do we have confidence that the coherent structures in the flow, the ones that will play an explanatory role in our account of how the star functions, are being depicted accurately enough to support our confidence in the explanation? If a simulation is being used in engineering and design, are the predictions made by the simulation reliable enough to sanction a particular choice of design parameters?

Given our long-standing preoccupation with issues of confirmation, it might seem obvious that philosophers of science would have the resources to easily approach these questions. Winsberg (1999a), however, argues that when it comes to topics related to the credentialing of knowledge claims, philosophy of science has traditionally concerned itself with the justification of theories, not their application. Most simulation, on the other hand, to the extent that it makes use of the theory, tends to make use of the well-established theory. The epistemology of simulation, in other words, is rarely about testing the basic theories that may go into the simulation, and most often about establishing the credibility of the hypotheses that are, in part, the result of applications of those theories.

More specifically, there are three features of the inferences made in simulation that, in combination, are not typical of the kinds of inferences that the philosophy of science has worried about licensing. These are outlined in Winsberg (2001): the knowledge produced by computer simulations is the result of inferences that are *downward, motley,* and *autonomous*.[4] They are *downward* in just the sense discussed before: they are inferences that are drawn (in part) from high theory, down to particular features of phenomena, rather than up from observations to theory. They are *motley* in that they draw on a wide variety of sources. These include theory, but also physical insight, extensive approximations, idealizations, outright fictions, auxiliary information, and the blood, sweat, and tears of much trial and error. Think, again, of our example of the simulation of the flow of water in a basin. Such a simulation is downward in the sense that what we begin with is the theory of fluids, and what we are after is a detailed description of the dynamics of the flow. It is motley in that, in addition to the theory of fluids, many other elements are added to the mix. In the simulation of fluids, for example, physical intuition might guide the choice of boundary conditions, or impose symmetry constraints. It is also not uncommon to add 'sub-grid parameterizations', such as eddy viscosity. These are mathematical functions added to the simulation to approximately recapture features of the dynamics that are lost inside the finite-sized cells of the discretization scheme. Finally, they are *autonomous* in the sense that the knowledge produced by simulation cannot be sanctioned entirely by comparison with observation. Simulations are usually employed to study phenomena where data are sparse. In these circumstances, simulations are meant to replace experiments and observations as sources of data about the world. Simulations of fluids like the kind described before might very well be used to study the inner convective structure of stars, or to determine the distribution of pressure and wind speed inside a supercell storm. Not all of the results of such simulations can be evaluated simply by being compared with the world. If a simulation reveals a particular pattern of convective flow inside a star, we have to be able to assess the trustworthiness of that information without being able to physically probe the inside of the star to check and see whether that result is confirmed by observation. In this sense, we can speak of simulations as being autonomously sanctioned. This simply means that if a simulation is to be useful, it must carry with itself some grounds for believing in the results it produces.

These three features were meant to be offered as conditions of adequacy; for which any adequate epistemology of simulation must account.[5] Against the background of the growing use of simulation in the sciences, an adequate epistemology for the philosophy of science needs to explain the fact that simulation results and computational models are often taken to be reliable despite these three features. Winsberg (2001) argues that simulation requires a new epistemology precisely because traditional stories in philosophy of science about how knowledge claims get credentialed cannot explain them.

Frigg and Reiss (2009) have argued, in response to this claim, that these are not unique features of simulation. None of these features, they argue, are

> specific to simulation and hence the need for a new epistemology is not forced upon us *only* in the face of simulations; rather it is owed to the fact that models are more complex than traditional philosophy of science allows and that we still do not have a worked out epistemology that accounts for this. The conclusion that a different epistemology is needed *could have been reached* by studying the practice of much of science apart from computer applications (my emphasis).

True enough. There are indeed other modeling practices in the sciences, practices that do not involve computer simulation, that have similar features. But there are two things we can say by way of defending the idea that simulation calls for its own epistemology. It is interesting to note, first of all, that Frigg and Reiss (2009) are making a counterfactual claim. While the previous conclusion *could have been reached* by studying other practices, it was not. To the extent that closely related points had been raised, they were raised primarily by philosophers of economics, who were examining contexts in which there is not a lot of underlying theory to begin with.

The second point that ought to be made in response to Frigg and Reiss (2009) is this: it is certainly true that careful attention to ordinary, non-computational modeling practices, might indeed have given rise to some of these same *general* concerns. Some of the more programmatic remarks that I made about the epistemology of simulation, in other words, were also applicable to other kinds of modeling. But there are also many specific features of the epistemology of simulation that are special and local to simulation.

3. Verification and Validation

Take, for example, the distinction between *verification* and *validation* that is widely employed in discussions of the sanctioning of simulation results. Here is how the distinction is supposed to go: the *validation* of a simulation is the process of assuring that the model equations that are the basis for the simulation represent the target system correctly. Verification, on the other hand, is the process of assuring that the numerical output of the simulation, and the conclusions drawn from them, are close enough to what the solutions of the original model equations would be if we were able to write them down.

At first blush, the conceptual clarity that is afforded by keeping these two activities distinct seems clear. It would seem, moreover, that it is precisely the activity of verification that is unique to simulation, as opposed to other modeling activities where the conclusions one can draw from the model can be inferred analytically. Finally, it might seem that questions of verification are entirely mathematical questions, without philosophical meat. This is indeed what Frigg and Reiss (2009) argue. But this conclusion comes too quickly.

Indeed Winsberg (1999a and forthcoming, 2010) argue that distinction between verification and validation is too clean and simple. It is true, of course, that simulationists do their best to show that their results are as close as possible to the real solutions of the equations that form the basis of their original models. The problem is that, in practice, the models they begin with are so complex, and rely so heavily on non-linear equations, that the arguments they can offer for these sorts of conclusions are incredibly weak. When models are sufficiently complex and non-linear, it is rarely possible to offer mathematical arguments that show, with any degree of force, that verification is being achieved. What simulationists are forced to do is to focus, instead, on establishing that the *combined* effect of the models they begin with, and the computational methods they employ, provide results that are reliable enough for the purposes to which they intend to put them. If we are simulating the global climate, it is almost certain that we will not be able to establish that our results bear any mathematical relationship to the ideal model of the climate. We will have to try to show, instead, that the combined effect of our choice of model and the computational methods produce predictions that are good enough for our policy-making needs. This, of course, is hard. It requires more finesse than we would expect if we thought that the activities of verification and validation could be kept separate.

One can think of this, in a sense, as a kind of a Duhem problem. When a computational model fails to account for real data, we do not know whether to blame the underlying model or to blame the modeling assumptions used to transform the underlying model into a computationally tractable algorithm. But to the extent that the Duhem problem is a problem about falsification – about where to assign blame when things go wrong – the present problem is more than a Duhem problem. When a computational model succeeds – when it provides results that are adequate for a particular purpose – it might not in fact be because either the underlying model is ideal, or because the algorithm in question finds solutions to that underlying model. It might rather be because of what simulationists sometimes call a 'balance of approximations'. This is likely the case when a model is deliberately tailored to counterbalance what are known to be limitations in the schemas used to transform the model into an algorithm. When success is achieved in virtue of this kind of back-and-forth, trial-and-error piecemeal adjustment, it is hard to even know what it means to say that a model is separately verified and validated.

4. Inspiration from the Philosophy of Experiment

Another unique feature of the epistemology of simulation is the ease with which it can draw inspiration from the epistemology of experiment. In this section, I want to discuss a few such themes in the literature: claims about the epistemology of simulations that have drawn inspiration from work in the philosophy of experiment.

In his work on the epistemology of experiment, Alan Franklin (1986, 1989) identified a number of strategies that experimenters use to increase rational confidence in their results. Weissart (1997), Winsberg (1999b, 2003), and (Parker 2008b) have all argued for various forms of analogy between these strategies and a number of strategies available to simulationists to sanction their results. The most detailed analysis of these relationships is to be found in Parker (2008b), where she also uses these analogies to highlight weaknesses in current approaches to simulation model evaluation.

Drawing inspiration from Mayo (1996), another philosopher of the experiment, Parker (2008a) suggests a remedy to some of these shortcomings in her. In this work, Parker suggests that Mayo's error-statistical approach for understanding the traditional experiment could benefit the epistemology of simulation. Taking an error-statistical perspective, the central question of the epistemology of simulation becomes, 'What warrants our taking a computer simulation to be a *severe test* of some hypothesis about the natural world? That is, what warrants our concluding that the simulation would be unlikely to give the results that it in fact gave, if the hypothesis of interest were false?' Parker believes that too much of what passes for simulation model evaluation lacks rigor and structure because it

> consists in little more than side-by-side comparisons of simulation output and observational data, with little or no explicit argumentation concerning what, if anything, these comparisons indicate about the capacity of the model to provide evidence for specific scientific hypotheses of interest.

Drawing explicitly upon Mayo's (1996) work, she argues that what the epistemology of simulation ought to be doing, instead, is offering some account of the 'canonical errors' that can arise, as well as strategies for probing for their presence.

Where Parker (2008a) draws on the analogy between simulation and experiment to make use of Mayo's (1996) analysis of experimental error, Winsberg (2003) appeals instead to Hacking's (1983, 1988, 1992) work on the experiment. One of Hacking's central insights about experiment is captured in his slogan that 'experiments have a life of their own'. He intends, I think, to convey two things with this slogan. The first is a reaction against the unstable picture of science that comes, for example, from Kuhn. Hacking (1992: 307) suggests that experimental results can remain stable even in the face of dramatic changes in the other parts of sciences. He also intends to convey that 'that experiments are organic, develop, change, and yet retain a certain long-term development which makes us talk about repeating and replicating experiments'.

Winsberg (2003) argued that some of the techniques that simulationists use to construct their models get credentialed in much the same way that Hacking says that instruments and experimental procedures and methods do; the credentials develop over an extended period of time and become deeply tradition-bound. In Hacking's language, the techniques and sets of assumptions that simulationists use become 'self-vindicating'. Perhaps a better expression would be that they carry their own credentials. The thesis I had in mind could perhaps be seen as a flip side of Parker's (2008a). Where Parker was focusing a critical eye on the existing simulation practice, and suggesting a better set of criteria under which we could be justified in taking their results to be reliable, Winsberg (2003) works from the assumption that some of the best examples of simulation *are* warrantably taken to be reliable, and asks how this could be the case. The conclusion is that this could only be the case if some of the numerical recipes used by simulationists have this property – the property of being capable of bringing their own credentialing power to the simulation table.[6]

5. The Identity Thesis

There is a different thread in the literature about the connection between simulation and experiment and its implications for epistemology. The central idea of this thread is that experiments are the canonical entities that play a central role in warranting our belief in scientific hypotheses, and that therefore the degree to which we ought to think that simulations can also play a role in warranting such beliefs depends on the extent to which they can be identified as a kind of experiment. This is quite different from the kind of theses discussed in the last section. There, the idea was more modest; it was simply to identify analogies between simulation and experiment in order to draw on the existing philosophical resources to be found in the experimental discussions. In this the work discussed in this section, what is at stake is an identity thesis, and the assumption (among some of the participants) that the power of simulation to warrant belief rises and falls with the degree to which the identity thesis can be established. In what follows, we will call the thesis that computer simulations literally are experiments the *identity thesis*, and we will call the thesis that the degree to which simulations warrant belief in hypotheses is proportional to the similarity between simulations and experiments the *epistemological dependence thesis*.

The earliest explicit argument in favor of the epistemological dependence thesis is in (Norton and Suppe 2001).[7] According to Norton and Suppe, simulations can warrant belief precisely because they literally are experiments.[8] They have a detailed story to tell about in what sense they are experiments, and how this is all supposed to work. According to Norton and Suppe, a valid simulation is one in which certain formal relations (what they call 'realization') hold between a base model, the modeled physical system itself, and the computer running the algorithm. When the proper conditions are met, 'a simulation can be used as an instrument for probing or detecting real world phenomena. Empirical data about real phenomena are produced under conditions of experimental control'.

There are, unfortunately, two problems with this story. The first is that the formal conditions that they set out are much too strict. It is unlikely that there are very many real examples of computer simulations that meet their strict standards. Simulation is almost always a far more idealizing and approximating enterprise. The second problem is that this story leaves entirely unanswered a central and crucial question. It is certainly true that if a simulation met the strict standard of realization, we would be warranted in believing the hypotheses it supported. But we want to know when and why we are warranted in believing that realization (or a weaker relation) obtains. All of the epistemological work of Norton and Suppe's (2000) story is borne by the assumption of realization, and none by the identity claim itself.[9]

The identity thesis itself has drawn criticism from Guala (2002) and Morgan (2002). Gaula begins by dismissing what he takes to be a poor argument against it. The poor argument goes something like this: simulations are not at all like real experiments because real experiments manipulate the real-world systems that are the very target of the investigation, while simulations merely manipulate *models* of the target system. What both Gaula and Morgan correctly point out is that it is, quite generally speaking, false. It is false that real experiments manipulate their targets of interest. In fact, in both real experiments and simulations, there is a complex relationship between what is manipulated in the investigation on the one hand, and the real-world systems that are the targets of the investigation on the other. In cases of both experiment and simulation, therefore, it takes an argument of some substance to establish the 'external validity' of the investigation – to establish that what is learned about the system being manipulated is *applicable* to the system of interest.

Still, both Guala (2002) and Morgan (2002) deny the identity thesis. Drawing on the work of Simon (1969), Guala argues that simulations differ fundamentally from experiments in that the object of manipulation in an experiment bears a material similarity to the target of interest, but in a simulation, the similarity between object and target is merely formal. Interestingly, while Morgan accepts this argument against the identity thesis, she seems to hold to a version of the epistemological dependency thesis. She argues, in other words, that the difference between experiments and simulations identified by Guala implies that simulations are epistemologically inferior to real experiments – that they have intrinsically less power to warrant belief in hypotheses about the real world.

A defense of the epistemic power of simulations against Morgan's (2002) argument could come in the form of a defense of the identity thesis, or in the form of a rejection of the epistemological dependency thesis. Parker (2009) employs the former strategy, while Winsberg (2009) employs the latter. Parker and Winsberg both discuss two problems with Guala's (2002) argument against the identity thesis. The first is that the notion of material similarity here is too weak, and the second is that the notion of mere formal similarity is too vague, to do the required work. Consider, for example, the fact that it is not uncommon, in the engineering sciences, to use simulation methods to study the behavior of systems fabricated out of silicon. The engineer wants to learn about the properties of different design possibilities for a silicon device, so she develops a computational model of the device and runs a simulation of its behavior on a digital computer. There are deep material similarities between, and some of the same material causes are at work in, the central processor of the computer and the silicon device being studied. On Gaula's distinction, this should mark this as an example of a real experiment, but that seems wrong. The peculiarities of this example illustrate the problem rather starkly, but the problem is in fact quite general: any two systems bear some material similarities to each other and some differences.

On the flip side, the idea that the existence of a formal similarity between two material entities could mark anything interesting is conceptually confused. Given any two sufficiently complex entities, there are many ways in which they are formally identical, not to mention similar. There are also ways in which they are formally completely different. Now, we can speak *loosely*, and say that two things bear a formal similarity, but what we really mean is that our best formal representations of the two entities have formal similarities. Winsberg (2009) argues that this is an insurmountable problem for the attack on the identity thesis that Guala (2002) and Morgan (2002) favor.

After discussing considerations like these, Parker (2009) characterizes 'simulation' and 'experiment' in such a way that computer simulation studies do qualify as experiments, and, moreover, many experiments involve simulation. Parker, in other words, accepts a very strong version of the identity thesis (Winsberg 2009), on the other hand, tries to salvage a meaningful distinction[10] between experiment and simulation in the light of the previous criticism of Guala's (2002) view. It argues that simulations fundamentally differ from experiments with regard to the background knowledge that is invoked to argue for the 'external validity' of the investigation. At the same time, it rejects the epistemological dependency thesis. The comparative epistemological power of a simulation and an experiment depends entirely on the quality of, and confidence that we have in, that background knowledge. The background knowledge, for example, that goes into a simulation of the solar system is so good that such a simulation produces more reliable knowledge than any experiment possibly could. It would seem that there are identifiable differences between ordinary experiments and simulations,[11] but there is nothing about these differences that makes one or the other intrinsically more epistemically powerful. Both the identity thesis and the epistemological dependency thesis are false.

6. *Simulations and Theories*

Not all of the philosophical issues that arise in the light of computer simulations are epistemological. Another philosophical claim that has been made about computer simulation, which is argued most forcefully by Humphreys (2004), is that computer simulations have profound implications for our understanding of the structure of theories; they reveal that both the semantic and syntactic views of scientific theories are inadequate. This claim has drawn sharp fire from Frigg and Reiss (2009) as well. They argue, correctly I think, that whether or not a model admits of analytic solution or not has no bearing on how it relates to the world. They use the example of the double pendulum to show this. Whether or not the pendulum's inner fulcrum is held fixed (a fact which will determine whether the relevant model is analytically solvable) has no bearing on the semantics of the elements of the model. From this, they conclude that the semantics of a model, or how it relates to the world, is unaffected by whether or not the model is analytically solvable. But this is perhaps not responsive to the most charitable reading of what Humphreys was pointing at.

The syntactic and semantic views of theories, after all, were not just accounts of how our abstract scientific representations relate to the world. They were also stories that had a lot to say about where the philosophically interesting action is when it comes to scientific theorizing. The syntactic theory suggested that scientific practice could be adequately rationally reconstructed by thinking of theories as axiomatic systems, and, more importantly, that logical deduction was a useful regulative ideal for thinking about how inferences from theory to the world are drawn. The semantic view of theories, on the other hand, urged that theories were non-linguistic entities. It urged philosophers not to be distracted by the contingencies of the particular form of linguistic expression a theory might be found in, say, in a particular textbook.

But computer simulations do seem to illustrate that both of these themes were misguided. It was profoundly wrong to think that logical deduction was the right tool for rationally reconstructing the process of theory application. Computer simulations show that there are methods of theory application that vastly outstrip the inferential power of logical deduction. The space of solutions, for example, that is available via logical deduction from the theory of fluids is microscopic compared with the space of applications that can be explored via computer simulation. On the flip side, computer simulations seem to reveal that, as Humphreys (2004) has urged, 'syntax matters'. It was wrong; it turns out, to suggest, as the semantic view did, that the particular linguistic form in which a scientific theory is expressed is philosophically uninteresting. The syntax of the theory's expression, it turns out, will have a deep effect on what inferences can be drawn from it, what kinds of idealizations will work well with it, etc. I would appeal again to the theory of fluids: whether we express that theory in Eulerian or Lagrangian form will deeply affect what, in practice, we can calculate and how; it will affect what idealizations, approximations, and calculational techniques will be effective and reliable in which circumstances. Hence, it seems right to me to suggest that computer simulations have revealed inadequacies with both the syntactic and semantic theories.

7. *Other Topics*

I want to conclude by highlighting two other topics that deserve philosophical attention. The first is the set of philosophical issues that arise in connection with computer simulations whose underlying models draw on inconsistent theoretical principles. Winsberg (2006a) argues that these sorts of simulations put pressure on two philosophical intuitions that are often taken for granted. The first intuition is that an inconsistent set of laws can have no models. Strictly speaking, this is of course true. But it is often assumed that the model of axiomatic logic and semantics is a sufficiently good rational reconstruction of theory application for it to follow from this that an inconsistent set of theoretical principles can give rise to no models. The examples discussed in Winsberg (2006a) show this to be false. The examples arise in the so-called multiscale simulations. These are simulations that rely, for example, on both quantum mechanics and classical molecular mechanics in order to overcome, at the same time, the computational intractability of the former and the predictive inadequacy of the latter. The second intuition is that the interesting relationships between theories at different levels of description are fully captured by the degree to which the higher-level theory is reducible to the lower-level theory in the ordinarily understood way. In fact, these same examples seem to show that it is a delicate and empirical matter how different theories will relate to each other in a successful and reliable model.

Finally, I end by simply mentioning some of the very recent philosophical work on simulations used to predict the future of the earth's climate. Work on this topic is just beginning, but it is likely to become an important area of debate in the near future. I have only the space here to list a few important pieces. For issues related to the confirmation of climate models, refer to Lloyd (forthcoming) and Parker (forthcoming). On the use of incompatible models to make multimodel ensemble predictions, see Parker (2006). On the ineliminable role of 'non-epistemic values' in climate science, see Biddle and Winsberg (forthcoming, 2010). Given the social importance of climate science, and the central role that simulations play in it, this is sure to be a topic that attracts a great deal of attention from philosophers in the coming years.

Short Biography

Eric Winsberg received his PhD in History and Philosophy of Science from Indiana University in 1999. After a postdoctoral fellowship in History and Philosophy of Science at Northwestern University, he joined the Philosophy Department Faculty at the University of South Florida in 2001, where he is now an associate professor. His principal interests are in the philosophy of science and philosophy of physics; especially in the role of computer simulations in the physical sciences, and in the foundations of statistical mechanics. Winsberg is the author of several articles on these topics that have appeared in journals such as *Philosophy of Science*, the *Journal of Philosophy*, *Studies in History and Philosophy of Modern Physics* and *Synthese*. He has recently held visiting fellowships at the Center for Interdisciplinary Studies (ZiF) at the University of Bielefeld in Germany, and the Institute of Advanced Study at the University of Durham in the United Kingdom. During that time, he wrote a monograph entitled *Science in the Age of Computer Simulation*, which will appear soon with the University of Chicago Press.

Notes

* Correspondence: 4202 E Fowler Ave, Tampa, FL 33701, USA. Email: winsberg@cas.usf.edu

[1] The earliest such work was of Humphreys (1991) and Rohrlich (1991), but genuine philosophical dialog about simulation began only in the very late 1990s.
[2] Inter alia, if it is a computer simulation then it is very likely that the reason we believe the computer has a similar dynamics to the target system is that we have programmed it following the procedure discussed in definition (ii). This is how the two seemingly very different definitions are related.
[3] Wendy Parker (2008a) has been the strongest and clearest proponent of this way of thinking about the confirmation of simulation models. She presses for this view in much of her work, but especially in this year.
[4] There are many features of science that are motley and autonomous, but it is the fact that they are downward inferences that are motley and autonomous that is remarkable.
[5] See Frigg and Reiss (2009) for a discussion of this.
[6] The argument for this claim is first developed in Winsberg (2003), but there are perhaps some better and clearer examples of the kinds of techniques I had in mind in Winsberg (2006b).
[7] Hughes (1999) argues for a form of the identity thesis, and there are strands of what look like the epistemological identity thesis, but support for the latter thesis is a best implicit. See Winsberg (2003) for a discussion of both of these pieces. One can also find an argument for the identity thesis, without epistemological implications, argued for in Humphreys (1995).
[8] To be fair, this is not exactly the dependence thesis, as it leaves open the possibility that there could be other reasons why a simulation might warrant belief. But, Norton and Suppe do not discuss other possibilities.
[9] See the discussion of Norton and Suppe's account in Winsberg (2003) for more details.
[10] In my view, it is perfectly legitimate to distinguish experiment and simulation in the way that Parker does. But there is an ambiguity in the word 'simulation' here. Parker clearly means the word simulation in the third sense that I defined it at the top of this essay. In this sense, it is correct to point out that almost all experiments involve simulation. But there is a narrower sense of simulation that is usefully distinguished from the experiment, and hence an important sense in which the identity thesis is false. See Winsberg (2009) for more details.
[11] Those differences, however, are not necessarily the obvious ones, and they require a surprising amount of subtlety to spell out.

Works Cited

Biddle, Justin, and Winsberg, Eric (forthcoming) 'Value Judgements and the Estimation of Uncertainty in Climate Modeling.' in *New Wares in Philosophy of Science*. Ed. P. D. Magnus and J. Busch. Palgrave Macmillan.

Franklin, Alan. *The Neglect of Experiment*. Cambridge: Cambridge University Press, 1986.

——. 'The Epistemology of Experiment.' *The Uses of Experiment*. Eds. D. Gooding, T. Pinch and S. Schaffer. Cambridge: Cambridge University Press, 1989. 437–60.

Frigg, Roman and Julian Reiss. 'The Philosophy of Simulation: Hot New Issues or Same Old Stew.' *Synthese: Models and Simulations* 169(3) (2009): 593–613.

Guala, Francesco. 'Models, Simulations, and Experiments.' *Model-Based Reasoning: Science, Technology, Values*. Eds. Lorenzo Magnani and Nancy Nersessian. New York: Kluwer, 2002: 59–74.

Hacking, Ian. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press, 1983.

——. 'On the Stability of the Laboratory Sciences.' *The Journal of Philosophy* 85 (1988): 507–15.

——. 'Do Thought Experiments have a Life of Their Own?' *PSA*, Vol. 2. Eds. A. Fine, M. Forbes and K. Okruhlik. East Lansing: The Philosophy of Science Association, 1992. 302–10.

Hughes, R. 'The Ising Model, Computer Simulation, and Universal Physics.' *Models as Mediators*. Eds. Mary Morgan and Margaret Morrison. Cambridge: CUP, 1999. 66–96.

Humphreys, Paul. 'Computer Simulations.' *Philosophy of Science PSA* 2 (1990): 497–506.

——. 'Computational Science and Scientific Method.' *Mind and Machines* 5 (1995): 499–512.

——. *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. Oxford: OUP, 2004.

Lloyd, Elizabeth. 'Varieties of Support and Confirmation of Climate Models.' *Proceedings of the Aristotelian Society*, forthcoming.

Mayo, D. *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press, 1996.

Morgan, Mary. 'Model Experiments and Models in Experiments.' *Model-Based Reasoning: Science, Technology, Values*. Eds. Lorenzo Magnani and Nancy Nersessian. New York: Kluwer, 2002. 41–58.

Norton, Steve and Frederick Suppe. 'Why Atmospheric Modeling is Good Science?' *Changing the Atmosphere: Expert Knowledge and Environmental Governance*. Eds. Clark Miller and Paul Edwards. Cambridge, MA: MIT Press, 2001: 67–105.

Parker, Wendy. 'Understanding Pluralism in Climate Modeling.' *Foundations of Science* 11:4 (2006): 349–68.

——. 'Does Matter Really Matter? Computer Simulations, Experiments and Materiality.' *Synthese: Models and Simulations* 169(3) (2009): 483–96.

——. 'Computer Simulation through an Error-Statistical Lens.' *Synthese* 163(3) (2008a): 371–84.

——. 'Franklin, Holmes and the Epistemology of Computer Simulation.' *International Studies in the Philosophy of Science* 22:(2) (2008b): 165–83.

——. 'Confirmation and Adequacy-for-Purpose in Climate Modeling.' *Proceedings of the Aristotelian Society*, (forthcoming).

Rohrlich, Fritz. 'Computer Simulation in the Physical Sciences.' *PSA* II (1991): 507–18.

Simon, Herbert. *The Sciences of the Artificial*. Boston: MIT Press, 1969.

Weissart, T. *The Genesis of Simulation in Dynamics*. New York: Springer-Verlag, 1997.

Winsberg, Eric. 'Sanctioning Models: The Epistemology of Simulation.' *Science in Context* (Summer) 12:02 (1999a): 275–92.

——. Simulation and the Philosophy of Science: Computationally Intensive Studies of Complex Physical Systems. PhD Dissertation. Bloomington: Indiana University, 1999b.

——. 'Simulations, Models, and Theories: Complex Physical Systems and Their Representations.' *Philosophy of Science* 68 (Proceedings) (2001): S442–54.

——. 'Simulated Experiments: Methodology for a Virtual World.' *Philosophy of Science* 70 (2003): 105–25.

——. 'Handshaking your Way to the Top: Inconsistency and Falsification in Intertheoretic Reduction.' *Philosophy of Science* 73 (2006a): 582–94.

——. 'Models of Success vs. the Success of the Models: Reliability Without Truth.' *Synthese* 152:1 (2006b): 1–19.

——. *Science in the Age of Computer Simulation Chicago: University of Chicago Press* (2010).

——. 'A Tale of Two Methods'. *Synthese: Models and Simulations* 169(3) (2009): 575–92.