

THE PHILOSOPHICAL NOVELTY OF COMPUTER SIMULATION METHODS

Paul Humphreys

pwh2a@virginia.edu

1. Introduction

Roman Frigg and Julian Reiss make two major claims in their paper 'A Critical Look at the Philosophy of Simulation'. The first is that computer simulation methods "raise few if any new philosophical problems."⁽²⁾¹. The second is that progress in understanding how simulations work would be enhanced by using results from the philosophical literature on models (passim). They are correct in their second claim. Their first claim is false and it reflects a radical failure to appreciate what is different about computational science.²

Here I shall lay out as clearly as I can some of the ways in which computational science introduces new issues into the philosophy of science. I shall also describe ways in which the traditional literature on models does not address these new issues and why, in many cases, it cannot. Because of space restrictions, I cannot address all of the issues that are relevant, but I hope that the essence of the philosophical points will be clear. I shall also focus on the criticisms that Frigg and Reiss make of my views, since the other authors mentioned in their article are

¹Numerals in parentheses refer to page numbers in Frigg and Reiss [2008].

²Stöckler [2000] argues that computer simulations do not represent a revolution in methodology. Schweber and Wächter [2000] contend that computational methods constitute a 'Hacking revolution'. I provide a different interpretation in Humphreys [unpublished]. A critical but fair review of some of the issues involved in Humphreys [2004] is Costas [forthcoming].

better positioned to defend themselves and, rather than restrict the discussion to computer simulations, I shall discuss the more inclusive area of computational science, the subject which has been the focus of my own research.

2. The Main Issue. Let me put the principal philosophical novelty of these methods in the starkest possible way: Computational science introduces new issues into the philosophy of science because it uses methods that push humans away from the centre of the epistemological enterprise. Until recently, the philosophy of science has always treated science as an activity that humans carry out and analyze. It is also humans that possess and use the knowledge produced by science. In this, the philosophy of science has followed traditional epistemology which, with a few exceptions such as the investigation of divine omniscience, has been the study of human knowledge. Locke's *Essay Concerning Human Understanding*, Berkeley's *A Treatise Concerning the Principles of Human Knowledge*, Hume's *A Treatise of Human Knowledge*, Reid's *Essays on the Intellectual Powers of Man* are but a few examples; the Cartesian and Kantian traditions in their different ways are also anthropocentric.³ In the twentieth century, the logical component of logical empiricism broke free from the psychologism of earlier centuries, but the empiricist component prevented a complete separation.⁴ Two of the great alternatives to

³A Kantian approach can be generalized to non-human conceptual categories, although the extent to which humans could understand those alien categories is then a version of one philosophical challenge faced by computational science.

⁴Carnap's *Aufbau* (Carnap [1928]) allows that a physical basis could be used as the starting point of the reconstruction procedure, but adopts personal experiences as the

logical empiricism, Quine's and Kuhn's epistemologies, are rooted in communities of human scientists and language users. Even constructive empiricism and its successor, the empirical stance, are firmly anchored in human sensory abilities. (van Fraassen [1980], [2004]). There are exceptions to this anthropocentric view, such as Popper [1972] and Ford, Glymour, and Hayes [2006], but the former's World 3 is too abstract for our concerns and the latter's artificial intelligence orientation does not address the central issues of computational science.

At this point I need to draw a distinction. Call the situation within which humans deal with science that is carried out at least in part by machines the *hybrid scenario* and the more extreme situation of a completely automated science the *automated scenario*. This distinction is important because in the hybrid scenario, one cannot completely abstract from human cognitive abilities when dealing with representational and computational issues. In the automated scenario one can, and it is for me the more interesting philosophical situation, but in the near term we shall be in the hybrid scenario and so I shall restrict myself here to that case.

For an increasing number of fields in science, an exclusively anthropocentric epistemology is no longer appropriate because there now exist superior, non-human, epistemic authorities. So we are now faced with a problem, which we can call the *anthropocentric predicament*, of how we, as humans, can understand and evaluate computationally based scientific methods that transcend our own abilities. This predicament is different from the older philosophical problem of understanding the world from a human perspective because the older problem involves representational intermediaries that are tailored to human cognitive capacities.

autopsychological basis. The overwhelming majority of the literature in the logical empiricist tradition took the human senses as the ultimate authority.

With the hybrid situation, the representational devices, which include simulations, are constructed to balance the needs of the computational tools and the human consumers. This aspect of computational science is nowhere mentioned by Frigg and Reiss yet it constitutes a major epistemological change that has significant consequences of a squarely philosophical nature for how we view theories, models, and other representational items of science.

In responding to a claim about a lack of philosophical novelty, two things have to be agreed upon. One is a criterion for what counts as philosophical; the other is a criterion for what counts as new. The only way that I can address the first issue is to draw parallels between existing issues that are generally agreed to be philosophical and those raised by computational science. Of course, this is an imperfect approach just because computational science introduces philosophical problems that may be distinctively different in kind from earlier problems, but I have no other answer to the question 'What counts as philosophical?' Concerning novelty, the reader will have to be the judge of what follows, but one has to keep in mind that nothing in philosophy is ever entirely new, in the sense that we can always trace a path of intellectual connections from a given position to earlier positions and that any serious claim in philosophy will have connections with other philosophical issues. The philosophical consequences of computer simulation methods are of course related to other developments in science and mathematics. There are issues about how computer assisted mathematics changes what is acceptable as a mathematical result, although in terms of priority the use of computer simulations antedates the use of computer assisted proofs. There is a small body of philosophical literature, mostly written by mathematicians, assessing the epistemological impact of computer assisted mathematics and automated theorem proving. (See e.g. Bailey and Borwein [2005], Thurston

1994])⁵ Scientific instruments that are quite different in their operation from optical microscopes and telescopes, such as radio telescopes and scanning tunneling microscopes, have been in use for decades. They too introduce philosophical challenges because these instruments have outputs that must be processed by other instruments, rather than having an output that is directly accessible to the human senses. Each of these philosophical issues is connected to philosophical aspects of computational science but the latter do not reduce to the former.

3. What is Philosophically New About Computational Science.

There are at least four specific philosophical issues related to the anthropocentric predicament which the study of computational science requires us to address. These are the essential epistemic opacity of most processes in computational science, the relations between computational representations and applications, the temporal dynamics of simulations, and the need to switch from in principle results to in practice considerations. Those four issues form a connected set and each will be treated in some detail below. I start with epistemic opacity.

1. Epistemic Opacity

One of the essential features of computational science, which is not mentioned by Frigg and Reiss, is the essential epistemic opacity of the computational process leading from the abstract model underlying the simulation to the output. Here a process is epistemically opaque relative to a cognitive agent X at time t just in case X does not know at t all of the epistemically

⁵It was unfortunate that the early philosophical discussions of computer assisted mathematics became distracted by focusing on the issue of whether a different sense of proof was involved. Perhaps this is why a recent 800 page reference work on the philosophy of mathematics (Shapiro [2005]) has no discussion of computer assisted mathematics.

relevant elements of the process. A process is essentially epistemically opaque to X if and only if it is impossible, given the nature of X, for X to know all of the epistemically relevant elements of the process.⁶ For a mathematical proof, one agent may consider a particular step in the proof to be an epistemically relevant part of the justification of the theorem, whereas to another, the step is sufficiently trivial to be eliminable. In the case of scientific instruments, it is a long-standing issue in the philosophy of science whether the user needs to know details of the processes between input and output in order to know that what the instruments display accurately represents a real entity.

Within the hybrid scenario, no human can examine and justify every element of the computational processes that produce the output of a computer simulation or other artifacts of computational science. This feature is novel because, prior to the 1940s, theoretical science had not been able to automate the process from theory to applications in a way that made the details of parts of that process completely inaccessible to humans.⁷ Many, perhaps all, of the features

⁶In my 2004, I used only the straightforward 'epistemically opaque' terminology. I now think that distinguishing between the weaker and stronger senses is useful. It is obviously possible to construct definitions of 'partially epistemically opaque' and 'fully epistemically opaque' which the reader can do himself or herself if so inclined. What constitutes an epistemically relevant element will depend upon the kind of process involved.

⁷I would put the two important turning points in the mid-1940s when Monte Carlo methods were first implemented on electronic computers and John Mauchly suggested that ENIAC could be used for difference equation calculations, rather than for just routine

that are special to simulations are a result of this inability of human cognitive abilities to know and understand the details of the computational process. The computations involved in most simulations are so fast and so complex that no human or group of humans can in practice reproduce or understand the processes. Although there are parallels with the switch from an individualist epistemology, within which a single scientist or mathematician can verify a procedure or a proof, to social epistemology, within which the work has to be divided between groups of scientists or mathematicians, so that no one person understands all of the process, the sources of epistemic opacity in computational science are very different.

I can illustrate some of the issues involved using as an example agent based simulations, which were only briefly discussed in my [2004]. The brevity of that discussion is a major deficiency of the book and the relative lack of attention I gave those methods can be misleading because agent based simulations are in certain ways very different from what one might call equation-based simulations. It is a common, although not universal, feature of agent based models that emergent macro-level features appear as a result of running the simulation, that these features would not appear without running the simulation, that new macro-level descriptions must be introduced to capture these features, and that the details of the process between the model and its output are inaccessible to human scientists. No traditional modeling methods address the first, second, and fourth features of these simulations. Let me elaborate a little on this point. The situation has been nicely captured by Stephen Weinberg: `After all, even

arithmetical operations. On the latter claim, see Metropolis [1993], p. 127. I do not vouch for the accuracy of Metropolis's recollections on this point although the exact historical turning point, if indeed `exact' ever makes sense in historical claims, is unimportant.

if you knew everything about water molecules and you had a computer good enough to follow how every molecule in a glass of water moved in space, all you would have would be a mountain of computer tape. How in that mountain of computer tape would you ever recognize the properties that interest you about the water, properties like vorticity, turbulence, entropy, and temperature?' (Weinberg [1987], p. 434). Many of these 'higher level' conceptual representations already exist in other theoretical representations; they are the starting point for what Ernest Nagel called inhomogeneous reductions. With other agent based models the situation is different because the simulation itself will, in some cases, construct a novel macro-level feature. These emergent patterns in computer simulations form the basis for what Mark Bedau has characterized as 'weak emergence' (Bedau 1997) and traditional human modeling techniques will not generate them from the agent base. They can only be arrived at by simulation.

2. '*Semantics*'. Philosophy of science has, as one of its concerns, how theories, models, and other representational devices are applied to real systems. One of the most common sources of frustration that scientists express about the philosophy of science is its detachment from the often brutal realities of getting theory into contact with data and how a scientific representation is applied to a real system involves considerably more than is included in the traditional semantical concerns of reference, meaning and truth. Although many philosophers will insist that providing a semantics for a simulation is necessary for it to be applied, and in the hybrid situation that is true, semantics and application are different because one can have a theory with a fully specified semantics that cannot be applied.⁸

⁸Paul Teller has argued that simulations cannot be treated as completely formal objects

Frigg and Reiss agree that the issue of application involves more than just semantics, but they almost immediately dismiss the most important aspect of how simulations are related to their applications. They say ‘In the broadest sense of application – meaning simply the entire process of using the model ... we use computational methods rather than paper and pencil to get the solutions of the equations that form part of the model. But if this is the claim, then this is just a restatement in ‘application jargon’ of the point of departure, namely that there are equations which defy analytical methods and have to be solved numerically.’ (9). The slide from ‘computational methods’ to ‘equations... that have to be solved numerically’ seriously under-describes how computational science deals with the application task. Each of the other three issues mentioned earlier – opacity, dynamics, possibility in practice – is relevant to the process of computationally applying a scientific representation to a real system.

In ([2004], section 3.12) I argued that syntax was an important element of computer simulations but that neither the syntactic account of theories nor the semantic account of theories was a suitable vehicle for representing simulations. Frigg and Reiss say that my claim is puzzling because ‘simulations by themselves do not clash with either the semantic or the syntactic view’ (10). The point is that both of these accounts of theories lack the resources to capture the essential features of computational science. The semantic account of theories, which requires that one abstract from the syntactic representation of the theory, is the wrong vehicle for capturing simulations because the specific syntactic representation used is often crucial to the solvability of the theory’s equations. Syntax is also inseparable from the dynamic implementation of

because of the problem of intentionality. (APA Pacific Division meetings, Spring 2005, unpublished talk.)

simulations. Computer simulations, because they are essentially dynamic processes taking place in time, involve the processing of linguistic strings on concrete machines. The only way to run a simulation is to run the code. Of course there are abstract representations of the algorithms involved, but it is a category mistake to claim that a collection of code in a file is a simulation. The program has the disposition to produce a simulation when run on a suitable machine but it only becomes a simulation when it is running.

However, we require more than the syntactic account of theories provides because the way in which simulations are applied to systems is different from the way in which traditional models are applied. As we described above, there is the issue of how, abstractly, the syntax is semantically mapped onto the world, on which the traditional syntactic account has a great deal to say, and then there is the different issue of how the theory is actually brought into contact with data. It is in replacing the explicitly deductive relation between the axioms and the prediction by a discrete computational process that is carried out in a real computational device that the difference lies.

Frigg and Reiss do occasionally address the specifically computational aspects of computational science, but when they do, they claim that the problems are of a mathematical nature, not philosophical. (8). The basis of their claim cannot be that mathematical results and techniques do not have philosophical consequences, because this is obviously false. The invention of non-Euclidean geometries, Gödel's incompleteness results, the initial proof of the Four Color Theorem, renormalization theory in mathematical physics, and the differences between classical and Bayesian statistical inference methods are only a few examples of how developments in mathematics can give rise to significant philosophical issues. So the claim must

be, once again, that no *new* philosophical issues are raised by technical issues in computational science.

Frigg and Reiss present three arguments against the need for a new relation between the representation and the system. The first appeals to the example of two pendulums, one a normal pendulum having a computationally tractable model and the other a double pendulum with a model that can only be solved numerically. They then ask 'Does this change [from tractable to intractable mathematics] change our understanding of how the equation relates to the world?', and answer in the negative: 'Nothing in our empirical interpretation of the terms of the equation changes in any way' (9). This response rests on accepting that the task is the narrow one of providing a semantics for the syntax, rather than the broader task of how that syntax is applied, and I have argued above that simulations connect models to applications in ways that go well beyond semantics.

The second argument presents a more serious challenge. One of the major unresolved issues in many areas of computational science is whether the invention of new mathematical techniques might eventually replace some of these computational methods. Frigg and Reiss (9) suggest that if we introduced a new class of functions that were solutions to the existing, currently intractable model, this would not change the way the model relates to the world. In fact it would, because with the availability of analytic solutions, the epistemic opacity of the relation between the model and the application would disappear. Moreover, even if this were to happen, the fact that the computational methods are, during our era, an unavoidable part of scientific method makes them of philosophical interest, just as the use of the Ptolemaic apparatus for computing planetary orbits is still of philosophical interest.

Their third argument is that, given the first two arguments, 'whether simulations make

either the syntactic or the semantic view of theories obsolete becomes a non-issue'. (9) This does not follow because there are aspects of computational science that are simply not addressed by either of the two philosophical accounts of theories. The traditional syntactic account of theories distinguished between some types of theories; those that were recursively axiomatizable, those whose axioms sets are only recursively enumerable, and a few other types. Computer scientists have since added to this classification, in moving from the simple issue of (Turing) computability to measures of theoretical computational complexity, such as P, NP, P-SPACE, and many others. This refinement can be incorporated within the syntactic account of theories. Other issues about the power of different computational architectures that are also relevant to computational science cannot be so incorporated. It is possible that if operational quantum or biological computers are built, a number of scientifically intractable problems will become tractable, opening up new areas of research. This is not an issue that is in any way addressed by traditional modeling techniques and although philosophical discussions of quantum computing have not been motivated much by issues in the area of simulations, the area is novel and is relevant to computational science. (See Mermin [2007]).

C. Temporal Dynamics

One of the features of computer simulations, although by no means the only one, that produces essential epistemic opacity, is the dynamic, temporal, nature of the computational process. Because in section 1 of their article, Frigg and Reiss set aside Stephan Hartmann's dynamic characterization of simulations, to which I subscribe in a rather different form, it is not surprising that they fail to address this issue.⁹ The issues that Frigg and Reiss bring out about the

⁹Their setting aside Hartmann's account on the grounds that 'those who put forward the

dynamics of simulations are important and this is an appropriate place to elaborate on my earlier treatment. In [2004] I distinguished between a core simulation and a full representation:

System *S* provides a core simulation of an object or process *B* just in case *S* is a concrete computational device that provides, via a temporal process, solutions to a computational model in the sense of section 3.14 [of Humphreys [2004]] that correctly represents *B*, either dynamically or statically. If in addition the computational model used by *S* correctly represents the structure of the real system *R*, then *S* provides a core simulation of system *R* with respect to *B*....In order to get from a core simulation to a full simulation, it is important that the successive solutions which have been computed in the core simulation are arranged in an [appropriate] order...the process that constitutes the simulation consists of two linked processes – the core process of calculating the solutions to the model, within which the order of calculations is important for simulations of the system but unimportant for simulations of its behavior, and the process of presenting them in a way that is either a static or a dynamic representation of the [behavior of the real system], within which the order of representation is crucial.’ ([2004], p. 110).

There are two distinct roles played by time in full simulations. The first is the temporal process involved in actually computing the consequences of the underlying model. The second is a temporal representation of the dynamical development of the system. The first process is always present in computational science. Depending on the application, the second may be

claims at issue here do not use this definition’ is startling, since section 4.1 of my 2004 is devoted to discussing and endorsing a version of Hartmann’s account that retains an essential dynamic element.

present or absent, as it will be in a static representation of the system's states. Frigg and Reiss recognize this double role played by time (10). Traditional, abstract representations of models and their relation to applications do not contain the temporal element involved in core simulations because they abstract from the implementation level. Here we can see another way in which philosophical assessments of computational science are different from the traditional modeling literature. When logicians deal with measures of computational complexity, it is how fast the number of computational steps grows relative to the input size that is important. These measures abstract from the actual time taken on a real machine and this is an appropriate abstraction for a logician. But for philosophers of science, the very nature of a prediction, as opposed to a deduction, rests on the ability to produce the result temporally in advance of the state of affairs being predicted. A model of weather dynamics would be useless for predictive purposes if any simulation based on it would take 10^6 years to run. It is because of this that in computational science technological considerations cannot be separated from philosophical considerations.

We can now see that the alternative suggestion made by Frigg and Reiss that 'the actual computational steps and the time they take to be executed somehow represent the time that the processes in the world take to unfold.' (10) fails to respect the distinction between the two roles played by time in simulations. Even with a static representation of the system's states, the core simulation will still involve a temporal dimension. There is no reason to require that the core simulation bears a mimetic relation to the system's temporal development – that is why a separate consideration of the output representation is crucial for a full simulation. The decision about the output representation is important in simulations; for example the difference between

synchronous and asynchronous updating schedules, which decide how often parts of the core simulation are used to update the system state, can make a crucial difference in the results obtained from agent based models. The principal philosophical point, however, is that the epistemic opacity, the dynamic aspects, the nature of the application process and the need to pay attention to what is possible in practice all depend on the real temporal nature of the core simulation.

D. *In Practice, Not In Principle*. A fourth novel feature of computational science is that it forces us to make a distinction between what is applicable in practice and what is applicable only in principle. Here the shift is, first from the complete abstraction from practical constraints that is characteristic of much of traditional philosophy of science, and second from the kind of bounded scientific rationality that is characteristic of the work of Simon and Wimsatt, within which the emphasis tends to be on accommodating the limitations of human agents. Ignoring implementation constraints can lead to inadvisable remarks. It is a philosophical fantasy to suggest, as Frigg and Reiss do that 'If at some time in the future we have a computer that can calculate all the states in no time at all' (11) or as Stöckler does that 'In principle, there is nothing in a simulation that could not be worked out without computers' ([2000], p. 368).¹⁰

In saying this I am not in any way suggesting that in principle results are not relevant in some areas. They clearly are; there are also other issues to which the philosophy of science needs

¹⁰ The first versions of Thomas Schelling's agent based models of segregation, and the first versions of Conway's Game of Life were done 'by hand', but almost all contemporary simulations require abilities that go so far beyond what is possible by the unaided human intellect.

to devote attention. One of the primary reasons for the rapid spread of simulations through the theoretically oriented sciences is that simulations allow theories and models to be applied in practice to a far greater variety of situations. Without access to simulation, applications are sometimes not possible; in other cases the theory can be applied only to a few stylized cases.

Within philosophy, there is a certain amount of resistance to including practical considerations, a resistance with which I can sympathize and I am by no means suggesting that the investigation of what can (or cannot) be done in principle is always inappropriate for the philosophy of science. One source of resistance to using in practice constraints is already present in the tension between descriptive history of science and normative philosophy of science, and in the tension between naturalistic approaches (which tend to mean different things to different people) and more traditional philosophy of science. But the appeal to in principle arguments involves a certain kind of idealization, and some idealizations are appropriate whereas others are not. A long-standing epistemological issue involves the limits of knowledge. Are there things that we cannot know, and if so, can we identify them? There surely cannot be any issue that this is a genuine philosophical problem. Of course it is not new – Kant famously gave us answers to the question. The question of what we can know in philosophy of science has been transformed by the rise of computational science and it is partly a question of what idealizations can legitimately be used for epistemic agents. We already have experience in what idealizations are appropriate and inappropriate for various research programmes. The move away from hyper-rational economic agents in micro-economics to less idealized agents within behavioral and experimental economics is one well-known example. For certain philosophical purposes, such as demonstrating that some kinds of knowledge is impossible even in principle, in principle

arguments are fine. But just as humans cannot in principle see atoms, neither can humans in principle

E. *Other Issues*

Frigg and Reiss claim, correctly, that much of what goes into constructing the models that lie behind many computer simulations involves issues that have been discussed in the previous philosophical literature on models. They write: 'But again, approximations, simplifications, idealisations, and isolations are part and parcel of all science and in no way specific to the use of computers in science.' True enough, but where is the list of things that *are* specific to computer simulations? These include constraints put on models by computational load issues, the problems of extending models when substantial chunks of existing code are written in legacy software, the choice of finite element decomposition, and the need for research teams to delegate substantial amounts of authority to programmers, to name just a few. Philosophers of science are free to abstract from all of these issues, but then in some areas of science their accounts will simply misrepresent how progress is made.

Even with idealizations, these computational features are relevant. Here is one particular example: Determining energy levels is a core interest for molecular chemists. Physical chemistry employs quantum mechanics as its basic theoretical apparatus, but *ab initio* calculations of the energy levels are impossible to carry out for any but the smallest molecules. The simple valence bond and molecular orbital models do not provide accurate predictions even for hydrogen molecules, so they have to be supplemented with dozens of extra terms to account for various features. They therefore employ multiple approximations and are heavily computational. So the approximations chosen in the Hartree-Fock self-consistent field approach, a central method of computational quantum chemistry, are inextricably linked with the degree to which those

calculations can actually be carried out in practice. On the other side, in 2004, I focused on the constraints that restricted computational resources place on computational science. There is now a growing sense that a different problem has arisen; that new techniques need to be developed to effectively exploit the massive computational power that is now available in many areas.¹¹

Conclusion

Frigg and Reiss attribute four claims to the contemporary philosophical literature on simulations, each of which they argue is either wrong or not new:

- a. A metaphysical claim: “Simulations create some kind of parallel world in which experiments can be conducted under more favorable conditions than in the ‘real world’”
- b. An epistemological claim: ‘Simulations demand a new epistemology’
- c. A semantic claim: “Simulations demand a new analysis of how models/theories relate to concrete phenomena”
- d. A methodological claim: ‘Simulating is a *sui generis* activity that lies ‘in between’ theorizing and experimentation.’”

I would add at least one more: A fifth aspect of simulations is that in the mathematically oriented sciences, progress is now inescapably linked to technological progress.

Frigg and Reiss’s claim about the metaphysical consequences of simulations is essentially correct and I have never subscribed to that metaphysical position myself. I have argued in this article that their second and third claims are incorrect. Computational science requires a new non-anthropocentric epistemology and a new account of how theories and models

¹¹ ‘Rationale for a Computational Science Center’, unpublished report, University of Virginia, March 2007.

are applied. These requirements are, to me, more than sufficient to justify the claim that computational science is a significantly new *sui generis* activity accompanied by new, recognizably philosophical, issues. Claims that these methods lie 'in between' theorizing and experimentation are, I believe, best interpreted metaphorically. The phrase indicates that computer simulations often use elements of theories in constructing the underlying computational models and they can be used in ways that are analogous to experiments. (For details of one of these ways, see Humphreys [1994]). Computational science has also made possible almost everything that takes place in complexity theory, itself a new area of science with its own methods that has powerful cross-disciplinary capabilities; classes dedicated to computational physics, chemistry, and biology along with textbooks dedicated to the topics have been introduced in those departments because the methods involved are different from those taught in their theory classes and in laboratory sessions.

Frigg and Reiss make many valuable points in their article; indeed, amongst the sceptical literature in this area, their arguments are the clearest that I have encountered. Nevertheless, these powerful new currents sweeping through the sciences bring with them philosophical challenges that older modeling frameworks cannot address. It is not a matter of lightly abandoning successful methods but of adapting to a different world.

References

- Bailey, David and Borwein, Jonathan [2005]: 'Future Prospects for Computer Assisted Mathematics', *Notes of the Canadian Mathematical Society* **37**, pp. 2-6.
- Bedau, Mark [1997]: 'Weak Emergence', *Philosophical Perspectives* **11**, pp. 375-399.
- Carnap, Rudolph [1928]: *Der logische Aufbau der Welt*. Berlin. English translation published as *The Logical Structure of the World*, Rolf George (translator). Berkeley: University of California

Press, 1967.

Costas, Horatio-Arlos [forthcoming]: Review of Humphreys [2004], *Mind*.

Ford, Kenneth, Clark Glymour, and Patrick Hayes [2006]: *Thinking About Android Epistemology*. Menlo Park, CA: AAAI Press.

Frigg, Roman and Julian Reiss [2008]: 'A Critical Look at the Philosophy of Simulation', *Synthese*.

Humphreys, Paul [1994]: 'Numerical Experimentation' pp. 103-118 in Paul Humphreys (ed) *Patrick Suppes: Scientific Philosopher, Volume 2. Philosophy of Physics, Theory Structure and Measurement Theory*. Dordrecht: Kluwer Academic Publishers.

Humphreys, Paul [2004]: *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. New York: Oxford University Press.

Humphreys, Paul [unpublished]: 'Computational Science and Its Effects', unpublished manuscript.

Mermin, N. David [2007]: *Quantum Computer Science*. Cambridge: Cambridge University Press.

Metropolis, Nicholas [1993]: 'The Age of Computing: A Personal Memoir', pp. 119-130 in *A New Era in Computation*, N. Metropolis and Gian-Carlo Rota (eds). Cambridge, Mass: The MIT Press.

Popper, Karl [1972]: 'Epistemology Without a Knowing Subject', pp. 106-152 in K. Popper, *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.

Redhead, Michael [1980]: 'Models in Physics', *British Journal for the Philosophy of Science* 31, pp. 145-163.

Schweber, Sam and Matthias Wächter [2000]: 'Complex Systems, Modeling and Simulation', *Studies in History and Philosophy of Modern Physics* 31, pp. 583-609.

Shapiro, Stewart (ed) [2005]: *The Oxford Handbook of Philosophy of Mathematics and Logic*. New York: Oxford University Press.

Stöckler, Manfred. [2000]: 'On Modeling and Simulations as Instruments for the Study of Complex Systems', pp.355-373 in Martin Carrier, Gerald Massey, and Laura Ruetsche (eds), *Science at Century's End: Philosophical Questions on the Progress and Limits of Science*. Pittsburgh: University of Pittsburgh Press.

Thurston, William [1994]: 'On Proof and Progress in Mathematics' *Bulletin of the American Mathematical Society* **30**, pp. 161-177.

Van Fraassen, Bas [1980]: *The Scientific Image*. Oxford: The Clarendon Press.

van Fraassen, Bas [2004]: *The Empirical Stance*. New Have: Yale University Press.

Weinberg, Stephen [1987]: 'Newtonianism, Reductionism, and the Art of Congressional Testimony', *Nature* 330, 433-437.