

The philosophy of simulation: hot new issues or same old stew?

Roman Frigg · Julian Reiss

Received: 5 December 2007 / Accepted: 17 October 2008 / Published online: 4 December 2008
© Springer Science+Business Media B.V. 2008

Abstract Computer simulations are an exciting tool that plays important roles in many scientific disciplines. This has attracted the attention of a number of philosophers of science. The main tenor in this literature is that computer simulations not only constitute interesting and powerful new *science*, but that they also raise a host of new *philosophical* issues. The protagonists in this debate claim no less than that simulations call into question our philosophical understanding of scientific ontology, the epistemology and semantics of models and theories, and the relation between experimentation and theorising, and submit that simulations demand a fundamentally new philosophy of science in many respects. The aim of this paper is to critically evaluate these claims. Our conclusion will be sober. We argue that these claims are overblown and that simulations, far from demanding a new metaphysics, epistemology, semantics and methodology, raise few if any new philosophical problems. The philosophical problems that do come up in connection with simulations are not specific to simulations and most of them are variants of problems that have been discussed in other contexts before.

Keywords Simulation · Models · Computer experiments · Representation · Epistemology of simulation

1 Introduction

Computer simulations are an exciting tool that plays important roles in many scientific disciplines, some of which owe their existence to increased computational power.¹

¹ Disciplines in which simulations are important include meteorology, nuclear physics, astrophysics, evolutionary biology, decision theory, chaos theory and parts of complexity theory.

Simulations increase the range of phenomena that are epistemically accessible to us in a way comparable to scientific instruments such as microscopes or telescopes. This is partly due to the fact that most models are specifically tailored to incorporate a particular piece of mathematics that we know how to handle. Computer simulations extend the class of tractable mathematics and thereby broaden the range of modelling tools that we can use. This extension of modelling tools leads to scientific progress because much of the success of modern science is due to what Ian Hacking called ‘calculation’ or, in contemporary parlance, ‘model building’, and computers make a huge difference to our ability to build models.

That the advent of computational techniques introduces new tools into science, opens new possibilities, and represents scientific progress is beyond dispute. Thus we endorse Humphreys’ conclusion that computational techniques ‘constitutes a significant and permanent addition to the methods of science’ (2004, p. 64).

This has attracted the attention of a number of philosophers of science. The main tenor in this literature is that computer simulations not only constitute interesting and powerful new *science*, but that they also raise a host of new *philosophical* issues. The protagonists in this debate claim no less than that simulations call into question our philosophical understanding of scientific ontology, theories and models, and submit that simulations demand a fundamentally new philosophy of science in many respects:

But physicists and engineers soon elevated the Monte Carlo above the lowly status of a mere numerical calculation scheme, it came to constitute an alternative reality—in some cases a preferred one—on which ‘experimentation’ could be conducted. Proven on what at the time was the most complex problem ever undertaken in the history of science—the design of the first hydrogen bomb—the Monte Carlo ushered physics into a place paradoxically dislocated from the traditional reality that borrowed from both experimental and theoretical domains, bound these borrowings together, and used the resulting bricolage to create a marginalized netherland that was at once nowhere and everywhere on the usual methodological map. (Galison 1996, pp. 119–120)

[...] I discuss some of the distinct epistemological characteristics that are exhibited by simulation studies and argue that these characteristics are novel to the philosophy of science. (Winsberg 2001, p. 443)

Computer simulations have a distinct epistemology [...] In other words, the techniques that simulationists use to attempt to justify simulation are unlike anything that usually passes for epistemology in the philosophy of science literature. (Winsberg 2001, p. 447)

[...] computer simulations are not simply number crunching techniques. They involve a complex chain of inferences that serve to transform theoretical structures into specific concrete knowledge of physical systems [...] this process of transformation [...] has its own unique epistemology. It is an epistemology that is unfamiliar to most philosophy of science [...]. (Winsberg 1999, p. 275)

[...] how we understand and appraise these new methods [i.e., computer simulation methods] is essentially different from the understanding and evaluation of traditional theories. (Humphreys 2004, p. 54)

Standard accounts of how theory is applied to the world have become inappropriate for these new methods [i.e., computer simulation methods] because, aside from failing to capture what is distinctive about them, they place constraints on the methods that are misguided. (Humphreys 2004, p. 57)²

[...] computer simulation provides [...] a qualitatively new and different methodology for the natural sciences, and [...] this methodology lies somewhere intermediate between traditional theoretical science and its empirical methods of experimentation and observation. In many cases it involves a *new syntax* which gradually replaces the old, and it involves *theoretical model experimentation* in a qualitatively new and interesting way. Scientific activity has thus reached a new milestone somewhat comparable to the milestones that started the empirical approach (Galileo) and the deterministic mathematical approach to dynamics (the old syntax of Newton and Laplace). Computer simulation is consequently of considerable philosophical interest. (Rohrlich 1991, p. 507, italics original)

Four claims emerge:

Metaphysical: Simulations create some kind of parallel world in which experiments can be conducted under more favourable conditions than in the ‘real world’.

Epistemic: Simulations demand a new epistemology.

Semantic: Simulations demand a new analysis of how models/theories relate to concrete phenomena.

Methodological: Simulating is a *Sui Generis* activity that lies ‘in between’ theorising and experimentation.

The aim of this paper is to critically evaluate these claims. Our conclusion will be sober. We argue that these claims are overblown and that simulations, far from demanding a new metaphysics, epistemology, semantics and methodology, raise few if any new philosophical problems. The philosophical problems that do come up in connection with simulations are not specific to simulations and most of them are variants of problems that have been discussed in other contexts before. This is not to say that simulations do not raise new problems of their own. These specific problems are, however, mostly of a mathematical or psychological, not philosophical nature. In sum, we agree that simulations introduce something new and exciting into *science*, but we doubt that this requires us to rewrite the *philosophy of science*.³

This is not tantamount to suggesting that the literature on simulation is without merit; quite to the contrary. But we see this literature as contributing to existing debates about, among others, scientific modelling, idealisation or external validity, rather than as exploring completely new and uncharted territory. The emphasis on novelty and the tendency to reason that the issues raised by simulations are completely unlike anything that philosophy of science has seen before prevents philosophers from appreciating what a discussion of simulations could contribute to ongoing debates. A focus on the continuity with existing debates would, we submit, be more beneficial than a

² Cf. also pp. 50–55 and 57–60, Humphreys (1991, p. 497), and (1995, p. 501). A somewhat ‘milder’ form of this view is expressed in Fox Keller (2003, p. 199).

³ Cf. Stöckler (2000).

forced attempt at compartmentalising the discussion. By contesting novelty claims we hope to pave the ground for an analysis of simulations that fits into the panorama of problems we are familiar with and thereby contributes to an advancement of these debates.

Before discussing the four specific claims about simulations in detail, let us briefly consider the term ‘simulation’. In everyday parlance, the term can mean quite a number of different activities or processes: we simulate a job interview by giving the candidate a mock interview, an illness by pretending to have its symptoms, work situations by using animated narrative vignettes, etc. Computer simulations are a narrower, but still vast and heterogeneous class, which contains processes as varied as that of flying a plane, the movement of people in public transport systems, strategic scenarios of warfare, and virtual reality scenarios of all kinds.⁴

Those who put forward the claims of the philosophical novelty of simulation, however, focus on using digital computers in situations involving analytically intractable equations (see [Winsberg 1999](#), pp. 275, 278; [2001](#), p. 444; [2003](#), pp. 107–108; [Humphreys 2004](#), p. 49). But even when narrowed down in this way, ‘simulation’ is used in two very different meanings in the literature:

In the *narrow* sense, ‘simulation’ refers to the use of a computer to solve an equation that we cannot solve analytically, or more generally to explore mathematical properties of equations where analytical methods fail (e.g., [Humphreys 1991](#), p. 501; [2004](#), p. 49; [Winsberg 1999](#), p. 275; [2001](#), p. 444).

In the *broad* sense, ‘simulation’ refers to the entire process of constructing, using, and justifying a model that involves analytically intractable mathematics (e.g., [Winsberg 2001](#), p. 443; [2003](#), p. 105; [Humphreys 1991](#), p. 501; [2004](#), p. 107). Following [Humphreys \(2004\)](#), pp. 102–104, we call such a model a ‘computational model’.

Since many of the claims about the philosophical significance of simulations are in fact about the construction of an analytically unsolvable model, in this paper we adopt the latter, broad meaning; the exception is Sect. 3.2 where we briefly come back to simulation in the narrow sense.

There are, of course, alternative definitions of ‘simulation’, most notably [Hartmann \(1996\)](#). On this definition, simulations are intimately connected to dynamical models, i.e., ones that include assumptions about the time evolution of a system. More specifically, a simulation amounts to solving the equations of a dynamical model, which accurately captures the time evolution of the target system; if the simulation is run on a computer it is a computer simulation. This definition enjoys the advantage of being able to distinguish simulation in a specific sense from numerical methods broadly construed (i.e., the computer assisted exploration of properties of mathematical objects in general). However, since those who put forward the claims at issue here do not use this definition (or only an altered version of it) and our criticisms do not turn on the (in)ability to distinguish between simulation in this specific sense and numerical methods broadly speaking, we will stick to the above, more widely adopted meaning in what follows.

⁴ For brief discussions of the term ‘simulation’ and its history see [Fox Keller \(2003\)](#), pp. 198–199, [Hartmann \(1996\)](#), pp. 83–84, and [Morgan \(2004\)](#).

2 Metaphysics: do simulations (and only simulations!) create ‘parallel worlds’?

Several philosophers, historians and scientists claim that simulations create ‘parallel’ (or ‘virtual’ or ‘artificial’) worlds (cf. Galison 1996, from which we quoted above, Edwards 2001, Sterman 2006). The most plausible interpretation of this idea, we think, is that the simulant investigates proximate systems that differ more or less radically from the systems he or she is ultimately interested in. This usually means that inferences about those latter systems, the ‘target systems’, are made in two analytically separable (though in practice not always thus separated) steps: first, a conclusion is established about the proximate system; second, the claim is exported from the proximate system to the target system. Because usually proximate and target systems differ from each other, and because the investigation of proximate systems is often methodologically involved itself, scientists focusing on the first inferential step can plausibly be said to investigate ‘parallel worlds’: they study one kind of system in order to ultimately learn about a different kind of system.

This is a crude but in our view entirely correct analysis is how reasoning with simulations works. What we disagree with is the further suggestion that simulations are alone or special in this respect. Whenever proximate systems are studied with the purpose of drawing inferences about differing target systems of interest ‘parallel worlds’ in the above sense are created and this kind of ‘inspection by proxy’ happens all over the place in the sciences. Most notably, studying mathematical models—whether or not analytically solvable—or in fact any kind of model used to draw inferences about natural or experimental systems involves the creation of ‘parallel worlds’ in precisely this sense (cf. Robert Sugden’s view that theoretical economic models depict ‘credible worlds’, Sugden 2000).

And the creation of ‘parallel worlds’ is not a prerogative of theoretical models. The great majority of experiments are conducted on proximate systems rather than the target systems themselves and therefore create ‘parallel worlds’ in our sense. For example, we infer from a small group of patients in a clinical trial to the general population; we experiment with mice to find out whether smoking causes cancer in humans; we examine the properties of a scale model of an aeroplane wing in a wind tunnel to learn about the properties of the aeroplane wing during a flight and so forth. In none of these cases is the proximate system identical to the target system. Importantly, proximate and target system differ by type, not just by token. That is, the problem of inferring from proximate to target system is not merely the traditional problem of induction, which concerns the inference among tokens of the same type, but rather a different, presumably more involved problem.

Moreover, many experimental systems are artificially created systems that are not found in nature. Nearly all experiments in the physical sciences create artificial environments that allow physicists to focus on the interaction of a small number of factors that cannot be found in nature in this pure form. It has also been said that few of the objects modern physics deals with can be found in ‘natural environments’ (Cartwright 1999, Part I; Sismondo 1999). Experiments on fruit flies, conducted in order to find out about the genetic constitution of organisms, are specially bred strains, grown to ensure the replicability of results. Laboratory mice, too, are standardised. When experimenting with chemical substances, standardised, pure forms are used, not those chemicals

found in nature. When economists experiment they frequently try to make sure that experimental subjects suppress all motives other than greed by giving them appropriate incentives. This too creates situations not normally found in ‘the field’ (Guala 2005, Chap. 11).

Studying proximate systems as stand-ins for target systems of interest is thus a practice that pervades all science. Simulations can rightly be said to create ‘parallel worlds’; but so can models and experiments of all kinds. The creation of ‘parallel worlds’ is therefore not a reason to believe that simulations require a new philosophy of science.

3 Epistemology: do simulations demand a new theory of knowledge?

Following our distinction between a narrow and a broad sense of ‘simulation’, let us first examine claims that have been made about the process of building a simulation, which involves the heavy use of *modelling assumptions* and which pertains only to the broad sense of ‘simulation’. After that we will look at simulation in the narrow sense, the process of solving intractable equations.

3.1 Modelling assumptions

Winsberg ascribes three features to the allegedly distinct epistemology of simulation: it is *downward*, *autonomous* and *motley*:

Downward: Simulation is about starting with theory and working your way down. This kind of epistemology is, to the philosopher of science, a curious beast. It is an epistemology that is concerned with justifying inferences from a theory to its application – an inference that most philosophy of science has assumed is deductive and consequently not in need of justification. (Winsberg 2001, p. 447)

Autonomous: Also typical of the notion of epistemology in philosophy of science is that it is founded on comparison. If you want to know if some representational structure is accurate or reliable compare it with the thing it is meant to represent. But simulation techniques are used precisely when data on the phenomena to be represented are conspicuously sparse. Because of this I describe the epistemology of simulation with term I borrow from Jeffrey Ramsey: autonomous [...] (ibid.)

Motley: Even though all simulation modeling [...] begins with theory [...] our theoretical knowledge is just one of several ingredients that simulationists use to produce their results. All of these resources and their influences need to be considered when justifying simulation results.’ (Winsberg 2001, p. 448) ‘Given limitations on computer speed and memory, these techniques [computer simulations] also invariably resort to other approximations, idealizations and even ‘falsifications’ – that is, model assumptions that directly contradict theory. (Winsberg 2003, p. 109; cf. 1999, p. 282)

Winsberg does not say explicitly what he means by ‘epistemology’ and so it is not immediately clear how to evaluate the claim that simulations have a distinct epistemology. We now discuss three different interpretations of this claim. In personal

communication Winsberg has pointed out to us what he had in mind is the third interpretation and that he does not subscribe the first two (we will discuss his views further below). However, as the first two seem to be at least *prima facie* plausible readings of the claim, it seems still worth pointing out that the first is reasonable but neither new nor specific to simulations, while the second fails to be convincing.

Following the time-honoured distinction between the context of discovery and the context of justification, let us first characterise two interpretations. We may say that, as a matter of fact, scientists reach certain conclusions by using computational models that are (1a) usually built starting from a piece of theory (though without *deducing* them from theory), (1b) not deduced from data (and background knowledge, say) and (1c) constructed using a variety of different sources and techniques including approximations, idealisations and so forth. That is, the *process of discovering a result with a computational model* is downward, autonomous and motley. Alternatively, we may say that scientists are *warranted* in believing or trusting the results of a simulation because the computational model (2a) has been constructed by using a (well-established) theory, (2b) is largely independently of direct tests of its assumptions against data and (2c) involves a variety of non-theory-driven elements, including approximations, idealisations and so forth. That is, the *process of justifying a simulation result* appeals to the fact that the computational model has been constructed in a downward, autonomous and motley way.⁵

Starting with the first reading, we think that the point is essentially correct but that these observations are neither novel nor simulation specific. (1a) That scientists start the construction process with a piece of theory but without deducing the model from it is interesting.⁶ Winsberg's example is that theory does not dictate what mathematical function one has to choose to account for the energy inside a cell when doing simulations using a grid (2001, p. 448). This is also something Humphreys points out (2004, pp. 67, 72–76, 88–95): computational models are not normally constructed by merely plugging special functions into general equations provided by well-established theories. But that models cannot be derived from theory in a mechanical way and that often elements that do not form part of the theory are used in the construction of a model is one of core claims of the 'models as mediators' project (see Morgan and Morrison 1999b, and the essays by Boumans, Suárez and Hartmann therein)⁷ and has been stated by Cartwright in the early 1980s (see in particular Cartwright 1983, Chap. 6). (1b) That computational models are constructed using empirical data without being deducible from these data is also a core concern of the 'models as mediators' project,

⁵ By simulation 'result' we mean a (usually complex) hypothesis established on the basis of the simulation, for example, 'the average global temperature will rise by 0.5°C by 2030'. In other cases the result may be a hypothesis about the causal mechanism responsible for patterns in the data.

⁶ Nevertheless we should point out that many simulations (and mathematical models!) do not take their starting point in a well-understood theory and are constructed bottom up. Cellular automata (CA), for instance, are based on relatively general hypotheses about the interaction between cells, which usually are not founded in theory (see Fox Keller 2003). Individual- or agent-based models can but don't have to be, and often are not, based on explicit theoretical considerations.

⁷ In this paper we focus on the 'models as mediators' project because this is the most recent point of reference in the models debate. Related points can also be found in the writings of Hesse, Black, Achinstein, and others; for a review see Frigg and Hartmann (2006).

where it has been argued using the case of macroeconomic models (see [Morgan and Morrison 1999a](#), 15ff. and the Chapter by van den Boogard). (1c) If simulations are autonomous of both theory and data, it is clear that other kinds of ingredients must play a part in the derivation of a simulation model, and it is correct that these ingredients form a diverse bunch (see also Winsberg's 'ad hoc modelling' in his [1999](#), p. 282). But again, approximations, simplifications, idealisations and isolations are part and parcel of many parts of science and in no way specific to the use of computers. For instance, we leave out factors all the time when we try to construct linear models and we often lump together many factors in some 'net effect' in order to obtain a simple model; so fudge factors really aren't anything simulation specific (just think of the liquid drop model of the nucleus, which we can handle with paper and pencil methods).

The more interesting and more contentious reading is that the process of justifying a simulation result appeals to the fact that the computational model on which the calculations are based has been constructed in a downward, autonomous and motley way. However, it is hard to see, at least without further qualifications, how justification could derive from construction in this way. There does not seem to be a reason to believe that the result of a simulation is credible just because it has been obtained using a downward, autonomous and motley process. In fact, there are models that satisfy these criteria and whose results are nevertheless not trustworthy. Real-business cycle (RBC) models in economics are a case in point. First, real business cycle models constitute a development of the earlier equilibrium strategy for analysing business cycles and thus all start with theory (viz., equilibrium theory). Second, for this and other reasons they are not pure models of data. Third, 'other ingredients' such as approximations, simplifications and so forth play an important role. There is a well-known trade-off between making the approximation to the analytical solution as simple as possible, in order to be easily computable, and not using assumptions that produce artefacts in relevant properties of the solution. Neither theory nor data analysis can yield information about how to trade off these desiderata best. And yet, RBC models, especially when the class is restricted to perfectly competitive representative agent models, are not generally regarded as trustworthy or reliable in their results (see for instance [Hartley et al. 1997](#)). For this reason (2a)–(2c), even suitably refined, do not provide criteria for justifying simulations.

A third and different reading of Winsberg's criteria understands downwardness, autonomy and motleyhood not as *features* that figure in a justification of a simulation result or a computational model, but rather as *conditions of adequacy* for a future, yet to be formulated, epistemology of simulation (this seems to be consonant with what he says in his [2001](#), p. 453). On that reading, Winsberg points out that such an epistemology has to explain the facts that simulation results and computational models (2a') are justified, at least in part, on the basis of the principles and techniques used in the construction of the computational model (cf. [Humphreys 2004](#), pp. 76–82), (2b') cannot be dismissed as unjustified because they cannot be compared directly with observational data and the target system, respectively, and (2c') are justified, at least in part, on the basis of the idealisations etc., that are used in the construction, even if these contradict the theory used.

It is correct, we believe, that these characteristics (probably along with others not mentioned here) have to be explained by an appropriate epistemology. But as the above

discussion of the context of discovery makes clear, at that level of generality, neither of these points is 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.⁸

This conclusion needs to be qualified in two respects. First, although the need to rethink the issue of justifying results obtained with the aid of complex models clearly arises as soon as the above mentioned features of models and model construction are recognised, the issue has not received much attention in the relevant debates.⁹ Winsberg's discussion of simulation draws our attention to this fact and highlights that what is missing from available analyses of how models work is the recognition that there is an issue about where models thus constructed get their credentials from.

Second, although the problem of justification is not specific to simulation when discussed at the general level, specific issues arise when we look into the details of computational models. Winsberg studies 'the use of computers for modeling very complex physical phenomena for which there already exist good, well-understood theories of the processes underlying the phenomena in question' (2001, p. 443), and argues—rightly, we think—that at least part of the justification of these models comes from the fact that they are based on such a theory. However, this source of credibility threatens to get lost, because in order to solve the equations of the model on a computer, these equations undergo all kinds of transformations which fundamentally change their mathematical characteristics (some of these transformations even result in so-called computationally ill-posed problems, i.e., ones of which we know that the solution never converges no matter how small a grid we choose). How, in the face of these transformations required to make them computationally tractable, do models preserve their theoretical credentials? And what reasons do we have to believe that such models, even if they work in one domain, can be successfully carried over to another domain? These are important and novel questions. The point that we would like to emphasise here is that, although these questions are novel and simulation specific, they belong (and Winsberg agrees) to the class of problems that arise in connection with complex models in general and do not representing a revolutionary departure from everything that philosophers were worried about in the past.

3.2 Equations and computation

Interestingly, the actual process of computing the solution to an equation and the methods that it involves receive only marginal attention in the philosophical literature on simulation. At first glance, this seems surprising, because it is at this point that the

⁸ Some advance in this respect has been made by recent contributions to the growing literature on experimentalism.

⁹ Noteworthy exceptions include Morton (1993) and some writers in the philosophy of economics; for instance Little (1995), Cartwright (2007), Boumans (1999), Sugden (2000) and Reiss (2007, Chap. 7).

new methods enter and this is what scientists engaging with computational science spend most of their time doing.

The central question that arises is whether the computer gives the correct solution or, more generally, informs us correctly about the properties of the solution of the equation under study. In other words, we want to know whether the simulation outcome is internally valid in the model provided by the modelling assumptions. Among the issues that arise in this connection are the following:

Approximation. How close to the actual solutions are the numerical solutions? This question is particularly pressing because the systems that we cannot solve are typically non-linear and often exhibit sensitive dependence on initial conditions (see [Smith 1998](#) for an accessible introduction to shadowing theorems, which are relevant in this context).

Truncation errors. The computer stores numbers only to a fixed number of digits. How do the resulting errors affect the result of the simulation? (cf. Humphreys, p. 118).

Reduction of dimensions. Scientists use reduction methods for high-dimensional problems to represent certain features of the solution in accessible form. One of the well-known techniques of that kind are Poincaré surfaces of section, which require an appropriate choice of section.

Interpretation. We have to cope with the fact that one loses generality and also the ability to read off general features (bifurcations etc.) from the solutions, and hence simulation results in many ways are less revealing than analytical solutions (cf. Humphreys, p. 65).

These questions are not easily answered and a good deal of the work in computational science addresses these and related concerns. But they are purely mathematical problems, which, as such, have to be solved at a formal level. Nevertheless, Winsberg argues that the issue of deriving valid results is not a mere mathematical problem:

Even though simulation is fundamentally about replacing analytical solution with calculation, which at first sight appears to be merely a mathematical transformation, the question of reliability of the results of simulation modeling goes beyond simple concerns about the reliability of the calculation, and extends to the entire simulation process and the conclusions that scientists reach when they use it. ([1999](#), p. 276; cf. [2001](#), p. 444 and [2003](#), pp. 110–111)

We should distinguish two different notions of reliability here, answering two different questions. First, are the solutions that the computer provides close enough to the actual (but unavailable) solutions to be useful? What ‘close enough to be useful’ means may well depend on contextual factors. But once the standards are set, this is a purely mathematical question and falls within the class of problems we have just mentioned. So, there is nothing new here from a philosophical point of view and the question is indeed one of number crunching (in the widest sense of the term—number crunching may involve complex mathematics). Second, do the computational models that are the basis of the simulations represent the target system correctly? That is, are the simulation results externally valid? This is a serious question, but one that is independent of the first problem, and one that equally arises in connection with models that do not

involve intractable mathematics and ordinary experiments. Moreover, the point that these need to be kept separate seems to be reinforced by the fact that computer simulations are also used in pure mathematics, where questions of the second sort do not arise at all.¹⁰ For instance, computers played a crucial role in the classification of finite simple groups and the construction of the largest of them, the so-called Fischer-Griess Monster, and they are heavily used in the investigation of fractal geometry.

But by distinguishing the justification of modelling assumptions and the derivation of a simulation result from a given set of equations, we commit exactly the sin Winsberg urges us to steer away from. Following the practitioners in the field, Winsberg calls the former process ‘validation’, the latter ‘verification’, and argues that the two are not separable in practice and can only be separated by the philosopher of science at the cost of presenting a seriously distorted picture of what happens in practice (especially in Winsberg forthcoming, Chap. 1). Although simulationists do their best to justify that the results of their computations are close to the solutions of the original equations, due to the complexities of the calculations these arguments are inevitably weak and inconclusive; in other words, verification is close to hopeless in practice. To get around this problem simulationist shift the focus of attention from validation and verification to the question of whether the *package* of the model and the calculation techniques used provides reliable information about the real-world target system. The most important way to do this is so-called benchmarking, which consists in comparing the relevant output of the package with known facts about the target (data, qualitative behaviour, known analytical results for special cases, etc.). If the two match up, then the package is deemed trustworthy. From this Winsberg concludes that the sanctioning of simulations cannot be neatly divided into verification and validation.

It is important not to confound the issues here. It is one thing to ask whether the process of validating a model relative to some target system of interest and the process of verifying a simulation result relative to a set of equations are distinct in practice, for example in that they are performed by different agents at different points in time with completely different methodologies. We agree with Winsberg that in this sense the two activities are not distinct. But it is a different thing to assert that it is not possible and worthwhile to distinguish two different questions we might ask with respect to the simulating activity.

And what is more, we think that even in practice it might be important to keep these two issues separate and to ask, on the one hand, whether a simulation result is faithful to the equations from which it is derived, and, on the other hand, whether the model used is a good (in whatever sense) representation of the target system at stake (interestingly, as also Winsberg points out, most practitioners in the field do in fact believe that these two issues should, ideally, be kept separate). Suppose, for example, a simulation fails to reproduce an observed pattern in the data. It would then be good to know whether this is due to a mistake in the calculation or due to bad modelling assumptions, because the responses would be different. Conversely, suppose the simulation does reproduce the pattern in the data. Then it would also be good to know whether the calculation is faithful to the equations used. The final aim is to apply

¹⁰ Humphreys seems to agree on this point, see pp. 110–111.

the simulation to areas for which there are no data, and if we knew that the model was both verified and validated we would have better reasons to trust the results of a simulation, than if we only knew that the net result of the package matched the facts used in benchmarking, which may well be due to it being the case that mistakes both in the model and in the calculations mutually cancel out.

But Winsberg insists that even if desirable in principle, the separation of validation and verification is impossible to achieve in practice. Where does this leave us? Put in the abstract, the problem we are facing here is that we reach a conclusion using different ingredients that we cannot test independently, because we have experimental access only to their joint net effect. This is of course a version of the Duhem problem. So the problem itself is not new. What is new, and here we agree with Winsberg, is that it occurs in the context of scientific modelling. When an analytically solvable model fails to account for the data we know that there must be something wrong with the modelling assumptions (assuming we haven't made a simple error when solving the equations). When a computational model fails to account for the data, we only know that there is something wrong either with the modelling assumption or with algorithm used to do the calculations, or both. This is a problem. If things go wrong we don't know where to put the blame, and hence we don't know where to start looking for a solution. And even if things work well, we don't know why they do (it might be that the effects of errors in the calculation and faulty modelling assumptions cancel out), which makes an extension of the domain of application of a computational model much more problematic than it is the case with analytically solvable models. In this sense there is indeed a new problem in connection with computational models as opposed to analytically solvable models: we somehow have to get how to get around the Duhem problem. This problem has no general solution and Winsberg is probably right that strategies to get around this difficulty in the context of simulations are different from strategies in other context.

4 Semantics: do simulations demand a new analysis of how models and theories relate to target systems of interest?

It has also been claimed that simulations urge us to rethink the representational function of models; i.e., the question how models relate to target systems of interest. For instance:

Standard accounts of how theory is applied to the world have become inappropriate for these new methods [i.e., computer simulation methods] because, aside from failing to capture what is distinctive about them, they place constraints on the methods that are misguided. (Humphreys 2004, p. 57)

In the broadest sense of application—meaning simply the entire process of using the model—this is of course true: 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.

However, the passage from Humphreys quoted above is immediately followed by a critical discussion of the syntactic and the semantic view of theories. Since these are philosophical accounts of what theories are and of how their semantics works, it seems clear that Humphreys wants to make a more substantial claim. We discuss two questions in this context: (a) Do computational models relate to target systems of interest in a way different from other models? (b) Does the dynamical nature of simulations make a difference to the semantics of computational models?

4.1 Do computational models relate to target systems of interest in a way different from other models?

Let us begin with an example. Consider the so-called double pendulum (Srivastava et al. 1990), a pendulum in which the bob is fixed not to a string but a hinge-like ‘arm’ consisting of two branches that are connected to each other with an axis and that enclose an angle α . If we block the hinge (and thereby keep α fixed) we have a normal pendulum whose equation is analytically solvable (even if we take friction into account). If we remove this constraint and let the angle between the two branches vary, then the equations of motion become non-integrable and solutions have to be calculated numerically. Does this change our understanding of how the equation describing the double pendulum relates to the world? We don’t think so. We still interpret α as the angle between the two branches, p as the momentum of the pendulum bob and x as its position. Nothing in our empirical interpretation of the terms of the equation changes in any way. The co-ordinates are, of course, different functions of time, but this is a mathematical not an interpretative difference.

Furthermore, if next year an ingenious mathematician introduces a new class of functions (that can be written down in the same way as sine and exponential functions) and shows that one of them is a solution to the equation of our (at present analytically insolvable) equation for the double-pendulum, this will constitute an important mathematical advance but not make us rethink how the model relates to the world.

This said, the question of whether simulations make either the syntactic or the semantic view of theories obsolete becomes a non-issue. There are many reasons to disagree with either of these views but the existence of computer simulations is not one of them. If these views can capture what is going on in, say, the normal pendulum, then they can also capture what is going on in the double-pendulum and other non-analytically solvable models.

4.2 Does the dynamical nature of simulations make a difference to the semantics of computational models?

In Chapter 4 of his book Humphreys suggest that simulations are special in two further respects: the role of time and the involvement of visualisation.

Time. Humphreys observes that simulations for the most part concern dynamical models and that time plays a special role in these simulations (p. 108).¹¹ An example

¹¹ Parenthetical references in the text of this section are to Humphreys (2004).

he discusses is the simulation of the orbit of a planet, a process consisting in ‘successive computations of the planet’s state (position and velocity) at discrete time intervals, using the mathematical equations that constitute a model of the orbital kinematics’ (p. 108). Time, according to Humphreys, is essential in two ways to this simulation. First, computer simulations are carried out in real time, even when parallel computing is used. This ‘temporal process of calculation’ (p. 109) is what makes simulations distinctive and what is essential to their understanding. Second, Humphreys emphasises the importance of the representation of the output, which he regards as ‘one of the key methodological features of simulation’ (p. 110). An output representation can be static (e.g., numerical tables or plots of a curve). The salient point now is that the outputs of simulations of dynamical processes are often dynamical themselves (‘the output is displayed dynamically via an elliptical motion’, p. 110). Humphreys explicitly admits that simulations are not always dynamical but he thinks that in the important and typical cases of modern computer simulation they are (p. 112). Let us discuss these points in turn.

The first point, according to Humphreys, ‘entails that none of the representational categories [...]—syntactic or semantic theories, models, research programmes, or paradigms—are able to capture what simulations can do’ (p. 109). This is puzzling. Simulations by themselves do not clash either with the semantic or the syntactic view. Furthermore, on any reasonable (i.e., not too narrow) understanding of models there is no problem either—after all, computations are based on a computational model, which is, by Humphreys’ own admission (p. 106), so broad a notion that even parts of antique Greek astronomy fall under that heading. Finally, research programmes and paradigms are surely flexible enough notions to accommodate simulations in one way or another.

But maybe Humphreys has a different point in mind when he emphasises the importance of time. He quotes Hartmann’s (1996, p. 83) characterisation of a simulation as a dynamical process that imitates another process. This suggests that the claim might be that the actual computational steps and the time they take to be executed somehow represent the time that processes in the world take to unfold. In other words, the claim seems to be that there is a mimetic relation between the time needed to do the calculations and the time needed for the target processes to develop. This interpretation is also supported by Humphreys’ emphasis on the fact that the computation has to be carried out on an actual machine and that a mere abstract representation of the calculation does not count (p. 109).

But if this is the claim, then it is false. First, as Humphreys admits, there are simulations of static objects (pp. 108 and 110–111) and one can (as he also admits on pp. 110–111) simulate purely mathematical objects that have no temporal component either (e.g., when we use Monte Carlo simulations to evaluate integrals). Hence these features are not features of simulations per se but only of particular kinds of simulations. Second, the fact that it takes time to calculate the successive states of the planet has nothing to do with the semantics of the states. For one, all that matters is that the computer provides states that come with a time index: the planet is in state x at time t . If at some time in the future we have a computer that can calculate all the states in no time at all, we surely don’t feel we lose anything. So the time needed to execute computation does not play any role in representing the process; what does matter is

the time index attached to the state. For another, depending on the model and the programme, calculations of temporally equidistant states (i.e., we set up a computer to calculate the planet's state at $t = 1, 2, 3, \dots$) may take different times, depending on the computational complexity at different states. But the 'real time' is the same—the steps are one time unit apart. So the time needed for calculation is not the same as, and not even necessarily proportional to, the time that elapses in the world.

Let us now turn to the second claim. Is a dynamical output representation the new distinctive feature (as Humphreys claims on p. 112)? This does not seem to be the case. A dynamical representation may (and often does) facilitate our understanding of things, but it contains exactly the same information as, say, an ordered list. It is true that we grasp things easier when presented with them in a 'dynamical' way, but this is a matter of psychology and not semantics.

Visualisation. On pages 112–114 Humphreys seems to suggest that the fact that the outputs of simulations are represented visually bears special weight. When using simulations, we don't just get a pile of numbers, we get graphs representing geometrical shapes, images resembling those we experience in other contexts and so on. This is surely correct, but the problem again seems to be that visual representations are not particular to simulations. We use photographs, and in recent years also video sequences to represent scientifically certain target systems even where no simulation is involved (e.g., microscope footage for processes in neurobiology). And even in 'traditional' disciplines visual representations are common: in our mathematics classes we were asked to draw functions, and in differential geometry we often use drawings to visualise the findings. The same is true when we use any sort of 'iconic' model such as ball and stick models in chemistry, the double helix model of DNA and so on. Moreover, it is, again, at best unclear whether this point has any epistemic import. It seems to be a matter of psychology rather than epistemology that we prefer visual to other forms of representation.

5 Methodology: does simulating constitute a *sui generis* activity, in-between theorising and experimenting?

Simulating is claimed to be an activity in-between theorising and experimentation:

[...] the computational methods of numerical experimentation constitutes a new kind of scientific method, intermediate in kind between empirical experimentation and analytic theory [...]. (Humphreys 1993, p. 103)

The central claim of this paper is that computer simulation provides (though not exclusively) a qualitatively new and different methodology for the physical sciences, and that this methodology lies somewhere intermediate between traditional theoretical physical science and its empirical methods of experimentation and observation. (Rohrlich 1991, p. 507)

[...] computer simulations [...] occupy an uneasy space between theory and experiment, between abstract and concrete, and often between the pressures of pure science and the needs of pragmatic action. (Sismondo 1999, p. 247)

By themselves phrases such as these are hard to understand: what does it mean for an activity or a methodology or a scientific practice to 'lie in between' two others? For

clarity we examine three different ways in which the idea of in-betweenness could be interpreted, one literal and two metaphorical. Our conclusion is that on a literal reading the claim that simulations are in-between theory and experiment is incorrect. On the two metaphorical readings the claim is correct, but it fails to support the thesis that simulations add something new to the philosophy of science because other objects of study, most notably thought experiments and mathematical models in general, fall in-between theory and experiment in the same way.¹²

Taken literally, using the locution ‘falls between’ or ‘lies between’ implies that there is an object, a substance or an event that has a property which is measurable on at least an ordinal scale: topaz lies in between apatite and diamond on Moh’s hardness scale, Easter between Christmas and Pentecost, Maida Vale between Kilburn and Paddington and the specific weight of gold is in between that of silver and that of platinum. What quantifiable properties do simulations and experiments share? Reliability, informativeness, and fruitfulness are but three that readily come to mind. For a want of space we here can only discuss reliability, but we believe that the same conclusion is reached when discussing other properties. Roughly speaking, the reliability of a procedure measures the frequency with which the procedure leads to successful outcomes, for instance in terms of the correctness of predictions made on the basis of the procedure. The claim then is that the reliability of simulations must be in between that of theoretical models and that of experiments.

The only *argument* to this effect we are aware of is due to Morgan (2003, 2005), who supports this claim by an appeal to the ‘materiality’ of the studied systems. *In nuce*, her argument is that inferences are less problematic when the system studied is made of the ‘same kind of stuff’ as the target system and that there is an inferential gap where this is not the case, which results in a lack of reliability.¹³ Her core idea seems to be that the reliability of inferences is grounded in the number of properties that are shared between model and target system. Model animals, such as laboratory mice, are made of the same stuff as other mice and thus share a large number of properties. By contrast, a mathematical model of an economy or a physical system necessarily abstracts from all material properties of the target system, which drastically diminishes the number of shared properties.

An example of a simulation considered by Morgan contains a computer-generated model of cow bones as its core. The model is constructed by making very thin slices of real cow bones, taking photographic digital images of the slices and assembling the slices to re-constitute the three-dimensional bone. In the simulation, experiments are conducted on this computerised image of the real bone. Morgan then likens the type of inference from this simulation to inferences made from model organisms such

¹² Some of the arguments that follow, especially those concerning the role of ‘materiality’ in drawing inferences, have been made independently by Parker (this volume) and Winsberg (this volume). We learned about their papers only after completing this section.

¹³ In all fairness one must say that Morgan only discusses mathematical models in economics which may have problems in their own right and for which her conclusions may well be true. However her distinction is based on the materiality (or absence thereof) of the model and thus applies *mutatis mutandis* also to mathematical models in physics. This point is reinforced by the fact that she takes the idea that mathematical models denote the systems they relate to from R.I.G. Hughes who, of course, introduced the idea in a physics context.

as laboratory mice and argues that these inferences are more reliable than those from mathematical models: '[The simulation experiment] relied on a particular hipbone to produce a model object that maintained (I suggest) enough material qualities necessary to establish valid experimental results about that one real bone' (p. 230). Since one could also establish that this particular hipbone was representative, one could use the results to make inferences about hipbones in general.

There are two problems with this example. First, if one wants to use this example to muster support for the thesis that simulations are in between theories and experiments in terms of reliability one has to believe that this example is representative for the bulk of computational models. Morgan is explicit about the fact that it is not, and this, we think, for good reasons. Paradigm cases of simulation such as climate models and models in solid state physics are not constructed on the basis of material components of these systems (let alone by slicing them up and using computer images of these slices). Second, it is not automatically the case that inferences are reliable whenever model and target are made of the same stuff or that inferences must be unreliable when model and target differ in their materiality. Laboratory organisms are highly artificial constructs, and inferences made on their basis cannot simply be taken outside the lab (see for instance, [Cartwright 1999](#), Chap. 4; for a radical point of view along these lines, see [Latour 1988](#)). In general, two organisms can share a great many properties and yet inferences are unreliable because they depend on those properties that are not shared or because shared properties interact with ones that are not shared and thus invalidate conclusions. A striking example for this feature of models comes from a paper on climate modelling by Norton and Suppe. They compare inferences about natural systems from wind-tunnel experiments with those made on the basis of computer simulations ([Norton and Suppe 2000](#), 70ff.). Surely a wind tunnel is more similar (in Morgan's sense) to the natural atmosphere it stands in for than a computer-generated process. However, wind tunnels produce many experimental artefacts that may invalidate conclusions: walls introduce turbulence affecting airflow past models and causing systematic airflow variations at different chamber locations, which can be ameliorated but not eliminated by experimental design and calibration corrections. There is simply no guarantee that a wind tunnel produces better results than a simulation that is, say, calibrated to measurements taken in more natural environments. Consequently, Norton and Suppe argue (p. 72): 'Suitably done, enhanced computer modelling of data introduces vicarious control every bit as good as, sometimes superior to, traditional experimental control'.

Another example is scale models, which are exact replicas of their target systems except for their size. This point has been made succinctly already by Black, who points out that a change of scale may upset that balance of factors so that the model fails to replicate the most salient feature of the target system: 'Too small a model of a uranium bomb will fail to explode, too large a reproduction of a housefly will never get off the ground ... Inferences from scale model to original are intrinsically precarious and in need of supplementary validation and correction.' ([Black 1962](#), p. 221)

It is also not clear that 'being made of the same kind of stuff' automatically means that a large number of *relevant* properties are shared. Perhaps it is structural properties that enable reliable inferences and such properties can be shared between physical systems and mathematical models. For instance, it turns out that the stability of the

ankle joint in humans only depends on certain geometrical properties of the joint and is by and large independent of the material features of the bones and the other components of the joint. These geometrical features can be captured in a simple mechanical model from which predictions can be derived that are well confirmed by clinical data (see [Frigg et al. 2007](#)).

Thus inferences based on theoretical models are often highly reliable, especially when target systems are experimental systems, while inferences made on the basis of real experiments often fail when the target system is a natural, non-experimental system. In brief, the materiality of an experimental model is immaterial to the reliability of the inferences made on the basis of the model. Simulations, thus, as a kind, are not ‘in between’ theoretical models and experiments in terms of reliability.

Of the two metaphorical readings the first says that simulations are hybrids of theoretical models and physical experiments. They are ‘in between’ models and experiments in the same way in which mermaids are ‘in between’ fish and women or the Minotaur is ‘in between’ a bull and a man. Simulating, then, is in between theorising and experimenting because, on the one hand, it involves, like theorising, the development of ideas and the non-material nature of the system on which manipulations are performed and, on the one hand, and it shares some essential characteristics with experimenting such as tinkering with the system, the involvement of tacit knowledge, the error-correction approach to checking the validity of results as well as the need for data reduction and analysis due to the huge data output.

There are, however, large and important classes of methods in science that are hybrids in just this sense but have nothing to do with the use of digital computers. Space limitations prevent us from discussing the matter more fully but let us very briefly consider thought experiments and mathematical models here. Both these activities are conducted in the mind or on paper, that is, on non-material model systems, but both share important characteristics with experiments. That thought experiments are experiments, albeit a peculiar form, has been argued variously before (see for instance [Sorensen 1992](#)). They are often conducted to test or develop a scientific theory; they use the method of controlled variation; they can be improved upon by tinkering with parameters; in a thought experiment, errors in reasoning are controlled by imagining a highly constrained, often artificial situation. That mathematical models, too, can be experiments is perhaps less obvious. But models also often use the method of controlled variation, they can be improved upon by tinkering with parameters and error-correcting (in deriving a result) sometimes plays an important role. At least in economics, this connection between mathematical models and experiments has been drawn explicitly (e.g., [Lucas 1982](#), p. 696; [Mäki 2005](#)).

The second metaphorical reading says that simulations are in between theoretical models and experiments because they are mediators. If party A wants or needs something from party B (and, perhaps, vice versa) but cannot or prefers not to get in touch directly, she uses a mediator to stand in between herself and B. In this sense bankers are in between creditors and debtors and estate agents in between buyers and sellers.

It is true, of course, that simulations in this way often mediate between more theoretical models or theories on the one hand and experiments, observations or measurements on the other. But this, again, is neither a new observation, nor is it specific to simulations. After all, the mediators idea arose in the context of scientific models,

not simulations (cf. the models as mediators project), and has since been extended to experiments (see for instance [Guala 1998](#)).¹⁴

6 Conclusion

We have argued that the philosophical problems raised by simulations have analogies in the context of modelling, experimenting or thought experimenting and are therefore not entirely new and far from demanding a novel philosophy of science. Those who claim that simulation is special may now answer that by slicing up the problem in the way we did and discussing the problems separately we missed the essential points because these arise precisely in how these problems interact with each other. This is an interesting thesis. It might, for instance, be the case that the availability of non-linear tractable models gives rise to new kinds of idealisations because so far we have most of the time linearised equations when constructing a theoretical model, and fitting such equations to the world requires certain types of idealisations. Now, it may well be that non-linear equations are brought to the world with different idealisations. If this was the case and if we had an account of what the idealisations in both cases are and of how they differ, this would indeed be an exciting new thesis and it would substantiate the claim that simulation adds something new to philosophy of science. However, we could not discover a claim of that sort in the literature.

We spent some twenty pages arguing for this conclusion not because we take malicious joy in spoiling the party; rather we think that it has important consequences for where the discussion is going in the future. Blinkered by the emphasis on novelty and the constant urge to show that simulations are unlike anything we have seen before, we cannot see how the problems raised by simulations relate to existing problems and we so forgo the possibility to have the discussions about simulation make contributions to the advancement of these debates. We suggest that we should rather benefit from synergies than getting hung up on compartmentalising the debate. For instance, if, rather than spilling much ink on convincing ourselves that simulations are unlike everything else, we recognise that the epistemological problems presented to us by simulations have much in common with the ones that arise in connection with models, we can take the insights we gain in both fields together and try to make progress in constructing the sought-after new epistemology. This, we think, is more likely to lead to progress than enviously defending one's own playground against intrusions from outside. We hope that future debates over the nature of simulation will engage more, and more constructively, with discussions surrounding the nature of models and experiments, and we believe that this will be to everybody's benefit.

Acknowledgements We would like to thank Stephan Hartmann, Paul Humphreys, Kevin Korb and Eric Winsberg for discussions and comments on earlier drafts.

¹⁴ [Sismondo \(1999, p. 254\)](#), acknowledges this when he observes that 'Simulations, like models, stand between theories and material objects, pointing in both directions'.

References

- Black, M. (1962). Models and archetypes. *Models and metaphors: Studies in language and philosophy* (pp. 219–243). Ithaca, NY: Cornell University Press.
- Boumans, M. (1999). Built-in justification. In M. Morgan & M. Morrison (Eds.), *Models as mediators* (pp. 66–96). Cambridge: CUP.
- Cartwright, N. (1983). *How the laws of physics lie*. Oxford: Oxford University Press.
- Cartwright, N. (1999). *The dappled world*. Cambridge: CUP.
- Cartwright, N. (2007). The vanity of rigour in economics: Theoretical models and Galileian experiments. *Hunting causes and using them* (pp. 217–235). Cambridge: Cambridge University Press.
- Edwards, P. (2001). Representing the global atmosphere: Computer models, data, and knowledge about climate change. In C. Miller & P. Edwards (Eds.), *Changing the atmosphere: Expert knowledge and environmental governance* (pp. 31–66). Cambridge, MA: MIT Press.
- Fox Keller, E. (2003). Models, simulation, and 'computer experiments.' In H. Radder (Ed.), *The philosophy of scientific experimentation* (pp. 198–216). Pittsburgh: University of Pittsburgh Press.
- Frigg, A., Frigg, R., Hintermann, B., Barg, A., & Valderrabano, V. (2007). The biomechanical influence of tibio-talar containment on stability of the ankle joint. *Journal of Knee Surgery, Sports Traumatology and Arthroscopy*, 15, 1355–1362.
- Frigg, R., & Hartmann, S. (2006, Spring). Models in science. In E. Zalta (Ed.), *Stanford Encyclopedia of philosophy*. Downloadable at: <http://plato.stanford.edu/archives/spr2006/entries/models-science>.
- Galison, P. (1996). Computer simulation and the trading zone. In P. Galison & D. Stump (Eds.), *Disunity of science: Boundaries, contexts, and power* (pp. 118–157). California: Stanford University Press.
- Guala, F. (1998). Experiments as mediators in the non-laboratory sciences. *Philosophica*, 62, 901–918.
- Guala, F. (2005). *The methodology of experimental economics*. Cambridge: Cambridge University Press.
- Hartley, J., Hoover, K., & Salyer, K. (1997). The limits of business cycle research: Assessing the real business cycle model. *Oxford Review of Economic Policy*, 13(3), 34–54.
- Hartmann, S. (1996). The world as a process: Simulation in the natural and social sciences. In R. Hegselmann, U. Müller & K. Troitzsch (Eds.), *Modelling and simulation in the social sciences from the philosophy of science point of view* (pp. 77–100). Dordrecht: Kluwer.
- Humphreys, P. (1991). Computer simulations. *Philosophy of Science PSA*, 19902, 497–506.
- Humphreys, P. (1993). Numerical experimentation. In P. Humphreys (Ed.), *Patrick Suppes: Scientific philosopher* (Vol. 2). Dordrecht: Kluwer.
- Humphreys, P. (1995). Computational science and scientific method. *Mind and Machines*, 5, 499–512.
- Humphreys, P. (2004). *Extending ourselves: Computational science, empiricism, and scientific method*. Oxford: OUP.
- Latour, B. (1988). *The pasteurisation of France*. Cambridge, MA: Harvard University Press.
- Little, D. (Ed.). (1995). *On the reliability of economic models: Essays in the philosophy of economics*. Dordrecht: Kluwer.
- Lucas, R. (1982). *Studies in business cycle theory*. Cambridge, MA: MIT Press.
- Mäki, U. (2005). Models are experiments, experiments are models. *Journal of Economic Methodology*, 12(2), 303–315.
- Morgan, M. (2003). Experiments without material intervention: Model experiments, virtual experiments, and virtually experiment. In H. Radder (Ed.), *The philosophy of scientific experimentation* (pp. 216–235). Pittsburgh, PA: University of Pittsburgh Press.
- Morgan, M. (2004). Simulation: The birth of a technology to create "evidence" in economics. *Revue d'Histoire des Sciences*, 57, 341–377.
- Morgan, M. (2005). Experiments versus models: New phenomena, inference and surprise. *Journal of Economic Methodology*, 12, 317–329.
- Morgan, M., & Morrison, M. (1999a). Models as mediating instruments. *Models as mediators: Perspectives on natural and social science* (pp. 10–37). Cambridge: Cambridge University Press.
- Morgan, M., & Morrison, M. (1999b). *Models as mediators: Perspectives on natural and social science*. Cambridge: Cambridge University Press.
- Morton, A. (1993). Mathematical models: Questions of trustworthiness. *British Journal for the Philosophy of Science*, 44, 659–674.
- Norton, S., & Suppe, F. (2000). Why atmospheric modeling is good science. In C. Miller & P. Edwards (Eds.), *Changing the atmosphere: Expert knowledge and environmental governance*. Cambridge, MA: MIT Press.

- Reiss, J. (2007). *Error in economics: Towards a more evidence-based methodology*. London: Routledge.
- Rohrlich, F. (1991). Computer simulation in the physical sciences. *PSA 1990, II*, 507–518.
- Sismondo, S. (1999). Models, simulations and their objects. *Science in Context, 12*, 247–260.
- Smith, P. (1998). *Explaining chaos*. Cambridge: Cambridge University Press.
- Sorensen, R. (1992). *Thought experiments*. Oxford: Oxford University Press.
- Srivastava, N., Kaufman, C., & Müller, G. (1990). Hamiltonian chaos. *Computers in Physics, 4*(5), 549–553.
- Sterman, J. (2006). Learning from evidence in a complex world. *American Journal of Public Health, 96*(3), 505–514.
- Stöckler, M. (2000). On modelling and simulations as instruments for the study of complex systems. In M. Carrier (Ed.), *Science at century's end: Philosophical questions on the progress and limits of science*. Pittsburgh: University of Pittsburgh Press.
- Sugden, R. (2000). Credible worlds: The status of theoretical models in economics. *Journal of Economic Methodology, 7*(1), 1–31.
- Winsberg, E. (1999, Summer). Sanctioning models: The epistemology of simulation. *Science in Context*.
- Winsberg, E. (2001). Simulations, models, and theories: Complex physical systems and their representations. *Philosophy of Science, 68*(Proceedings), S442–S454.
- Winsberg, E. (2003). Simulated experiments: Methodology for a virtual world. *Philosophy of Science, 70*, 105–125.