

Simulated Experiments: Methodology for a Virtual World*

Eric Winsberg^{†‡}

This paper examines the relationship between simulation and experiment. Many discussions of simulation, and indeed the term “numerical experiments,” invoke a strong metaphor of experimentation. On the other hand, many simulations begin as attempts to apply scientific theories. This has led many to characterize simulation as lying between theory and experiment. The aim of the paper is to try to reconcile these two points of view—to understand what methodological and epistemological features simulation has in common with experimentation, while at the same time keeping a keen eye on simulation’s ancestry as a form of scientific theorizing. In so doing, it seeks to apply some of the insights of recent work on the philosophy of experiment to an aspect of theorizing that is of growing philosophical interest: the construction of local models.

1. Introduction. Many complex systems in the physical sciences are studied by developing models of their underlying physics on a computer, and by using computationally intensive methods to learn about the behavior of those systems. These methods are called “simulations,” or “numerical experiments”; names that strongly evoke the metaphor of experimentation. At the same time, the mathematical models that drive these particular kinds of simulations are motivated by theory. *Prima facie*, they are nothing but applications of scientific theories to systems under the theories’ domain. So where “on the methodological map”¹ do techniques of computer

*Manuscript received December 2001; revised July 2002.

[†] To contact the author, please write to Eric Winsberg, Department of Philosophy, The University of South Florida, FAO 226, 4202 Fowler Ave., Tampa FL 33620; e-mail: winsberg@chuma1.cas.usf.edu.

[‡] I thank Arthur Fine, Mathias Frisch, Karen Darling, and Joanne Waugh for their comments and criticisms, as well as many people who offered helpful comments when I presented earlier versions of this paper at the University of South Florida and at Wichita State University. I am greatly indebted to three anonymous reviewers—especially one, whose efforts went well beyond the norm.

1. See quotation from Galison below.

Philosophy of Science, 70 (January 2003) pp. 105–125. 0031-8248/2003/7001-0009\$10.00
Copyright 2003 by the Philosophy of Science Association. All rights reserved.

simulation lie? What features do simulations actually share with laboratory experiments?

While there is nothing like a received view of the relation between simulation and experiment, I will piece together and critically examine three perspectives from the scientific, philosophical, and other “STS” literature. I will argue that while each of these views can contribute to an understanding of simulation, none of them is quite suitable for understanding how simulation can have methodological and epistemological features in common with experimentation, while still playing the role of a form of scientific theorizing. This examination will apply some of the insights of recent work on the philosophy of experiment to understand how, and to what extent, theoretical principles guide the construction of locally representative models.²

My examination will both amplify and illuminate Mary Morgan and Margaret Morrison’s recent characterization of “autonomous models”:

One of the points we want to stress is that when one looks at examples of different ways that models function, we see that they occupy an autonomous role in scientific work. We want to outline . . . an account of models as *autonomous agents*, and to show how they function as *instruments* of investigation. . . . It is precisely because the models are partially independent of both theories and the world that they have this autonomous component and so can be used as instruments of exploration in both domains (Morrison and Morgan, 1999, p.10, original emphasis).

For our purposes, the term “autonomous models” is somewhat misleading. A better term would be “semiautonomous.” The claim frequently made by Morrison and Morgan that models are autonomous or independent of theory is meant to emphasize the fact that there is no algorithm for reading models off from theory. While models generally incorporate a great deal of the theory or theories with which they are connected, they are usually fashioned by appeal to, by inspiration from, and with the use of material from, an astonishingly large range of sources: empirical data, mechanical models, calculational techniques (from the exact to the outrageously inexact), metaphor, and intuition. But to call these models completely “autonomous,” at least in this context, is to deny the obvious and strong connections these models have to theory.

Examining the methodological role and epistemic status of simulation techniques will help to clarify the relevant notion of a “semiautonomous model.” Clarifying Morgan and Morrison’s notion of “autonomy” is a special case of a wider objective: in simulation studies, model building is

2. For example, see Morgan and Morrison 1999, Cartwright 1999, and Giere 1999.

carried out in the open. That is, unlike in other scientific domains where modeling assumptions are often swept under the rug with a phrase like “it can easily be shown” in simulation studies, modeling assumptions are often much more explicitly motivated and scrutinized by the practitioners themselves. Simulation provides a clear window through which to view the model making process.³

2. Motivating Models.

Equation (1.2–9) is a second order, nonlinear, vector, differential equation which has defied solution in its present form. It is here therefore we depart from the realities of nature to make some simplifying assumptions . . . (Bate, Mueller, and White 1971)

Computer simulations are techniques for studying mathematically complex systems. They have applications in almost every field of scientific study—from quantum chemistry to meteorology and from paleontology to the study of traffic flow patterns. Simulations are typically classified according to the type of algorithm that they employ. “Discretization” techniques transform continuous differential equations into step-by-step algebraic expressions. “Monte Carlo” methods use random sampling algorithms even when there is no underlying indeterminism in the system. “Cellular automata” assign a discrete state to each node of a network of elements, and assign rules of evolution for each node based on its local environment in the network.

In this paper, I will be focusing on simulations that employ methods of discretization. This method begins with a mathematical model that depicts the time-evolution of the system being studied in terms of rules of evolution for the variables of the model. For discretization techniques to apply, these rules of evolution come in the form of differential equations. In some cases (typically in the physical sciences) the mathematical model is motivated by well-established theoretical principles. In other cases, the model may be based on speculation, the consequences of which need to be studied. In this paper, I will focus on theoretically motivated models.

When a theoretically motivated model is to be discretized, the original differential equations transformed into a computable algorithm that can be run on a digital computer, and the evolution of the dependant variables

3. A note on language: In this paper I move back and forth between talk of models and talk of simulation. I use the term “simulation” to refer to comprehensive process of building, running, and inferring from computational models. Simulations are based on models, they incorporate model assumptions, and they in turn produce models of phenomena. For a more thorough discussion of the role different kinds of models play in the simulation process, see my 1999.

on the computer is said to “simulate” the evolution of the system in question. Sophisticated graphical techniques are often used to transform the output into graphics and videos that sometimes resemble images such as might be produced by laboratory instruments trained upon the system in question.

We can see here the importance of emphasizing the semiautonomy of models. In these simulations, the principle model is hatched directly out of theory. But as the quotation at the beginning of this section suggests, it is one thing for theory to directly yield a model, and quite another for theory to be able to yield a model for which we can find solutions.

In fact, for many of the systems that are of interest in the computationally intensive sciences, the models that are suggested directly by theory consist of second-order, non-linear differential equations; and finding useful and reliable solutions for these models, even using numerical methods, is an unrealistic goal. Successful numerical methods, therefore, invariably require of the simulationists that they transform the model suggested by theory in significant ways. Idealizations, approximations, and even self-conscious falsifications are introduced into the model. In the end, the model that is used to run the simulation is an offspring of the theory, but it is a mongrel offspring. It is also substantially shaped by the exigencies of practical computational limitations and by information from a wide range of other sources.

A concrete example might help to clarify these remarks. Suppose we are interested in understanding how vortices are generated when a shock wave passes through an inhomogeneous medium at supersonic speeds. One way to study the problem is in the laboratory. Skilled experimenters have conceived of idealized experiments involving the interaction of shock waves in air with spherical and cylindrical bubbles of various gases of various densities. The purpose of these experiments is to help shed light on the general dynamical features of these kinds of shock waves. Shadowgraph images of these experiments can reveal a great deal of valuable information. (See Winkler et al. 1987)

But by “simulating” these same experimental setups on a computer, a much wider variety of initial conditions can be studied and much more detailed data can be gathered. In order to run such a simulation, the simulationist begins by choosing a principle model—a model that characterizes the system in terms of both the arrangement of its constituent parts, and the rules of evolution that govern the changes of state that the system undergoes over time. In our case, the choice of how to do this is motivated by our understanding of the experimental setup, and by our best theories of the flow of compressible fluids—in this case the Euler equations for an inviscid compressible gas. These equations describe a fluid that is compressible but they describe neither the effects of viscosity nor of heat conduction.

Since these equations are analytically unmanageable under all but the most symmetric and time-independent conditions, simulation requires computational methods, that, given reasonable computer resources, can accurately trace out the patterns of flow. This is not a matter of simply taking continuous differential equations and transforming them into discrete algebraic equations. Given limitations on computer speed and memory, these techniques also invariably resort to other approximations, idealizations, and even “falsifications”—that is, model assumptions that directly contradict theory. Making the simulation work, and making it produce results that the simulationist is willing to sanction as reliable, is a skill that has been developed in a lengthy period of trial, error, and comparison with both theory and known results from physical experiments. In sum, by the semiautonomy of a simulation model, one refers to the fact it starts from theory but one modifies it with extensive approximations, idealizations, falsifications, auxiliary information, and the blood, sweat, and tears of much trial and error.

In our example, the primary difficulty in modeling these supersonic fluid flows is capturing the discontinuities—abrupt changes in value in the flow variables. It is in overcoming such difficulties that extra-theoretic modeling “tricks” are invoked. In particular, in simulations of supersonic flows, there are important effects in and around shock fronts that take place at length and time scales that are much too small to be captured in even the most finely grained discretizations. Thus, some kind of falsification must be employed in the model in order to resolve these key effects (Woodward and Collela, 1984). But this raises an important epistemological question; one that needs to be reflected on if we are to take the notion of an autonomous model seriously. What, if not theory, can serve to sanction these models? This question will be lurking in the background as we investigate the relation between theory and experiment. It will return to center stage later on.

3. It’s All Metaphorical. As mentioned, there are three different views in the literature about how to methodologically classify simulation *vis à vis* its experimental cousins. One answer is that all talk of “simulation” and “numerical experiments” is purely hyperbolic or metaphorical—simulation is nothing more and nothing less than using brute-force computational means to solve analytically intractable equations. A second view, in which the terms “simulation” and “numerical experiment” are taken quite literally, a simulation is a stand-in, or mimic, of a real-world system, and can therefore be experimented on just like any other experimental target. Finally, there is the view that simulation is a brand new “third mode” of science, neither experimental nor theoretical. In what follows, I will weigh the merits of these three views, emphasizing how much each of

these views can contribute to understanding how models are to be evaluated.

I should point out that I know of no published references that themselves argue for exactly the view that simulation is just a fancy word for using brute-force computational means to solve analytically intractable equations. It is, however, a not-uncommon opinion expressed in conversation.

One perspective from which to see clearly that this view is very wide of the mark has been offered by the philosopher R. I. G. Hughes, who has urged a distinction between “the use of computer techniques to perform calculations, on the one hand, and [genuine] computer simulation, on the other” (Hughes 1999, 128). For Hughes, the distinction hinges on the “thoroughly realist mode of description” that is used to describe the results of genuine simulations, and to the images that are produced by simulation, which often resemble photographs of material physical models. Thus, what distinguishes genuine simulations from mere number crunching is that simulations have genuinely “mimetic” (Hughes’ word) characteristics.

There is certainly something right about this. The “mimetic” characteristics of an algorithm—one that uses sophisticated graphics and that is treated realistically by its users—are surely an important consideration that lead practitioners to call them simulations rather than mere computations. So psychologically, at the very least, working with a simulation is much more like doing an experiment if the simulation produces life-like images reminiscent of laboratory photographs.

Of course, we are looking for something deeper than a superficial, psychological resemblance. In fact, if this were all that simulations, or “numerical experiments”, had in common with experiments, then one might indeed be tempted to agree that talk of simulation as numerical experimentation was no more than mere hyperbole.

To illustrate this point, we need look no further than the example of mere number crunching offered by Hughes; the use of a computer to calculate the orbits of planets in a three-body problem. In such a case, there is nothing whatsoever that prevents creators of such an algorithm from imbuing it with high quality graphical output, and from characterizing their results in as realistic a mode of description as they like. Since no distinction of great philosophical import could hinge on whether or not “mere” computation is embellished with graphic presentation, we are going to have to do more than cite the fact that some algorithms use graphics and some don’t.

There are in fact two characteristics of true simulations that I argue meaningfully distinguish them from mere brute-force computation, in ways that connect them to experimental practice in an interesting fashion:

- Successful simulation studies do more than compute numbers. They make use of a variety of techniques (most of which, pace Hughes, involve imaging in one way or another) to draw inferences from these numbers.
- Simulations make creative use of calculational techniques that can only be motivated extra-mathematically and extra-theoretically.⁴As such, unlike simple computations that can be carried out on a computer, the results of simulations are not automatically reliable. Much effort and expertise goes into deciding which simulation results are reliable and which are not.

Let me elaborate on these two themes. First, simulations are interestingly like experiments, because they involve, in essence, data analysis. Suppose that we are confronted with a fluid flow problem like the one discussed above. In such a case, we think that we know the governing principles (in this case, a form of the Euler equations), but we still do not know what kind of local behavior these principles predict. One approach would be to use a tractable, approximative, analytical technique. For example, analysts might use a series expansion, truncate all but two or three of the terms and end up with an equation for which one can write down a solution. Or, one might put to work the kind of more indirect simulation methods indicated in the last section. Both of these techniques employ approximations, both involve creative work, and both will, in turn, give rise to issues of justification.

One important difference between the two methods, however, is this: the analytic method will produce an algebraic expression as its result. That expression, in turn, can represent the behavior of a general class of systems. Various functional dependencies and patterns of behavior can easily be read off from a closed form expression.

Numerical methods, on the other hand, result in a big pile of numbers. If simulationists want to learn about the general qualitative features of a class of systems, then they must apply all the usual tools of experimental science for analyzing data: visualization, statistics, data mining, etc. If they want to discover functional dependencies, then they must also run a barrage of trials, looking at the results across a wide range of parameters. It is without a doubt this aspect of simulation that carries the most obvious methodological characteristics of experimental work. Furthermore, it is

4. I note that a similar point about simulations has also been made by Paul Humphrey in his 1995. He refers to the techniques that motivate the construction of these models as heuristics. He explicitly contrasts ‘heuristics’ with determinate methods for moving from theory to model. See also Winsberg 1999 for more details on how model construction is often guided, but not determined, by theory.

the need to draw inferences from these piles of data that brings about what Hughes calls the “mimetic” aspects of simulation—their “thoroughly realistic mode of description.”

In a review article on computational methods, a group of researchers describe the situation this way:

The most common method for observing the behavior of laboratory flows is to make photographs using a variety of techniques that bring out specific features. . . . Computational fluid dynamicists naturally want to use similar techniques to display their results. Displays not only make comparisons with laboratory data much easier, but also are useful at getting at the fundamental physics of wave interactions, surface instabilities, vortex generation and other phenomena that may be involved in the flow. (Winkler et al. 1987, 29)

Recall Hughes’ insistence that genuine simulations be distinguished from mere step-by-step calculations on the basis of the former’s mimetic qualities. It seems to me that what underlies this distinction is really the difference between computational methods that are used to make specific quantitative predictions (for example, where will the fourth planet in a five-body system be after 72 days) on the one hand and computational methods used to “get at the fundamental physics”⁵ on the other. It is only in the latter case that imaging techniques become *de rigueur* and that the computational model takes on a mimetic quality, so that techniques familiar to the laboratory observer can be applied. But whether or not imaging techniques per se are used, simulations that aim to get at global, diachronic features of the systems they model will invariably not only produce numbers, but also draw inferences from those numbers in a manner that is richly analogous to experimental data analysis.

The second feature of simulations with respect to which they have a great deal in common with experimental practice is the constant concern with uncertainty and error. As I have already emphasized, although simulation often is initially motivated by well-established theory, the end model that drives the computations generally incorporates modeling assumptions that are not theoretically motivated. The results of the simulation, therefore, do not automatically come with a stamp of approval that carries the

5. Obviously, Winkler et al. are using the expression “fundamental physics” in a manner that is quite different than that to which philosophers might be accustomed. I take it that they are using the expression to signify not the fundamentals laws or theories of the system, but rather the emergent structural features of the dynamics like vortices, waves, and surfaces and their interactions—the important, dynamically significant long-term and long-range features of the dynamics that give insight into its time-evolution. This notion of “fundamental” physics and its connection to accounts of scientific explanation is an area worthy of further research.

full faith and credit of the governing theory's epistemic credentials. Quoting again Winkler—"Unless uncertainties are kept under control, the computational approach cannot uncover *new* physical phenomena" (Ibid. 29; original emphasis).

That simulations and laboratory experiments both have a need to manage uncertainties is perhaps interesting in and of itself. But if we are interested in epistemological issues surrounding autonomous models, the really interesting question is whether or not they do so in analogous ways. If so, then perhaps some of the insights from recent work in the philosophy of experiment might shed some light on the sanctioning of models.

I will have more to say about these questions later on. For now it will suffice to address the challenge of this section with the following point: while having a mimetic quality is not in itself what gives a simulation interesting methodological and epistemological features, it is an important sign of other features which do. The extensive use of realistic images in simulation is a stepping stone that simulationists use in order to make inferences from their data. It is also a tool they use in order to draw comparisons between simulation results and real systems; a move that is part of the process of sanctioning their results. It is the drawing of inferences and sanctioning of results that give rise to interesting philosophical connections between simulation and experimental practice.

4. The Computer as Experimental Target. Another view on how to locate simulation methodologically has been to interpret expressions like "simulation" and "numerical experiment" literally. The idea here is to interpret the simulation algorithm as literally "mimicking" the system or systems of interest and to understand what scientists do as performing experiments on the computer or computer algorithm, which acts as a stand in for, or probe of, the system in question. Computer simulations are experiments, and the computer is the target of the experiment.

Practitioners of simulation often emphasize this point of view themselves, especially in popular or semi-popular presentations that are geared towards expounding the virtues of numerical methods. For example:

A simulation that accurately mimics a complex phenomenon contains a wealth of information about that phenomenon. Variables such as temperature, pressure, humidity, and wind velocity are evaluated at thousands of points by the supercomputer as it simulates the development of a storm, for example. Such data, which far exceed anything that could be gained from launching a fleet of weather balloons, reveals intimate details of what is going on in the storm cloud. (Kaufmann and Smarr 1993, 4)

The germ of this idea probably comes from the father of computational methods himself, John von Neumann—although he actually expressed the

idea in reverse. Working on highly intractable fluid dynamical problems at Los Alamos, von Neumann lamented the fact that he and his colleagues often had to perform difficult experiments just to determine facts that should, in principle, be derivable from well-known underlying principles and governing equations:

The purpose of the experiment is not to verify a proposed theory but to replace a computation from an unquestioned theory by direct measurement. . . . Thus wind tunnels are used . . . as computing devices . . . to integrate the nonlinear partial differential equations of fluid dynamics.” (quoted in Winkler et al. 1987, 28)

Once von Neumann’s dream became ostensibly realized, and the wind tunnel was replaced by the computer, it becomes not-entirely unnatural to view the resulting activity as performing experiments in a virtual wind tunnel. Something like this view of computer simulation has received the most attention by philosophers of science. Paul Humphreys (1994), for example, argues for a view like this *vis à vis* Monte Carlo simulations, on the grounds that when a program runs on a digital computer, the resulting apparatus is a physical system. Any runs of the algorithm, therefore, are just experimental trials on physical target. This is also roughly the view that is espoused by Hughes. According to Hughes, there is no reason that we should resist thinking of computer simulations as experiments since they lie on a “slippery slope” that makes them conceptually inseparable from experiment. The slope looks something like this:

1. We experiment on a model water wheel in order to learn about water wheels in general.
2. We experiment on an electrical damped harmonic oscillator in order to learn about mechanical damped harmonic oscillators that are governed by a structurally identical equation.
3. We run experiments on cellular automata machines in order to learn about systems with identical “symmetries and topologies.”
4. Why not also say, then, that we perform experiments on computers running algorithms designed to simulate complex physical systems? (Hughes, 1999, 138f)

In a similar vein, Steve Norton and Frederick Suppe argue on an explicitly epistemological basis for the claim that simulation is a form of experiment. According to Norton and Suppe, a valid simulation is one in which certain formal relations hold between a base model, the modelled physical system itself, and the computer running the algorithm (Norton and Suppe 2001). Glossing, the relation is that of *realization*, where a system S_1 *realizes* a system S_2 just in case there is a many-one behavior-

preserving mapping from the states of S_2 onto the states of S_1 . When the proper conditions are met, a simulation

can be used as an instrument for probing or detecting real world phenomena. Empirical data about real phenomena are produced under conditions of experimental control. (Norton and Suppe 2001)

Simulation modeling is just another form of experimentation, and simulation results are nothing other than models of data. (Ibid., emphasis in original)

In general then, the claim being made by Hughes, and by Norton and Suppe is that the methodological structure of simulation is like that of experimentation because simulation proceeds in the following way:

1. Create an algorithm that accurately mimics a physical system of interest.
2. Implement the algorithm on a digital computer.
3. Perform experiments on the computer, which will tell us about the system in question—like experiments on an electrical damped harmonic oscillator tell us about mechanical damped harmonic oscillators.

The real problem with this sort of story, from the point of view of the present project, is that it begs the question of whether or not, to what extent, and under what conditions, a simulation reliably mimics the physical system of interest. We are interested in applying some of the insights of the philosophy of experiment to methodological and epistemological issues *vis à vis* simulations. But to identify the methods of simulation with the methods of experiment in this way is to tuck away all of these important questions into step 1 and then to focus exclusively on step 3 as the step where the connection lies between simulation and experiment.

Hughes himself at least partially recognizes this problem. His solution is to say that computer “experiments” reveal information about actual, possible, or impossible worlds. To know that we are finding out about actual worlds, according to Hughes, requires an extra step. “Lacking other data, we can never evaluate the information that these experiments provide” (Hughes 1999, 142). This is not so much a problem for Hughes since his analysis is not intended to be epistemological. In fact, in the last quoted passage, Hughes is essentially disavowing any epistemological force to his analogy. But we must face this issue if we want to offer the analogy between simulation and experiment as a way of understanding what makes simulation results reliable. One of the central epistemological questions about simulation is “how do we evaluate the information that these simulations provide when other data *are* lacking?”

This state of affairs reminds me of a passage of Wittgenstein’s, in which he is criticizing Ramsey’s theory of identity:

Ramsey's theory of identity makes the mistake that would be made by someone who said that you could use a painting as a mirror as well, even if only for a single posture. If we say this, we overlook that what is essential to a mirror is precisely that you can infer from it the posture of a body in front of it, whereas in the case of the painting you have to know that the postures tally before you can construe the picture as a mirror image. (Wittgenstein P.R., sec. 121, quoted in Marion 1998, 70)

For us to adopt this particular analogy between simulation and experiment would be to make the same mistake as the one identified by Wittgenstein. Simulation is a technique that begins with well-established theoretical principles, and through a carefully crafted process, creates new descriptions of the systems governed by those principles. It is a technique that, when properly used, will provide information about systems for which previous experimental data is scarce. Paraphrasing Wittgenstein, "What is essential to a [simulation] is precisely that you *can* infer from it" what some system in the real world is like even when other data *are* lacking. If in our analysis of simulation we take it to be a method that essentially begins with an algorithm antecedently taken to accurately mimic the system in question, then the question has been begged as to whether and how simulations can, and often do, provide us with genuinely new, previously unknown knowledge about the systems being simulated. It would be as mysterious as if we could use portraits in order to learn new facts about the postures of our bodies in the way that Wittgenstein describes. We need to understand how, as Kaufmann and Smarr suggest, we can reliably learn about storms from simulations, even when data about such storms are conspicuously sparse.

Furthermore, simulations often yield sanctioned and reliable new knowledge of systems even when nothing like the stringent conditions required by Norton and Suppe are in place. In practice, simulationists need not suppose—nore even begin to suppose—that their simulations perfectly mimic any particular physical system in order to convince themselves that certain qualitative properties of their results can reliably be attributed to the systems being studied.

If there is a useful analogy to be made by philosophers between simulation and experiment, then that analogy ought to help make methodological and epistemological connections. These connections, in turn, should help us to apply some of the insights of recent work in the philosophy of experiment towards gaining an understanding of the conditions under which we should take simulation results to be accurate representations of real systems. Thus far, this proposal fails to do this because it assigns experimental qualities only to those aspects of simulation reason-

ing that occur *after* it is assumed that the simulation algorithm ‘realizes’ the system of interest.

5. A Third Mode. A third view on simulation’s methodological geography is that simulation represents an entirely new mode of scientific activity—one that lies between theory and experiment. For example, Fritz Rohrlich, a physicist and philosopher of science, writes:

Computer simulation provides . . . a qualitatively new and different methodology for the physical sciences . . . This methodology lies somewhere intermediate between traditional theoretical physical science, and its empirical methods of experimentation and observation. (1990, 507)

Historians, sociologists, and even leading practitioners of the techniques themselves have expressed something like this view. In an introductory essay for a special issue of *Physics Today* on “computational physics,” Norman Zabusky wrote:

Supercomputers with ultrafast, interactive visualization peripherals have come of age and provide a mode of working that is coequal with laboratory experiments and observations, and with theory and analysis. (1987, 25)

Deb Dowling, a sociologist of science, has also suggested a similar interpretation of simulation. According to her, simulation is like theory in that it involves “manipulating equations” and “developing ideas” but is like experiment in that it involves “fiddling with machines,” “trying things out,” and “watching to see what happens” (Dowling 1999, 264). The historian Peter Galison takes a similar view:

[Simulation] ushered physics into a place paradoxically dislocated from the traditional reality that borrowed from both the experimental and theoretical domains, bound these borrowings together, and used the resulting bricollage to create a marginalized nether land that was at once nowhere and everywhere on the usual methodological map. (Galison 1996, 120)

Locating simulation as lying *between* theory and experiment provides a natural perspective for both historians and sociologists. In Galison’s case, it is a useful way of expressing how simulation could provide a trading zone between experimentalists and theoreticians and across disciplines. Dowling, in turn, has argued that simulation’s ambiguity with respect to the experiment/theory dichotomy can play important social and rhetorical functions.

From the point of view of the project of this paper, however, it is not

clear what we gain by saying that simulation “lies between theory and experiment”. What is of interest philosophically is to understand (a) how it is that what is at root a theoretical enterprise, takes on characteristics of experimentation, (b) what those characteristics are—at the abstract, reconstructed level, (c) what consequences there are of such a hybrid for our understanding of the nature of modeling, theorizing, and experimenting, and d) how simulation produces knowledge and what kind of knowledge that is. For these purposes, making simulation out to be an entirely new methodology and urging that it lies between theory and experiment is, at best, a good place to start.

6. Speculation vs. Calculation. One way in which we might be able to make this approach do more work for us is to be clearer about the concept of “theorizing” when we say, for example, that simulation is an activity that lies between theorizing and experimenting. A good place to start is with Ian Hacking’s repudiation of the traditional dichotomy of theory and experiment. He urges that we replace it with a tripartite division: speculation, calculation, and experimentation:

[By calculation] I do not mean mere computation, but the mathematical alteration of a given speculation, so that it brings it into closer resonance with the world. (Hacking 1983, 213–14)

The point that Hacking is making with his neologisms is that there are really two quite distinct activities that we often naively lump together under the label of “theorizing.” The first activity is that of laying out basic theoretical principles: Maxwell’s equations, Newton’s laws of motion, Einstein’s field equations, etc. This is the activity Hacking calls “speculation.” The second activity is what Thomas Kuhn long ago called “theory articulation.” This is the hard work that is involved in making the aforementioned theoretical principles apply to the local, concrete systems that make up the real world. Hacking calls this activity “calculation”, but I prefer the simple expression “model building” to Hacking and Kuhn’s phrases because it emphasizes that this is an activity that often bring us *beyond* the original theoretical principles themselves. The expression “model building” also is meant to emphasize that the model being built is not properly seen as a component of the theory proper. The models constructed are offspring of theories, but they emerge as semiautonomous agents. In this respect, Hacking’s pair of terms “speculation” and “calculation”, are somewhat misleading, since in conjunction they suggest that there is nothing speculative about the process of building models under the guidance of theory, when nothing could be further from the truth.

The word “theorizing,” as it is *naively* used, expresses some amalgam of these two distinct activities, and it effectively collapses a valuable dis-

tion. This sloppiness of language has persisted, in no small part, because most commentators on science, especially philosophers, have woefully underestimated the importance of theory articulation, or model building.

It is this under appreciation of the importance of model building that Nancy Cartwright has labeled the “vending machine view of theories” (Cartwright 1999, 179–183). On this view, criticized by Cartwright, the ability of a theory to represent the world is captured precisely by the set of those conclusions that can be drawn deductively from the theory—drawn, moreover, with the ease with which we can extract candy from a vending machine.

For Cartwright, theories do not have by themselves the power necessary to represent real, local states of affairs. Only what she calls representative models are fully able to “represent what happens, and in what circumstances.” (Ibid, 180) It is these models that represent “the real arrangements of affairs that take place in the world (Ibid, 180). On her view, the process of creating representative models from theory is complex and creative.

Having widened the field to three activities instead of the traditional two, it is a simple point that simulation is a form of theory articulation or “model building.” The example of a simulation described above is a complex and creative mathematical alteration of the theory of fluids designed to create a representative model of what motions occur when two fluids collide at supersonic speeds. Using Cartwright’s language, simulationists work *with* theories (among other elements) but what they build are “representative models.”⁶ The fact that simulation work is such a creative and such an epistemologically delicate process is grist for Cartwright’s mill. The difficulty of obtaining reliable simulation results is a testament to the claim that theories do not dispense representative models as easily and conveniently as a vending machine dispenses candy.

Nevertheless, I do not want to simply make the point that simulation—rather than being either theorizing, or experimenting, or some midway point between the two—is actually a form of what Hacking calls “calculation”. Even with a tripartite distinction in hand, there are still aspects of simulation that, both methodologically and epistemologically, seem to have characteristics more commonly associated with experimental practice rather than with the pencil and paper varieties of calculation. Nevertheless, I think it is very important to keep a keen eye on the distinction between the laying out of theoretical principles and the construction of local models when we throw around the word “theorizing.”

In fact, I would argue, something like what Hacking and Cartwright

6. What I called “models of phenomena” in Winsberg 1999.

say about theorizing must be true if anything at all can be a significant hybrid of experiment and theory. In order to avoid the appearance of there being anything strange or paradoxical about a practice that straddles the terrain between the theoretical and the experimental, we need to recognize that while simulation is, in the general sense, a form of what we once naively called theorizing, it is the kind of theorizing that has only recently begun to attract philosophical attention—construction of local, representative models. In the end, the point is simply this: philosophers like Hacking, Cartwright, and Giere⁷ have afforded us the insight that we can have good, reliable theoretical knowledge in a particular domain and still have a lot of difficult, creative work to do in building local models under that domain. This insight is crucial if we are going to understand what goes on in simulation.

7. The Experimental Qualities of Simulation. Hacking, Cartwright, Giere, Morrison and Morgan and others have shown how models can function semiautonomously from theory. What has perhaps been lacking from their analysis is an understanding of where semiautonomous models get their credentials. There is an enormous and controversial philosophical literature on how theories get credentialed. But even when it is established that the theory for a given domain is credible and reliable, how do we come to the conclusion that the local, autonomous models of the phenomena in that domain are reliable? As we have seen, it is not simply a matter of the model's fidelity to theory, since, in simulation, model construction often involves steps that go beyond, or even contradict, theory. It is also not simply a matter of fidelity to real-world data, since we often run simulations in order to learn more about the world than our observations will allow. The literature on the philosophy of experiment may help with these questions.

A starting point for this project is to point out that simulationists and experimenters both need to engage in error management. In simulations, errors can arise the as a result of transforming continuous equations into discrete ones and of transforming a mathematical structure into a computational one. All discretization techniques present the possibility of roundoff errors or instabilities creating undetected artifacts in simulation results. At a deeper level, any modeling assumption that goes into the creation of a simulation algorithm can have unintended consequences. Developing an appreciation for what sorts of errors are likely to emerge under what circumstances is as much an important part of the craft of the simulationists as it is of the experimenter. Precision, accuracy, error analysis and calibration are concepts that we typically associate with experi-

7. For example, Giere (1999).

mentation, and not with theorizing, but they are also very much a part of the vocabulary of the simulationist. There is indeed a great deal of similarity and analogy between the actual techniques that experimenter and simulationist each use to manage uncertainty.

In his work on the epistemology of experiment, Alan Franklin (1986) has outlined a list of common sense techniques that experimenters use to augment our reasonable belief in the results of their work. It is a straightforward exercise to go through this list and see that many, if not all, of these techniques apply directly or by analogy to the sanctioning of simulation results.⁸ Here, I will focus on a deeper aspect of experimental epistemology—found in the writings of Hacking and Peter Galison (1997)—according to which various experimental techniques and instruments develop a tradition that gives them their own internal stability, or, put most provocatively, that “experiments have a life of their own” (Hacking, 1988).

8. A Life of Their Own.

I [once] wrote that experiments have a life of their own. I intended partly to convey the fact that experiments are organic, develop, change, and yet retain a certain long-term development which makes us talk about repeating and replicating experiments . . . I think of experiments as having a life: maturing, evolving, adapting, being not only recycled, but quite literally, being retooled. (Hacking 1992, 307)

The quoted passage comes from a piece in which Hacking argued that thought experiments, unlike real experiments, do not have a “life of their own.” I want to argue here that techniques of simulation do. More specifically, some of the techniques that simulationists use to construct their models get credentialed in much the same way that Hacking says that instruments and experimental procedures and methods do; the credentials develop over an extended period of time and become deeply tradition bound. In Hacking’s language, the techniques and sets of assumptions that simulationists use become “self-vindicating.” Perhaps a better expression would be that they carry their own credentials.

By the “techniques” of simulation, I am here referring to the whole host of activities, practices, and assumptions that go into carrying out a simulation. This includes assumptions about what parameters to include or neglect, rules of thumb about how to overcome computational difficulties—what model assumptions to use, what differencing scheme to employ, what symmetries to exploit—graphical techniques for visualizing data,

8. See Winsberg 1999a or Weissert 1997, 122–124 for more details.

and techniques for comparing and calibrating simulation results to known experimental and observational data.

Whenever these techniques and assumptions are employed successfully, that is, whenever they produce results that fit well into the web of our previously accepted data, our observations, the results of our paper and pencil analyses, and our physical intuitions, whenever they make successful predictions or produce engineering accomplishments, their credibility as reliable techniques or reasonable assumptions grows.

That is, the next time simulationists build a model, the credibility of that model comes not only from the credentials supplied to it by the governing theory, but also from the antecedently established credentials of the model building techniques developed over an extended tradition of employment. That is what I mean when I say that these techniques have their own life; they carry with them their own history of prior successes and accomplishments, and, when properly used, they can bring to the table independent warrant for belief in the models they are used to build. In this respect, simulation techniques, and indeed many pre-computer calculational modeling techniques as well⁹, are much like microscopes and bubble chambers as Hacking and Galison describe them (Hacking 1988; Galison 1997).

Consider, as an example, a particular computational technique now commonly known as the “Piecewise Parabolic Method” (PPM). The PPM is an algorithm that has been shown to be well suited to simulating fluid flows that contain significant shock discontinuities, such as the example given at the beginning of this paper. Different versions of the algorithm have been used to simulate a wide variety of physical phenomena, ranging from simple laboratory setups like the one described above to such complex astrophysical systems as supernova explosions, heat convection in red-giant stars, gas accretion disks, supersonic jets, and models of the development of the entire cosmos.

The PPM begins as a discretization of the Euler equations but we should not think of it simply as a purely mathematical transformation of those equations. Recall that the primary difficulty in modeling supersonic fluids is in dealing with shock discontinuities. The difficulty arises because real fluids do not have discontinuities. In other words, the fundamental theories of fluid dynamics always describe continuous variables. Instead of discontinuities, there are very thin regions of very steep gradients.

In principle, these thin regions could be accurately modeled by including terms for viscosity and heat conduction in the principle equations. In

9. For a brief discussion of the comparative autonomy of simulation modeling techniques and such analytic modeling methods as perturbation expansions and other methods from celestial mechanics see Winsberg 1999, 288–289.

practice, however, viscous momentum transport and molecular heat transport take place on extremely small length scales. No computer algorithm that is computationally tractable could ever hope to capture effects on these scales.

The earliest approach to solving this problem was suggested by von Neumann and Richtmyer. Their solution was to artificially increase the coefficients of viscosity and heat conduction until the point at which the effects manifest themselves at length scales sufficiently large to be resolved on a reasonable computer grid. The flow can then be tweaked in just the right way to spread the shocks over a few grid cells.

The von Neumann-Richtmyer method represents the earliest attempt at overcoming these difficulties and it is relatively simple and easy to describe. Other methods that have emerged over the last fifty years have become progressively more complex and elaborate. The PPM represents the state-of-the-art method. It differs fundamentally from its predecessors in the following way. Most differencing methods are derived from Taylor series expansions of the terms in the differential equations. This move essentially assumes that the solution is smooth. While so called “shock discontinuities” are not truly discontinuous in the theory of fluids, it is nevertheless not a good assumption to treat them as smooth. So the PPM abandons this assumption, and instead of piecing together continuous, linear solutions, it pieces together discontinuous, non-linear solutions. Since superposition fails to apply to these non-linear terms, this technique then requires a special solver to compute the non-linear interactions between the discontinuous states. The construction of this solver requires outside knowledge of the propagation and interaction of non-linear waves (Winkler et al. 1987).

Exactly how all of these methods are achieved is a subject that is more than a little arcane, and the details are not that important here. What is important to note are some of the features of these methods and their development.

Perhaps most important, the presence of “shocks” in the flow of a system prevents any straightforward attempt to hammer the Euler equations into discrete form from being effective. Thus, there is nothing in the Euler equations or in the fundamental theory of fluids that tells you how to proceed if you are unable to capture the steep gradients in the flow that the theory predicts. The theory of fluids has been a useful guide in the development of these methods but neither has it come anywhere close to sufficing on its own, nor has it certified the end product.

Another interesting observation can be made by tracing the history of these methods from von Neumann’s time to the present. The history of a simulation technique is very much like the history of a scientific instrument. It begins with a relatively crude and simple technique for attacking

a relatively small set of problems. Over time, the instrument or technique is called upon to attack a larger set of problems or to achieve a higher degree of accuracy. In order to achieve this, the technique needs to be improved upon, reconfigured, and ever radically revised. In each case, the knowledge relied upon to devise and sanction the tool or method can come from a wide variety of domains.

The PPM has gained recognition as a reliable method for simulating discontinuous flows over a fairly long history of use. The results of simulations making use of this technique have been evaluated and found to be reliable in a wide variety of applications. Just like scientific instruments and experimental techniques, the PPM has “matured, evolved, been adapted and not only recycled but [not quite literally] retooled.” And just like instruments and techniques of their use, the trust that we place in the PPM as a reliable method has grown with every maturation, evolution, and retooling in which it has been successfully applied. Just like the microscope and the bubble chamber, the Piecewise Parabolic Method for calculating discontinuous compressible flows has had its own independent history. Beginning with the von Neuman-Richtmyer method, it has matured, evolved, been adapted, recycled, and retooled. It has had a life of its own.

REFERENCES

- Bate, R., J. Mueller, and D. White (1971), *Fundamentals of Astrodynamics*. New York: Dover.
- Cartwright, Nancy (1983), *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- (1999), *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Dowling, D. (1999), “Experimenting on Theories”, *Science in Context* 12 (2): 261–73.
- Franklin, Allan (1986), *The Neglect of Experiment*. Cambridge: Cambridge University Press.
- Galison, Peter (1996), “Computer Simulations and the Trading Zone”, in P. Galison and D. Stump (eds.), *The Disunity of Science: Boundaries, Contexts, and Power*. Stanford: Stanford University Press.
- (1997), *Image and Logic: A Material Culture of Microphysics*. Chicago: University of Chicago Press.
- Giere, Ronald N. (1999), *Science Without Laws*. Chicago: University of Chicago Press.
- Hacking Ian (1983), *Representing and Intervening: Introductory Topics in the Philosophy of Science*. New York: Free Press.
- (1988), “On the Stability of the Laboratory Sciences”, *The Journal of Philosophy* 85 (10): 507–15.
- (1992) “Do Thought Experiments Have a Life of Their Own?”, in A. Fine, M. Forbes, and K. Okruhlik, (eds.), *PSA 1992*, Vol. 2. East Lansing, MI: The Philosophy of Science Association, 302–310.
- Hughes, R. (1999), “The Ising Model, Computer Simulation, and Universal Physics”, in Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*. Cambridge: Cambridge University Press.
- Humphreys, Paul (1991). “Computer Simulation”, in A. Fine, M. Forbes, and L. Wessels (eds.), *PSA 1990*, Vol. 2. East Lansing MI: The Philosophy of Science Association, 497–506.
- (1994), “Numerical Experimentation”, in Paul Humphreys (ed.), *Patrick Suppes: Sci-*

- entific Philosopher*, Vol. 2 of *Philosophy of Physics, Theory Structure, Measurement Theory, Philosophy of Language, and Logic*. Dordrecht: Kluwer Academic Publishers.
- (1995), “Computational Science and Scientific Method”, *Minds and Machines* 5 (1): 499–512.
- Kaufmann, W.J., and L.L. Smarr (1993), *Supercomputing and the Transformation of Science*. New York: Scientific American Library.
- Marion, Mathieu (1998), *Wittgenstein, Finitism, and the Foundations of Mathematics*. Oxford: Clarendon Press.
- Morgan, Mary, and Margaret Morrison (eds.) (1999), *Models as Mediators*. Cambridge: Cambridge University Press.
- Newell, A. and H. A. Simon (1976), “Computer Science as Empirical Enquiry”, *Communications of the ACM* 19: 113–126.
- Norton, S., and F. Suppe (2000), “Why Atmospheric Modeling is Good Science”, in C. Miller and P. Edwards (eds.), *Changing the Atmosphere: Expert Knowledge and Environmental Governance*. Cambridge, MA: MIT Press.
- Rohrlich, Fritz (1991), “Computer Simulation in the Physical Sciences”, *PSA 1990*, Vol. 2. East Lansing, MI: The Philosophy of Science Association, 507–18.
- Weissert, T. (1997), *The Genesis of Simulation in Dynamics: Pursuing the Fermi-Pasta-Ulam Problem*. New York: Springer.
- Winkler, K., J. Chalmers, S. Hodson, P. Woodward, N. Zabusky (1987), “A Numerical Laboratory”, *Physics Today* 40 (10): 28–37.
- Winsberg, Eric (1999), “Sanctioning Models: The Epistemology of Simulation”, *Science in Context* 12(2): 275–292.
- (1999a) *Simulation and the Philosophy of Science: Computationally Intensive Studies of Complex Physical Systems*, Ph.D. Dissertation, Bloomington: Indiana University.
- (2001) “Simulations, Models, and Theories: Complex Physical Systems and Their Representations”, *Philosophy of Science* 68 (Proceedings): S442-S454.
- Woodward, P., and P. Collela (1984), “The Numerical Simulation of Two-Dimensional Gas Flows with Strong Shocks”, *The Journal of Computational Physics*. 54: 115–43.
- Zabusky, N. (1987), “Grappling with Complexity”, *Physics Today*. 40 (10): 25–27