Sismondo, Sergio. 1997. "Deflationary Metaphysics and the Construction of Laboratory Mice." *Metaphilosophy* 28:219–32.

—. 1999. "Island Biogeography and the Multiple Domains of Models." *Biology and Philosophy*, forthcoming.

Sluijs, Jeroen van der, Josée van Eijndhoven, Simon Shackley and Brian Wynne. 1998. "Anchoring Devices in Science for Policy: The Case of Consensus around Climate Sensitivity." *Social Studies of Science* 28:291–323.

Department of Philosophy
Queen's University
Kingston, Canada, K7L 3N6

DEBORAH DOWLING

Experimenting on Theories

The Argument

This paper sets out a framework for understanding how the scientific community constructs computer simulation as an epistemically and pragmatically useful methodology. The framework is based on comparisons between simulation and the loosely-defined categories of "theoretical work" and "experimental work." Within that framework, the epistemological adequacy of simulation arises from its role as a mathematical manipulation of a complex, abstract theoretical model. To establish that adequacy demands a detailed "theoretical" grasp of the internal structure of the computer program. Simultaneously, the pragmatic usefulness of simulation arises from its role as a "virtual laboratory." That role is made possible by black-boxing the internal structure of the program, such that the scientist can interact with the computer in an intuitive, "experimental" manner. Thus simulation is rendered authoritative, opening up encoded theories to a novel, "experimental" type of manipulation.

Introduction

A scientist running a computer simulation performs an experiment upon a theory.

An abstract, mathematical model of a physical system is implemented on a concrete machine. Through that machine, the model can be manipulated as if it were a physical experimental target.[1] The mathematical model can then be approached and analyzed using skills traditionally associated with experimental work: visual observation, "tinkering" with the machine, and intuition about the behavior of the concrete system.

This "new mode of doing scientific work" (Galison 1996, 137) is inseparable from the technology used to perform it: the digital computer. Without the appropriate medium it is not possible to interact with mathematical models in this way. Technological change has been accompanied by changes in scientific heuristics and practices.

---

[1] I take the phrase "experimental target" from Hacking 1992. He defines it loosely as that which (together with various instruments) constitutes "the matériel of the experiment." It is clear from the context that he uses the phrase for the physical system toward which the experimental activities of the scientist are directed — for whatever belongs at the end of the sentence "the scientist did an experiment on a . . . ."

Although this style of scientific practice has spread across almost all fields of scientific research, the methodology of simulation has rarely been discussed in the large literature describing scientists' epistemic and practical activities. With that in mind, I take a broad view of simulation as it is practiced across multiple scientific disciplines. Other articles in this issue focus on particular instances of simulation, providing more detailed insight into the local construction of the methodology in a specific setting. This study, however, focuses more on what is common to simulation-based science, offering a general framework for describing the new methodology by analogy to more traditional scientific practices.

Rather than attempting a justification or criticism of simulation-based research, this paper examines how simulation is presented as useful and authoritative by the practitioners themselves (Pickering 1992; Latour 1987). The paper draws strongly on scientists' accounts of their own practices, focusing particularly on the categories of "theory," "experiment," and "computer simulation." These categories provide a framework to elucidate how simulation is commonly located on scientists' maps of their own methodologies. Competent use of simulation draws on many typically "experimental" skills and practices, while also demanding a "theoretical" stance with respect to the abstract, mathematical computer program.

The paper is based on thirty-five interviews with scientists who use computer simulation as part of their professional research. The respondents work in a wide range of fields, including physics, meteorology, chemistry, physiology, artificial life, and computer science. It is also based on textbooks about scientific computer simulation, and on numerous recent scientific publications reporting simulation-based results from a wider range of fields, among them sociology, aerospace engineering, botany, robotics, and artificial intelligence.

A Useful Ambiguity

Peter Galison, in his historical study of "Monte Carlo" simulations, argues that this novel scientific practice is "at once nowhere and everywhere on the usual methodological map" (Galison 1996, 120). The "usual methodological map" drawn by Galison in this context has two clear landmarks: "theoretical work" and "experimental work." Much of his paper is devoted to locating simulation with respect to these two landmarks.

A similar methodological map is sketched by philosopher of science Fritz Rohrlich:

> 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. . . . Computer simulation is consequently of considerable philosophical interest. (Rohrlich 1990, 507)

The ambiguous position of simulation with regard to "theory" and "experiment" is widely acknowledged in the scientific literature and by the respondents described above. Simulation can be aligned with whichever methodological category suits the local circumstances. In this regard, computer simulation plays different roles according to the requirements of the narrative.

In one interview, for example, a biologist argued: "Look, simulation must be theoretical." That comment was made in the context of discussing the intellectual aspects of designing a simulation, and his simulation's basis in biological knowledge. In a subsequent interview, while discussing the process of observing how the system works, he took the opposite position: "I think we have to say that simulation is experimentation."

Sometimes the rhetorical positioning of simulation serves a social function (Hine 1995, 120). Labeling a research technique as "theory" or "experiment" makes a difference to its status, to publications, and to collaborations. One respondent, from the biological sciences, presented his computer as a remarkably fast and efficient theorist, to support a grant application. He hoped to structure the way the funding body conceptualized his work, suggesting that it was similar to the type of theoretical research they routinely supported. Another respondent, in physics, described the difference in social status between the "theorists" and "experimentalists" in his department. To align his work with the higher-status "theory," he drew an epistemic distinction between his own "theoretical" style of computational physics and the more "experimental" style of standard simulation. "I call my work 'simulation' on bad days," he said, on the basis that his work, like "theory," was elegant, analytical, and broadly applicable, whereas most simulation, like "experiment," was "clunky," repetitive, uncontrolled, and narrowly applicable.

As well as serving social functions, different positionings of simulation on a methodological map support different aspects of application and analysis of the technique. In contexts where the accuracy and manipulation of the encoded equations are at issue, it is appropriate to draw on simulation's function as symbolic representation, to locate it with "theory." In contexts where analogies are being drawn between the output of a simulation and the predicted behavior of an "equivalent" physical system, it is more helpful to draw on the actions associated with simulation, to locate it with "experiment."

For example, a respondent in pharmaceutical chemistry routinely compares his computer simulation with more abstract "theoretical" models of molecular systems. While describing these comparisons he argued that his simulation is similar to, but different from, "theory": "It's a theory, but it's not a fundamental theory."

Under other circumstances, the same scientist compares the results of the same simulation with experiments. As such, his simulation is also presented as similar to, but not the same as, experiment: "It's an experiment, . . . although it's a computer experiment."

Simulation is constructed as both "theoretical" and "experimental" not only in scientists' choice of language but also in the ways they design and interact with computer simulation programs. This paper describes in detail the types of arguments that scientists use to locate simulation with either "theory" or "experiment," and it describes the types of "theoretical" and "experimental" practices that simulators draw on as they construct simulation as a useful methodological hybrid.

Defining "Simulation" with Respect to "Theory" and "Experiment"

The categories "theory" and "experiment" are rarely imagined to be either mutually exclusive or clearly defined. However, they are broadly understood across the scientific community and are commonly used as landmarks in locating the methodology of "simulation."

Although many, sometimes contradictory, definitions of "theory" and "experiment" arose in interviews, two distinctions were broadly agreed on and frequently raised. The first is based on *what the scientist works with*: the "theorist" generates ? and manipulates "representations," or "descriptions," while the "experimentalist" works with "things," or "reality." The second emphasizes the *type of activity the scientist engages in*. The "experimentalist's" work involves such activities as adjusting things, testing things, and waiting for data to emerge. The "theorist's" work is characterized as the analytical, relatively predictable manipulation of equations and development of ideas.

For example, one respondent (a physiologist) defined "experiment" using the latter criterion of the *type of work involved*, then defined "theory" using the former criterion, of *working with descriptions*: "You have experimentalists and theoreticians, right? Now the experimentalist does experiments to test things. And the theorist uses mathematical descriptions."

	What the scientist works with	*The type of activity involved*
"Theory"	"Representations"	• Manipulating equations • Developing ideas
"Experiment"	"Things" or "reality"	• Fiddling with machines • Trying things out • Watching to see what happens
So ...	Simulation is like theory	Simulation is like experiment

Figure 1: The Distinction between "Theory" and "Experiment"

These two common distinctions between "theory" and "experiment" provide a useful framework for characterizing computer simulation (see figure 1). The first can be used to present simulation as "theoretical": it is a technique for manipulating "mathematical descriptions" or "representations."[2] As a programmer expressed this argument: "What you're working with in a simulation is a theory. You're computing a theoretical equation and the effects thereof." A respondent from physics made the same point more emphatically: "Of course it's theory! It's not real!"

The second distinction can be used to present simulation as "experimental": the scientist running a simulation also "does experiments to test things." The unpredictability of a simulation run makes the activity of simulation similar to the activity of performing a physical experiment. The scientist prepares the system, sets initial conditions, then takes a relatively passive role, waiting to find out how the system will respond (Pickering 1993, 23).

Analogies between simulation and this type of activity are often presented in simulation texts: "Simulation is essentially an experimental approach to solving problems" (Payne 1982, 158). Such analogies were also regularly expressed by respondents. For example a biochemist argued: "I think [computer simulation] has all the characteristics of what people would classically call an experiment. . . . Varying parameters, and just seeing . . . whether or not your output is consistent with your assumption. You do that whether or not you're on a computer, you do that whether or not you're doing it with petri dishes." In an extension of this analogy between "simulation" and "experiment," it is commonplace to present the output of a simulation as "observations," "samples," and "data."

Most respondents characterized simulation as similar to "theory" in terms of its relationship to "reality," and characterized simulation as similar to "experiment" in terms of the types of activity it entails. Simulation is thus presented as a hybrid of traditional scientific practices, facilitating "experiments" on "theories."

Strategic Black-Boxing

To "do an experiment" with on a computer simulation, it must be characterized as an opaque, unpredictable entity, with which one can interact in an "experimental" manner. Such a characterization of the machine is significantly different from presenting it as a large calculator or an electronic mathematician.

A physiologist, for example, found it appropriate to present his simulation as an experimental target similar to those he was accustomed to "prodding" in the laboratory. "You don't give the [simulation] system a problem. You prod the system and it reacts in the way it's been constructed to react."

Lucy Suchman has argued that it is the complexity of a computer program that invites the user to treat the system as a unitary entity, and to *interact* with it, the way one would interact with a person or object:

---

[2] Some respondents also used a second, more narrow sense of the word "theory," defined by the type of mathematics used. Such secondary definitions do not contribute significantly to the argument presented here.

> Insofar as the machine is somewhat predictable, in sum, and yet is also both internally opaque and liable to unanticipated behaviour, we are more likely to view ourselves as engaged in interaction with it than as just performing operations upon it, or using it as a tool to perform operations upon the world. (Suchman 1987, 16)

The results of a simulation are unpredictable. They are typically based on calculations that are analytically intractable, and the computer's numerical computation of the solutions is far too rapid to be followed by an individual scientist.

The complexity and unpredictability of a simulation program invite the scientist to present it as a black box, interacting with it in an experimental manner, trying things out to see what will happen. A respondent from chemistry argued: "It's experimental in the sense that when you set up your simulation you're waiting for the result, and when the result comes out, it can genuinely be something that you don't expect — because of this remoteness of the computer processes." The computer processes are "remote": they are not, and cannot be, grasped or followed by a scientist during a simulation run.

However, this black-boxing of the simulation program differs from the opacity of many nonscientific computer programs,[3] in that the black-boxing is only *temporary*. To analyze the manipulations of the "theory" performed by the computer, the technology must be characterized less as an opaque, interactive entity and more as a transparent calculating machine.

While it is pragmatically convenient to suspend interest in the internal mechanisms of a working simulation, it is professionally reprehensible to be ignorant of those workings. Both designers and users of simulations argue that a competent scientist should have a clear, analytic grasp of the mathematics built into the program. A simulation programmer exclaimed, "people should get into the code and say OK, what is really going on here? I mean it's just crazy not to do that." Similar expectations were found in interviews across all the fields sampled, with researchers strongly aware that "these methods are not foolproof. They are still in development. To really use them efficiently I think you have to know what they are doing, what are the strengths, what are the approximations, what are the difficult points, the pits you can fall into." A pharmaceutical chemist argued emphatically that a responsible user of simulation must "try to break down that sense that, because this number came out on the printed form, it's got to be right. You have to constantly check your premises and estimate things on the back of an envelope to make sure." Opening the black box in this way is necessary to the establishment of a simulation's legitimacy. It is the code inside the box that determines the authority with which the program can claim to mimic an experimental target.

The process of computer simulation requires both an analytical understanding of the mathematical principles programmed into the machine, and a temporary suspension of interest in those principles, in order to interact with the computer as if it were a black box.

why?

These two modes of interaction with computer simulation are reminiscent of the modes of human-computer interaction identified by Sherry Turkle. Turkle argues that there are two primary "modes" of interacting with computers: the "calculation" mode, where one is conscious of, and interacting with, the internal mechanisms behind the machine's outward behavior, and the "simulation" mode, where one negotiates the interface of the computer, without paying attention to its underlying programming (Turkle 1995, 36–43). The competent professional use of scientific computer simulation demands both levels of interaction. It requires knowledge of the mathematical principles driving the machine's behavior, *and* it requires temporary, strategic suspension of interest in those principles, in order to treat the computer as an interactive entity and perform an "experiment" on it.

Combining "Theoretical" and "Experimental" Aspects

This temporary black-boxing of the computations involved in simulation allows scientists to interact with theories as if they were entities that could be adjusted, observed, and measured. Thus "theoretical" projects, manipulations of representations, can be approached through "experimental" procedures.

On the basis of simulation's role as a manipulation of a representation, many applications of simulation are proposed and analyzed as theoretical projects. Simulations are often compared to experimental results, in a Popperian check of the "validity of the representation," or to predict the behavior of "represented" systems. Analysis of simulation results draws attention to such traditionally "theoretical" issues as the underdetermination of the model by the available data, or the degree to which the representation is "biased" (Taylor 1989; De Landa 1994; Oreskes et al. 1994; Suppe 1996).

What is novel about these "theoretical" projects is the "experimental" procedures used to tackle them. The integration of theory and experiment can be seen in this excerpt from an article in the *Journal of Computational Physics*. It describes interactions with a particular simulation.

> If one adds a damping term, the oscillation decays. However, if one switches on a constant bias (that is, a constant is added to the right side of [an equation]) for an appropriate time interval, a tunable periodic oscillation is obtained. (Zabusky 1981, 220)

---

[3] In commercially available software, see for example Winograd and Flores 1987, 5: "'What is a word processor?'. . . It is a medium for the creation and modification of linguistic structures that play a role in human communication. For the purchaser of the word processor, this is the relevant domain. The word processor exists as a collection of hardware or programs only when it breaks down." In education, see for example Cabrera 1990, 82: "Students can then investigate the structure of these physical models in an interactive environment, even though the mathematical sophistication of the models is often beyond their grasp." See also the literature on virtual reality — e.g., Sherman and Judkins 1992; Woolley 1992; Heim 1993; Casti 1997; Morse 1994.

There is a strong sense of experiment in the description of this procedure. Things are "added" or switched on and off, and responses are recorded. Given that the oscillation is presented as a graphic on a computer screen, and that the mechanisms behind the behavior of the oscillation occur too fast to be observed, there is a strong sense of trial and observation. However the things that are added or switched are not chemicals or currents: they are mathematical terms and constants. And the process of switching on "for an appropriate time interval" can be translated into the mathematical addition of a constant to an equation.

Another example of this experimental approach to traditionally "theoretical" projects is the mathematical work of Walter Fontana, on complexity. Fontana's computer program chooses mathematical functions at random and applies them to one another. (For example the function *f[g]* is "plugged in" to the function *h[f]*, thus transforming *f* [according to *h*] into a new function.) Fontana carefully constructs an analogy between his program and a physical system, presenting his work as practical experimentation with a concrete, unpredictable population of entities.

The work that goes into constructing that metaphor is not empty poetry. The metaphor helps to "set up a model that also provides a workbench for experimentation" (Fontana 1991, 166). The "results and discussion" section of Fontana's work bears little resemblance to a mathematical paper. His results are presented in terms of observations of the system's behavior, such as: "Most of the computer experiments exhibit very complicated short-lived states that reduce to simpler cooperative metastable transients"; and, "Polymers consisting of monomers nested into one single tree branch have been observed" (ibid., 186, 193). In this way the analytical process of manipulating a mathematical model can be approached through a negotiation with the computer interface. Mathematical calculation and experimental practice come together in an unprecedented form.

Tinkering, Noticing, and Intuition

Fontana, Zabusky, and other scientists can cross the boundaries of figure 1 to manipulate "representations" using "experimental" practices, as a result of the strategic, partial black-boxing of their computer simulations. The program is "open" at a level through which a general analytical grasp of the theory is possible. At another level the program is sufficiently "closed" that (as Papert and Turkle argue) the abstract mathematical theory can be manipulated as if it were a concrete, physical object (Lévi-Strauss 1972; Turkle and Papert 1991; Wilensky 1991):

> The computer stands betwixt and between the world of formal systems and physical things; it has the ability to make the abstract concrete. In the simplest case, an object moving on a computer screen might be defined by the most formal of rules and so be like a construct in pure mathematics; but at the same time it is visible, almost tangible, and allows a sense of direct manipulation. (Turkle and Papert 1991, 162)

A sense of direct manipulation encourages simulators to develop a "feel" for their mathematical models with their hands and their eyes, by *tinkering* with them, *noticing* how they behave, and developing a practical *intuition* for how they work. Because the simulation is presented as an experimental target, the researcher can interact with it as if it were a "real" target, drawing on the physical skills of recognition and reaction.

The quick generation of data through simulation allows scientists to patch ideas together, to explore possibilities, to build up a broad base of experiences, and thus to develop a more intuitive, "hands-on" understanding of a mathematical model. Many formal applications of simulation use a systematic process of parameter variation in order to "explore the solution space" (Kaufmann and Smarr 1993, 4–7). The author of a journal article on simulation in physics notes:

> It is clear that interactive and rapid turn-around computing provides an opportunity to concentrate deeply and develop a special intuitive "feel" for the results. This noninterrupted mode augments the innovative process. This is also easy to achieve for systems with a few dependent variables. A similar remark also applies to processing and visualizing the data base obtained from long-duration runs of large-scale simulations. (Zabusky 1981, 232)

A large element of skill and tacit knowledge is involved in developing this "intuitive feel" for a computer simulation. The same author claims that he has "found it difficult to relate this mode of working via lectures. Perhaps this mode is still an art form understood by committed practitioners in benign computer environments and learned only by apprenticeship" (ibid., 196).

A respondent from inorganic chemistry made a similar claim about the way he used his own simulation:

> I wouldn't say at the moment that simulation is strictly a science. It's sort of more like an art. . . . It's just, you have to have a feel. In this case, the reason why things are not coming to equilibrium is because the run isn't long enough or it's not short enough, or I haven't had enough interchange attempts. It's something you pick up, I suppose, with experience.

An important part of the exploration of a simulation's behavior is the typically "experimental" skill of noticing visual patterns generated by an inscription device. As Hacking has argued, "The good experimenter is often the observant one who sees the instructive quirks or unexpected outcomes of this or that bit of the equipment" (Hacking 1983, 167). The good simulator requires a similar ability.

Most computer simulations are equipped with data analysis routines and graphic interfaces, so that the output appears in the form of tables, graphs, and diagrams that are intended to be as intuitively clear to the user as possible. This way, the scientist can make intuitive sense of the program's behavior, and be alert to patterns and oddities. A biochemist, for example, can interpret the output of a computer program in terms of the three-dimensional shape of a molecule: "You can see things with different representations that you can't just see out of the numbers. So you can see if the DNA is bent." Similarly, graphic presentation can allow a physicist to interpret a series of mathematical results as a moving wave: "In a film depicting the time evolution of an initially sinusoidal solution of the KdV equation, made by Gary Deem at Bell Laboratories under the direction of Kruskal and Zabusky, there are eight clearly distinguishable 'solitons,' each pursuing its own path" (Lax 1975, quoted in Zabusky 1981, 227).

These skills of tinkering and noticing have become important practices in the development of scientific theory. Unstructured interaction with the computer builds a base of experience of a mathematical model's behavior, which can strengthen a general understanding of the model or suggest new directions for research. As one respondent put it: "Simulation is useful when you don't know what to do next, and you just want to mess around and see what happens, and look for something unexpected and interesting."

> Interplay between simulation results and intuitions is the core of the process of model invention, and occupies the lion's share of the development time. ... We have often found it to be useful to pick a single point in the parameter space, generate predictions (perhaps somewhat noisy), think about the results, choose another single point, etc., thereby learning not anything about a best fit but something about the logic of the model. (Shiffrin and Nobel 1996, 4, 5)

A more directed exploration of such a semi-opaque system can generate "data sets," "observations," and "discoveries" regarding the system's behavior. A neurophysiologist related the discoveries he had made by "experimenting" with different values for particular parameters in his simulation:

> For example when I started mucking around with conduction velocities, which is the rate the signal gets from one cell to the next, I had expected that would be a very critical parameter. It turns out to have some role, but very very little in relationship to things that I thought it would be critical for. And that was something of a shock. On the other hand I hadn't thought the threshold would be important, but when we started mucking with that we found that that actually produces more variation than varying the conductance velocity. Which is again a surprise.

Experimental practices such as "mucking around" and noticing surprising aspects of a computer's behavior can generate important changes in scientists' understanding of the systems they study.

Conclusion

Comparisons of scientific computer simulation with "theory" and "experiment" provide a convenient framework for examining the construction of simulation-based methodology. Within that framework, computer simulation is presented as useful and authoritative insofar as it opens an epistemically justified "theory" to skilled "experimental" manipulation. The authority of scientific simulation relies on an analytical knowledge of the underlying "theoretical" mathematics, while the pragmatic utility of simulation relies on the scientist's ability to black-box the program, to delegate the computations to the machine and to interact "experimentally" with the surface.

Many authors in the social studies of science emphasize the close interrelationships between technological and epistemic concerns in scientific practice, showing how "articulated knowledge and machinic performances are reciprocally tuned to one another" (Pickering 1995, 29; see also Clarke and Fujimura 1992). In the scientific use of computer simulation, an abstract model is manipulated through a digital machine, entwining the epistemic and the technical inextricably together. This technique of embodying a representation in a machine has facilitated a "new method for extracting information from physical measurements and equations" (Galison 1996, 120).

By combining an analytical grasp of a mathematical model with the ability to temporarily "black-box" the digital manipulation of that model, the technique of simulation allows creative and experimental "playing around" with an otherwise impenetrable set of equations, to notice its quirks or unexpected outcomes. The results of a large and complex set of computations are thus presented in a way that brings the skills of an observant experimenter to the development of mathematical theory.

The usefulness of computer simulation thus depends on the construction and maintenance of this methodological ambiguity. In their everyday interactions with the computer, and in their choice of language in varied narrative contexts, scientists strategically manage simulation's flexible position with respect to "theory" and "experiment." In this way they facilitate a significantly novel, and highly productive, mode of scientific work: creative experimentation with a mathematical model.

References

Cabrera, B. 1990. "Early Experiences with Physics Simulations in the Classroom." In *Computers in Physics Instruction*, 77–83. Raleigh: Addison-Wesley.

Casti, J. L. 1997. *Would-be Worlds*. New York: Wiley.

Clarke, A., and J. Fujimura, eds. 1992. *The Right Tools for the Job: At Work in Twentieth-Century Life Sciences*. Princeton, N.J.: Princeton University Press.

De Landa, M. 1994. "Virtual Environments and the Emergence of Synthetic Reason." In *Flame Wars: The Discourse of Cyberspace*, edited by M. Dery, 263–85. Durham, N.C.: Duke University Press.

Dowling, D. 1998. "Experiments on Theories: The Construction of Scientific Computer Simulation." Ph.D. diss., University of Melbourne. *Dissertation Abstracts International* 59:9A.

Fontana, W. 1991. "Algorithmic Chemistry." In *Artificial Life II: Proceedings of the Workshop on Artificial Life Held February 1990 in Santa Fe, New Mexico*, edited by C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, 159–209. Redwood City, Calif.: Addison-Wesley.

Galison, P. 1996. "Computer Simulations and the Trading Zone." In *The Disunity of Science: Boundaries, Contexts, and Power*, edited by P. Galison and D. J. Stump, 118–57. Stanford, Calif.: Stanford University Press.

Hacking, I. 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.

——. 1992. "The Self-Vindication of the Laboratory Sciences." In *Science as Practice and Culture*, edited by A. Pickering, 29–64. Chicago: University of Chicago Press.

Heim, M. 1993. *The Metaphysics of Virtual Reality*. New York: Oxford University Press.

Hine, C. 1995. "Information Technology as an Instrument of Genetics." *The Genetic Engineer and Biotechnologist* 15(2–3):113–24.

Kaufmann, W. J. I., and L. L. Smarr. 1993. *Supercomputing and the Transformation of Science*. New York: Scientific American Library.

Latour, B. 1987. *Science in Action*. Milton Keynes: Open University Press.

Lévi-Strauss, C. 1972. *The Savage Mind*. London: Weidenfeld and Nicolson.

Morse, M. 1994. "What Do Cyborgs Eat? Oral Logic in an Information Society." In *Culture on the Brink: Ideologies of Technology*, edited by G. Bender and T. Druckrey, 157–89. Seattle: Bay Press.

Oreskes, N., et al. 1994. "Verification, Validation and Confirmation of Numerical Models in the Earth Sciences." *Science* 263:641–46.

Payne, J. A. 1982. *Introduction to Simulation: Programming Techniques and Methods of Analysis*. New York: McGraw-Hill.

Pickering, A., ed. 1992. *Science as Practice and Culture*. Chicago: University of Chicago Press.

——. 1993. "The Mangle of Practice: Agency and Emergence in the Sociology of Science." *American Journal of Sociology* 99(3):559–89.

——. 1995. *The Mangle of Practice: Time, Agency and Science*. Chicago: University of Chicago Press.

Rohrlich, F. 1990. "Computer Simulation in the Physical Sciences." *Philosophy of Science Association* 2:507–18.

Sherman, B., and P. Judkins. 1992. *Glimpses of Heaven, Visions of Hell: Virtual Reality and Its Implications*. London: Hodder and Stoughton.

Shiffrin, R. M., and P. A. Nobel. 1996. *The Art of Model Development and Testing*. Report 167, Research Report Series, Cognitive Science Department, Indiana University.

Suchman, L. 1987. *Plans and Situated Actions: The Problem of Human Machine Communication*. Cambridge: Cambridge University Press.

Suppe, F. 1996. "The Epistemology of Simulation Modeling" — draft copy.

Taylor, P. 1989. "Revising Models and Generating Theory." *Oikos* 54(1):121–26.

Turkle, S. 1995. *Life on the Screen: Identity in the Age of the Internet*. New York: Simon & Schuster.

Turkle, S., and S. Papert 1991. "Epistemological Pluralism and the Re-valuation of the Concrete." In *Constructivism: Research Reports and Essays 1985–1990*, edited by I. Harel and S. Papert, 161–91. Norwood, N.J.: Ablex.

Wilensky, U. 1991. "Abstract Meditations on the Concrete and Concrete Implications for Mathematics Education." In *Constructivism: Research Reports and Essays 1985–1990*, edited by I. Harel and S. Papert, 193–203. Norwood, N.J.: Ablex.

Winograd, T., and F. Flores 1987. *Understanding Computers and Cognition: A New Foundation for Design*. Reading, Mass.: Addison-Wesley.

Woolley, B. 1992. *Virtual Worlds: A Journey in Hype and Hyperreality*. Oxford: Blackwell.

Zabusky, N. 1981. "Computational Synergetics and Mathematical Innovation." *Journal of Computational Physics* 43:195–249.

Department of History and Philosophy of Science
University of Melbourne