Evelyn Fox Keller

10 / Models, Simulation, and "Computer Experiments"

> As soon as an Analytical Engine exists, it will necessarily guide the future course of the science.
>
> Charles Babbage, *Passages from the Life of a Philosopher,* 1864

1. Introduction

As Nelson Goodman famously observed, "Few terms are used in popular and scientific discourse more promiscuously than 'model'" (Goodman 1968, 171). Writing more than thirty years later, much the same might be said of the term "simulation." Yet this was not always the case. Both words have ancient histories, but until very recently, the meaning of "simulation," at least, was manifestly stable: it invariably implied deceit. Usages offered by the Oxford English Dictionary (OED) prior to 1947 include "false pretence"; "A Deceiving by Actions, Gestures, or Behaviour" (1692); "a Pretence of what is not" (1711). Evidence provided by the OED, in short, suggests that it was only after World War II that the word took on the meaning that brings it into its current proximity with models: "The technique of imitating the behaviour of some situation or process . . . by means of a suitably analogous situation or apparatus, especially for the purpose of study or personnel training." Here, the valence of the term changes decisively: now productive rather than merely deceptive,[1] and, in particular, designating a technique for the promotion of scientific understanding. The shift

reflects a crucial change not only in the perceived value of simulation but also, as others have already noted, in the means of production of scientific knowledge (see, for example, Rohrlich 1991; Humphreys 1991; Galison 1996; Winsberg 1999). Furthermore, it is this new sense of the term that encourages its use in much of the current historical and philosophical literature as either interchangeable with the term "model," or as one part of a single composite noun (as in "models and simulations"). An obvious question arises, however, and it is this: Do the actual uses of simulation in contemporary scientific practice in fact warrant such facile assimilation? Or, to pose the question somewhat differently, does the use of simulation in post–World War II science add significantly new features to the range of practices that had earlier been subsumed under the term "modeling"? My answer is yes, but I argue that the novelty has been multilayered and cumulative in its effects, requiring a more nuanced history than has yet been made available.

The rise of simulation in post–World War II science is not exclusively associated with the advent of the computer—in fact, the earliest invocations of the term relied primarily on the use of electrical and electronic analogue devices designed to mimic the behavior of real-world phenomena.[2] However, it was the introduction of the digital computer that provided the major impetus for the adoption of simulation techniques in scientific research, and for that reason, my discussion will be confined to what has come to be known as "computer simulation." Very crudely, it might be said that the immediate effect of these new techniques on scientific practice was to radically extend the range of problems amenable to quantitative analysis. They did so in a variety of ways, however, and with widely varying implications. Indeed, even the term "computer simulation" covers so complex a range of activities that some sort of taxonomy would seem to be in order. What kind of taxonomy? We might start with a division along disciplinary lines—distinguishing the uses of simulation in the physical sciences from those in the biological sciences, in cognitive science, in economics or management. But proceeding with a canonical evolutionary tree will clearly not serve, for such a structure misses the cross-structures needed for and resulting from ongoing hybridization. On the other hand, differences in aims, interests, and tradition are plainly evident, and they bear critically on subsequent historical developments. I suggest, therefore, that it is useful to follow the history of simulation along quasi-disciplinary lines while at the same time remaining alert to the extensive cross-disciplinary traffic of technical innovations that has been so much a part of this history. In this chapter, I focus primarily on the physical sciences, and I argue that, even within such a major disciplinary category, important subdivisions need to be demarcated. Furthermore, each of these subdivisions pushes the boundaries of the initial disciplinary divide in distinctive ways.[3]

2. Computer Simulation in the Physical Sciences

In one of the first attempts to bring the novel features of computer simulation to the attention of philosophers of science, the physicist Fritz Rohrlich put forth the claim for "a qualitatively new and different methodology" lying "somewhere intermediate between traditional theoretical physical science and its empirical methods of experimentation and observation. In many cases," he wrote, "it involves a *new syntax* which gradually replaces the old, and it involves *theoretical model experimentation* in a qualitatively new and interesting way. Scientific activity has thus reached a new milestone somewhat comparable to the milestones that started the empirical approach (Galileo) and the deterministic mathematical approach to dynamics (the old syntax of Newton and Laplace)" (Rohrlich 1991, 907). Others have argued in a similar vein. Peter Galison (1996), for example, draws a sharp distinction between the new computer simulations and the earlier analogue simulations. He suggests that, while the latter can be readily assimilated into a long history of analogue models (including not only ship models and wind tunnels but also nineteenth-century models built out of "pulleys, springs, and rotors to recreate the relations embodied in electromagnetism" [Galison 1996, 121]), the new techniques of computer simulation effected a radical epistemological transformation in the physical sciences, ushering "physics into a place paradoxically dislocated from the traditional reality that borrowed from both experimental and theoretical domains" and creating a "netherland that was at once nowhere and everywhere on the methodological map" (Galison 1996, 120). Such claims have become familiar, and they are generally taken to be uncontroversial. Over the last half a century, a new domain of physical science has come into being that is widely recognized as different from the older domains of both theoretical and experimental physics, and that has accordingly warranted a new designation, namely "computational physics." "Computational physics" is simply a term referring to the use of computer simulation in the analysis of complex physical systems, and, as such, it is unquestionably both new and distinctive.

Controversy arises only in response to the question, What exactly is it that is so distinctive about this new endeavor? In one sense, the answer is obvious: computer simulation opened up the study of complex systems—that is, it brought a range of phenomena that had hitherto been mathematically intractable into analytic reach. Until the advent of computers, the primary tool physicists had at their disposal for representing their theoretical understanding of the mechanics and dynamics of material systems had been the differential equation, and their principal task was to relate the solutions of these equations to observed experi-

mental effects. But differential equations are notoriously difficult to solve once they depart from the linear domain, and especially so when representing the interactions of many bodies. Thus, prior to the computer, the study of complex, nonlinear phenomena by physicists had been limited to what could be achieved by perturbation methods, simplifying models/approximations (for example, "effective-field" approximations), or paper-and-pencil schemes for numerical approximation,[4] and the first and most obvious use of "computer simulations" (in the widest sense of the term) was to provide mechanized schemes of calculation that vastly expanded the reach of available methods of analysis.[5] While computers were not capable of giving exact solutions of the equations already provided—either by established theoretical principles or by the various models that had been developed to make these principles more tractable—they could give approximate solutions to high degrees of accuracy, and with astonishing rapidity, and this capacity in itself clearly transformed both the domain and the practice of physical science.

Yet even so, there remains the question of epistemological novelty. How does the availability of high-speed computation qualitatively alter the epistemic character of what numerical analysts had already been doing, albeit on a manifestly smaller and slower scale? Indeed, in what sense can these early computational schemes be said to be *simulations?* What are they simulations *of?* Finally, although it might be easy to see how numerical analysis falls outside the range of what is conventionally regarded as "theory," what is it about these techniques that brings them into the domain of "experiment"? Taking these questions in order, I will argue that what we have now come to see as the epistemological novelty of computer simulation in fact emerged only gradually—not as a consequence of the introduction of any single technique, but as the cumulative effect of an ever-expanding and conspicuously malleable new technology: a technology that may originally have been designed to meet existing needs but that was, from its inception, already generating new opportunities and new needs. Just as with all the other ways in which the computer has changed and continues to change our lives, so too, in the use of computer simulations—and probably in the very meaning of science—we can but dimly see, and certainly only begin to describe, the ways in which exploitation of and growing reliance on these opportunities changes our experience, our sciences, our very minds.

Computer simulation may have started out as little more than a mechanical extension of conventional methods of numerical analysis, where what was being "simulated" were the precomputer, handwritten equations and where the early deprecatory sense of the term was still very much in place, but such methods rapidly grew so effective that they began to challenge the status of the original,

soon threatening to displace the very equations they were designed to simulate. Over the course of time, evolving practices of computer simulation generated qualitatively different ways of doing science in which the meaning as well as the site of "theory," of "modeling," and eventually of "experiment" and "data" all came in for similar dislocations: Simulation came to lose its earlier sense of ontological inferiority, its status of "pretender," but also its sense of epistemological inferiority, at first nothing more than a mechanization of the lowliest form of scientific work, numerical computation. Paraphrasing Galison (1996, 119), we might agree that simulation eventually came to constitute "an alternate reality." Yet no single technical innovation can be held responsible. The transformation to which Galison and others refer emerged out of the collective and cumulative successes of many different effects in which new technical developments built on older ones in ways that might look seamless from afar but that, upon closer inspection, reveal a number of several more-or-less distinct stages (or branches), each bearing its own marks of epistemological novelty and its own disturbances to traditional notions of "theory," "experiment," and "data."

Provisionally, I suggest three such stages: (1) the use of the computer to extract solutions from prespecified but mathematically intractable sets of equations by means of either conventional or novel methods of numerical analysis; (2) the use of the computer to follow the dynamics of systems of idealized particles ("computer experiments") in order to identify the salient features required for physically realistic approximations (or models); (3) the construction of models (theoretical and/or "practical") of phenomena for which no general theory exists and for which only rudimentary indications of the underlying dynamics of interaction are available. With the growing success of these practices, use of the new techniques (as well as reliance upon them) increased steadily, inevitably enhancing the perceived epistemological and even ontological value of the simulation in question. But the originals whose privileged status was thereby threatened, and that may even have been put at risk of being supplanted by the simulation, were of widely different kinds in these three different practices. What is most directly called into question by the first case is the traditional status of the differential equation as the primary tool of theoretical physics; in the second, it is the nature of modeling and its relation to the construction of theory; in the third, it is both the meaning and the goals of explanation that come in for transformation. Interestingly, the roots of all these practices, as well as the first invocation of computers as "experimental" tools, can be found in the work of the mathematician Stanislaw Ulam (1909–1984) at the Los Alamos National Laboratories.

2.1. Computers as Heuristic Aid: "Experiments in Theory"

The immediate impetus for the development of computer simulation techniques came from research at Los Alamos from 1946 to 1952 on the feasibility of various proposals for building effective thermonuclear weapons.[6] Here, the first and foremost need was to bypass the mathematical intractability of equations conventionally used to describe the highly nonlinear phenomena that were involved (for example, neutron diffusion, shock waves, and "multiplicative" or branching reactions). Ulam, working with Von Neumann, Fermi, and others, originated a number of novel approaches to existing computational procedures that have since become staples in the analysis of complex systems. Ulam's main contributions were not, however, dependent on the computer, but rather on the deployment of methods of combinatorial analysis and statistical sampling for exploring the solution space of conventional differential equations.

The most famous of these, the Monte Carlo method, has been extensively discussed by Galison (1996), and its introduction (Richtmyer and Von Neumann 1947; Metropolis and Ulam 1949) is sometimes taken as synonymous with the origin of simulation (see, for example, Mize and Cox 1968, 1). In point of fact, however, the epistemological novelty of Monte Carlo (at least as it was first introduced) had little to do with the computer. Its application to differential equations depended on the formal isomorphism of such equations with certain equations in probability theory. The first novelty of the method lay in inverting the customary use of that relation (that is, in exploiting the probability relations to solve the differential equations rather than using the differential equations to analyze the probability relations); the second lay in replacing the computation of the combinatorial possibilities for all sequences of events (or individual trajectories) by estimates of successful outcomes obtained by sampling a number of different "experimental" trajectories, trials, or "games." As Ulam wrote, "Given a partial differential equation, we construct models of suitable games, and obtain distributions on solutions of the corresponding equations by playing those games, i.e., by experiment" (Ulam 1952, 267). The method's first published use was in "solving" prespecified (Boltzmann-type) equations for neutron diffusion. Ulam explained the procedure as follows: "[D]ata are chosen at random to represent a number of neutrons in a chain-reacting system. The history of these neutrons and their progeny is determined by detailed calculations of the motions and collisions of these neutrons, randomly chosen variables being introduced at certain points in such a way as to represent the occurrence of various processes with the correct probabilities. If the history is followed far enough, the chain reaction thus represented may be regarded as a representative sample

of a chain reaction in the system in question. The results may be analyzed statistically to obtain various average quantities of interest for comparison with experiments or for design problems" (Ulam 1990, 17).

What lent the computer its importance to this application was simply its ability to perform the required "detailed calculations" on a scale and at a speed exceeding anything that could have been done by hand or by other mechanical devices. In other words, the difference made by the computer in the application of this method was identical to that which it made to more conventional kinds of numerical analysis. It was because of its speed (leading in turn to its versatility) that Ulam described the electronic computer as enabling "one to make, as it were, 'experiments in theory' by computations, either of the classical or of the 'Monte Carlo' type computations" (1990, 122). What did he mean by "experiments in theory"? I submit that, for Ulam, computer simulations were "experimental" in the same sense in which a thought experiment was "experimental" (or in which repeated games of chance were "experimental"), different only in that the computer permitted the working out of the implications of a hypothesis so rapidly as to rival the speed of thought and was certainly vastly faster than any of the traditional means of computation that had been available. They extended the mathematician's powers of analysis and, as such, ought to have been as valuable for solving problems in pure mathematics as in mathematical physics. In no sense were they to be confused with actual experiments ("experiments in practice"?) on which confirmation of theory depended. Nor were such simulations to be confused with the "design problems" posed by the physical materials with which the engineers were working—problems that had perforce to be solved before a device with real explosive power could be detonated. When it came to building actual bombs, nowhere is there any evidence of either confusion or slippage between simulation and the real thing.[7]

As is perhaps inevitable, such "experiments in theory"—especially as they proved ever more successful—began to take on a life of their own. Indeed, their very success brought conspicuous pressure to bear on the primacy of (or need for) more conventional mathematical tools, most notably, the differential equation. Similarly, that success added new legitimacy to the practice of numerical analysis that had been around for so long. But, while both these effects can claim certain kinds of epistemological novelty (for example, challenging both the hegemony and the realism of continuous variable representations and demanding "an empirical epistemology, and not merely a mathematico-deductive one" [Winsberg 1999, 290]),[8] neither, it seems to me, yet alters the perception of the basic aim of matching theoretical predictions with experimental findings. Computer simulation is at this stage still directed toward eliciting the implications of well-formulated theoretical models.

2.2. "Computer Experiments" in Molecular Dynamics

A rather different sense of "experiment"—one that is noticeably closer to physicists' understanding of the term than to that of a mathematician—enters the literature on computer simulations in the mid to late 1950s. Where, in the earlier sense of the term, what was to be simulated were the equations of the traditional theoretical physicist, and the aim was to obtain approximate solutions of these equations, in the new practice of simulation it was an idealized version of the physical system that was to be simulated, the aim of which was to produce equations (or models) that would be both physically realistic and computationally tractable.[9] In fact, the practice of "computer experiments" (as such techniques soon came to be called) deployed two levels of simulation: first, substitution of the actual physical system by an "artificial" system, and, second, replacement of the equations to which the first level of simulation gave rise by computationally manageable schemes for numerical analysis. Thus, they were "experimental" in two senses of the term—not only in Ulam's sense of "experiments in theory" but also in the sense we might call "experiments in modeling"—and they were aimed at redressing theoretical intractability on two corresponding levels: descriptive and computational. As two of the early advocates of this approach described it, "This half-way house between elegant theory and experimental hardware, our programmed version of the physical laws and boundary conditions, we call a "computer experiment." It differs from a typical computation of a theoretical result in that we do not evaluate mathematical expressions derived from the laws of nature, but we make the computer simulate the physical system" (Buneman and Dunn 1965, 4).

To understand the need out of which this new use of simulation arose, we need to recall the state of "theory" in the mid-1950s for the macroscopic (thermodynamic) properties of liquids, gases, and solids. "Theory," in this context, means statistical mechanics, and its aim is to derive the equilibrium and nonequilibrium behavior of many-body systems from the molecular dynamics of the component particles, and the obvious problem is how to deal with so large a number of particles. Clearly, some form of simplification and successive approximation is required. For example, one might begin by ignoring all interactions between particles (the ideal gas approximation); a next step, and slight improvement, would be to treat the particles as quasi-independent, each moving in some average potential due to all the other particles in the system. But neither of these approximations is adequate at high density, nor for describing the phenomena of greatest interest, namely phase transitions; for this, one needs a more realistic representation of the effects of molecular interactions. The crucial next step (often said to be the origin of modern liquid theory) was taken in 1935 when J. G.

Kirkwood rewrote the equations in terms of pairs of particles moving in an effective potential (due to all other particles). Thus rewritten, one needed only two functions: the intermolecular pair potential and the radial (or pair) distribution function; yet, without knowledge of the interatomic or intermolecular forces, neither of these functions could be specified. Various models for the pair potential were available (for example, hard spheres; hard spheres plus square well; Lennard-Jones), but until the mid-1950s, the only available access to the radial distribution function (representing the distribution of distances between atoms or molecules) was from X-ray or neutron diffraction patterns in simple fluids. Such measurements were not only cumbersome (having to be redone for each change in density or temperature) but also limited to finite ranges of frequency, and even in those ranges dependent on data that were often fraught with ambiguities. Finally, and perhaps most important, such empirical derivations were theoretically unsatisfying in that they offered no insight at all into the molecular dynamics responsible for the shape or behavior of the function.

Computers enter the history of this field with the development of an alternative approach to the problem by two physicists working at the Livermore National Labs, Berni Alder and Ted Wainwright, in the late 1950s. Building on the Monte Carlo computations of N-body systems that the Los Alamos group had pioneered, and using the high-speed computers available at Livermore, Alder and Wainwright were able to follow the behavior of systems of a finite number of particles (ranging between 32 and 500) idealized as hard spheres under conditions of varying density and temperature. As they wrote, "With fast electronic computers it is possible to set up artificial many-body systems with interactions which are both simple and exactly known. Experiments with such a system can yield not only the equilibrium and transport properties at any arbitrary density and temperature of the system, but also any much more detailed information desired. With these 'controlled' experiments in simple systems it is then possible to narrow down the problem as to what analytical scheme best approximates the many-body correlations" (Wainwright and Alder 1958, 116). Others—both in studies of classical fluids and in plasma physics—soon picked up on the method, as well as on the nomenclature, and ran with it. In the first of a series of papers on computer experiments, in which Alder and Wainwright's approach was extended to systems of particles interacting through a Lennard-Jones potential, Loup Verlet explained, "The 'exact' machine computations relative to classical fluids have several aims: It is possible to realize 'experiments' in which the intermolecular forces are known; approximate theories can thus be unambiguously tested and some guidelines are provided to build such theories whenever they do not exist. The comparison of the results of such computations

with real experiments is the best way to obtain insight into the interaction between molecules in the high-density states" (Verlet 1967, 98).

Here, the aim of computer experiments is clearly stated: to test "approximate theories" where they exist and to provide guidelines for building such theories where they do not. Here, also, use of the word "exact" highlights the difference between the Monte Carlo simulations initiated at Los Alamos for the purpose of computation and the new methods: in contrast to the "exact" machine computations, it is the theories (or models) to be computed by these methods that are now acknowledged as "approximate." In this shift of usage, the purpose of computation (conventionally associated with the application of a theory) is tacitly subordinated to another kind of aim—namely, that of building theory—and it is in pursuit of this latter aim that the simulation (or "artificial system") serves as an "experimental" probe. The simulation is a trial theory, and the role of machine computation is to render the test of that theory "unambiguous." As the final sentence (in the passage quoted above) makes clear, however, the ultimate power of arbitration was still seen as residing in "real experiments," especially in the real experiments that provided measurements of macroscopic properties.

That physicists were well aware of the irregularity of this use of the term "experiment"—and equally of the threats such usage raised for traditional understandings of "theory" and "experiment"—is well attested to by the recurrence of discussions of the matter to be found in the literature throughout the 1960s. In one of the earliest such discussions, the authors, Buneman and Dunn, focus their attention on the relation between "computer experiments" and "theory." They begin by observing, "we are at the threshold of a new era of research." Computer experiments "yield surprising and significant answers"; they permit the deduction of "a qualitative or even an analytic theory" and allow one to "guess what are the significant effects and what is the correct way of looking at a problem" (Buneman and Dunn 1965, 56). In just the brief time they'd been around, use of this new mode of analysis had already increased dramatically. From a retrospective search of the literature performed in 1966, another observer, Charles Birdsall, estimated the growth rate of articles on computer experiments over the years since 1956 by an exponential factor of 1/3. The rapid increase, Birdsall wrote, "shows the strong entry of computer experiments into at least junior partnership with theory, analysis and laboratory experiment" (Birdsall 1966, 4). It was obvious that the trend would continue and equally obvious that questions —both about their epistemological and their professional status—would arise and had in fact already arisen. In response to such questions, Buneman and Dunn rose to the defense of their methods, and they did so in a way that makes

manifest a certain already existing heterogeneity regarding the meaning of "theory." They wrote: "One encounters, at times, a prejudice against computer experiments. Partly, such prejudice is based on mathematical snobbery (the formal description of the skin effect in Bessel-functions of complex argument enjoys higher prestige than a few graphs showing how it actually goes!). But often one hears the complaint that a computer can at best say "this is *how* it happens" and never "this is *why* it happens." The examples produced here should suffice to answer this complaint. The mere fact that the computer was able to produce the 'how' has, many times, told us the 'why'" (Buneman and Dunn 1965, 56).

By the early 1970s, however, contestation had spread (if not shifted) to the relation between at least some real experiments and computer experiments. Computations of the radial distribution function based on X-ray scattering remained fraught with technical difficulties, but, by contrast, the prowess of the simulators, their machines, and those who programmed the machines increased rapidly. As a consequence, confidence in the reliability of computer "observations" soon came to rival (if not overtake) confidence in the reliability of observations based on experimental measurements. Not only were the former easier to obtain (that is, more economical), but also the repertoire of internal consistency checks available to the "simulators" soon granted their "observations" a trustworthiness that the methods of the experimentalists were unable to inspire.[10]

Nevertheless, it remains the case that it is primarily in the domain generally referred to as "theory" that "computer experiments" have had their major triumph. Today, with the accumulation of four decades of experience, their value as theoretical tools—that is, for building theory—has been amply vindicated, and defense of the sort that Buneman and Dunn had earlier felt called upon to make seems no longer to be required. Or so, at least, one might conclude from their prominence in the literature. For a contemporary assessment, I quote from the contribution to the current *Encyclopaedia Britannica* on the molecular structure of liquids. Here John M. Prausnitz and Bruce E. Poling write: "Since 1958 such computer experiments have added more to the knowledge of the molecular structure of simple liquids than all the theoretical work of the previous century and continue to be an active area of research for not only pure liquids but liquid mixtures as well" (Prausnitz and Poling 1999).

2.3. Cellular Automata and Artificial Life

The third class of computer simulation I want to discuss departs from the first two in at least one crucial respect: It is employed to model phenomena that lack a theoretical underpinning in any sense of the term familiar to physicists—phenomena for which no equations, either exact or approximate, exist (as, for

example, in biological development), or for which the equations that do exist simply fall short (as, for example, in turbulence). Here, what is to be simulated is neither a well-established set of differential equations (as in Ulam's "experiments in theory") nor the fundamental physical constituents (or particles) of the system (as in "computer experiments"), but rather the phenomenon itself. In contrast to conventional modeling practices, it might be described as modeling from above.

Perhaps the most conspicuous example of this use of simulation is to be found in A-Life studies, officially christened as such at a conference held at Los Alamos in 1987 and organized by Christopher Langton, at that time a member of the Theoretical Division of the Los Alamos National Laboratory (Langton 1989). I include A-Life under the category of "simulation in the physical sciences" for the simple reason that, despite its explicitly biological allusion, it was developed by—and for the most part has remained in the province of—physical scientists. Furthermore, Langton may have been responsible for introducing (as well as for popularizing) the term "Artificial Life" as a label for computer simulations of biological evolution,[11] but the basic project of simulating biological processes of reproduction and natural selection on the computer is in fact of much longer standing: indeed, it has its origins in the same context (and in the work of the same people) from which the first use of computer simulation for numerical analysis arose.

Von Neumann is the man most frequently credited as the "father of Artificial Life," and his contributions to the field arose directly from his preoccupations with a question that might be regarded as the oldest and most fundamental of all questions about simulation, namely, how closely can a mechanical simulacra be made to resemble an organism? More specifically, he asked: Is it possible to construct an automaton capable of reproducing itself? Beginning in the 1940s, Von Neumann worked with a kinematic model of automata afloat in a sea of raw materials but never fully succeeded in capturing the essential logic of self-reproduction. The breakthrough came with the suggestion of his close colleague, Stanislaw Ulam, that a cellular perspective (similar to what Ulam was using in his Monte Carlo computations)—in which the continuous physical motion required in the kinematic model would be replaced by discrete transfers of information—might provide a more effective approach. "Cellular automata," as they have since come to be called, have no relation to biological cells (and, indeed, from the beginning they were also invoked for the analysis of complex hydrodynamic problems), but they did suggest to Von Neumann a way of bypassing the problems posed by his kinematic model. Here, all variables (space, time, and dynamical variables) are taken to be discrete: An abstract space is represented as a lattice with a cellular automaton (a mathematical object—that is,

"a finite-state" machine) located at each node of the lattice. Each such automaton is connected to its nearest neighbors, and it evolves in time by reading the states of its neighbors at time t_n and, according to prespecified and simple rules, moving to a new state at time t_{n+1}. Ulam and Von Neumann reasoned, and indeed soon proved, that the collective dynamics resulting from such simple rules might bear a formal resemblance to the biological process of self-reproduction and evolution.

Von Neumann's initial construction in the early 1950s was cumbersome (requiring 200,000 cells with twenty-nine states for each automaton), but it made the point. The story of its subsequent development (and dramatic simplification)—from John Conway's "Game of Life" (see Gardner 1970) to Chris Langton's even simpler self-reproducing "loops" (1984)—has been recounted many times and hardly needs repeating here.[12] Somewhat less well known is the history of the use of cellular automata in the modeling of complex physical phenomena (for example, turbulence, crystallization, etc.)[13]—an activity that, like "Artificial Life," also exploded with the appearance of "super-computers" in the 1980s. Indeed, the very first conference on "Cellular Automata" was also held at Los Alamos (preceding the A-Life conference by four years), and while it provided the occasion for Langton's initial foray into artificial life, the primary focus of the earlier conference was on the physical sciences (Farmer, Toffoli, and Wolfram 1984).[14] A proper account of this part of the history of cellular automata remains to be written by historians, but my focus here is not so much historical as it is conceptual: that is, to try to identify what is distinctive about this new kind of simulation and to capture its epistemological novelty. What follows is at best a very rough (and necessarily brief) characterization.

Cellular automata are simulations par excellence: they are artificial universes that evolve according to local rules of interaction that have been prespecified. Change the initial conditions, and you change the history; change the rules of interaction, and you change the dynamics. In this sense, the analogy with differential equations is obvious. Also obvious are many of the differences between CA and DE's: the universe of CA is discrete rather than continuous; its rules generally describe interactions that are local (for example, nearest neighbor) rather than long range, and uniform rather than spatially variable; the temporal evolution of CA systems is exactly computable for any specified interactions (given enough time) while DE's are rarely susceptible to exact analytic solutions and only approximately computable when they are not.[15] But more important by far are the differences in the uses to which they are put, in the processes by which they are crafted, and in the criteria by which they are judged.

CA have a home in A-Life studies precisely because of the unavailability of differential equations for the processes they simulate; similarly, they lend them-

selves to the simulation of excitable media, turbulence, and earthquakes because the equations that do exist are not adequate to describe the phenomena of interest. And indeed, in some of their uses, CA models might be viewed simply as an alternative to DE's in which exact computability enables unambiguous tests of approximate theories—that is, just as it was claimed to do in molecular dynamics, only this time around without the quotation marks. More often, however, they are employed in a radically different spirit, aimed more at producing recognizable patterns of "interesting" behavior in their global or macrodynamics than in their microdynamics. As Stephen Wolfram writes,

> Science has traditionally concentrated on analyzing systems by breaking them down into simple constituent parts. A new form of science is now developing which addresses the problem of how those parts act together to produce the complexity of the whole.
>
> Fundamental to the approach is the investigation of models which are as simple as possible in construction, yet capture the essential mathematical features necessary to reproduce the complexity that is seen. CA provide probably the best examples of such models. (Wolfram 1986, v)

Several points bear emphasizing here, and they are related: one has to do with the process by which CA models are constructed, another with their synthetic capacities (in both senses of the word), and a third, with the focus on formal similarity between the outcomes they yield and the "overall behavior" of the processes they are designed to mimic (physical, biological, economic, or other).[16] Toffoli and Margolus's introduction to the subject is instructive, and I quote it at length:

> In Greek mythology, the machinery of the universe was the gods themselves. . . . In more recent conceptions, the universe is created complete with its operating mechanism: once set in motion, it runs by itself. God sits outside of it and can take delight in watching it.
>
> Cellular automata are stylized, synthetic universes. . . . They have their own kind of matter which whirls around in a space and a time of their own. One can think of an astounding variety of them. One can actually construct them, and watch them evolve. As inexperienced creators, we are not likely to get a very interesting universe on our first try; as individuals we may have different ideas of what makes a universe interesting, or of what we might want to do with it. In any case, once we've been shown a cellular-automaton universe we'll want to make one ourselves; once we've made one, we will want to try another one. After having made a few, we'll be able to custom-tailor one for a particular purpose with a certain confidence.
>
> A cellular automata machine is a universe synthesizer. Like an organ, it has keys and stops by which the resources of the instrument can be called into action, combined, and reconfigured. Its color screen is a window through which one can watch the universe that is being "played." (Toffoli and Margolus 1987, 1)

The seductive powers of CA are obvious, and many (including Toffoli) have been seduced. Their successes at simulating global effects have encouraged them to shift not only the meaning of simulation (and of model), but, at least in some writings, the status (or even locus) of what had earlier been taken to be the original, the real thing. Thus, for example, G. Y. Vichniac (1984) proposes "[c]ellular automata as original models of physics," and he suggests the possibility that the physical world really is a discrete space-time lattice of information bits evolving according to simple rules, an enormous CA running with one of many possible sets of rules. This view was represented in a number of presentations at the 1983 conference, and it has since gained considerable legitimacy in the world of computational physics (occasionally referred to as "synthetic physics"); in fact, however, it had already been advocated by some (notably, Ed Fredkin) ever since the 1960s.[17] Moreover, from such claims about the physical universe to Langton's (1989) arguments for "A-Life" ("we expect the synthetic approach to lead us not only to, but quite often *beyond,* known biological phenomena; beyond *life-as-we-know-it* into the realm of *life-as-it-could-be*") is a short step. The point to be noted is that, in both Vichniac's and Langton's proposals, the very cellular automata that had originally been invoked as explanatory crutch, as simulation of something prior, of features of a world assumed to be simultaneously more fundamental and more "real," have somehow metamorphosed into entities with ontological primacy in and of themselves.

So radical an inversion of conventional understandings of the relation between simulation and reality are not yet widespread—either in the physical or the biological sciences (indeed, they have yet to make any noticeable impact on the majority of biologists)—but the very fact that they have become thinkable, and in certain circles even acceptable, is surely worth noting. Minimally, it provides an indication of the power of CA models to subvert conventional distinctions between real and virtual, or between real and synthetic, and hence of their efficacy in establishing an "alternate reality." The epistemological novelty of CA modeling is in this sense quite different both from that of the Monte Carlo techniques first introduced by Ulam and Von Neumann (however much it may owe these early inventions for its technical development) and from that novelty associated with the computer experiments of molecular dynamics. Where one extended the meaning of "mathematical," and the other the range of "theory," the primary novelty of CA modeling may well lie in extending the range of the "real." I would argue, however, that the principal route by which CA modeling achieves this extension is to be found less in its capacity to present visually compelling images of synthetic (that is, artificial or virtual) objects than in its synthetic powers in the other sense of that term, namely, through its utility in synthesizing

new kinds of objects that are unambiguously real. This, I suggest, is especially evident in Artificial Life studies.

Despite initial hopes in the value of CA modeling in promoting better theory —in particular, a better understanding of biological principles—Artificial Life studies have made little impression on practicing biologists. Far more significant has been their influence on engineers. In a recent book entitled *Creation: Life and How to Make It,* Steve Grand writes, "Research into artificial life is inspiring a new engineering discipline whose aim is to put life back into technology. Using A-life as an approach to artificial intelligence, we are beginning to put souls into previously lifeless machines. . . . The third great age of technology is about to start. This is the Biological Age, in which machine and *synthetic* organism merge" (Grand 2002, 7–8).

Synthetic life forms that are real objects in the sense that they are made from material components and assembled in real space and time are clearly being built, and in ways that draw directly from work on "lifelike" simulations in cyberspace. Engineering is a science that specializes in negotiating the gap between symbol and matter, and robotic engineers, like their colleagues in allied disciplines, have well-developed techniques for translating from one domain to the other, for realizing the metaphors of simulation in the construction of material objects. Computer simulations of biological organisms may well be "metaphorical representations," but they are also models in the time-honored sense of guides or blueprints: in the hands of skillful engineers, they can be, and are, used as guides to construction in an altogether different medium. Here, the simulated organisms of cyberspace are used to guide the synthesis of material objects mimicking the behavior of biological organisms in real space. Without doubt, these entities are real. But another question immediately arises: are they "alive"? This is a question that worries many philosophers, but, as I argue elsewhere (Keller 2002, chap. 9), it may well be a question that belongs more properly in the realm of history than in that of philosophy.

NOTES

1. While the shift from deception to instruction is undoubtedly worthy of study in itself, my aim in this chapter is merely to examine its impact rather than the process by which it occurred.

2. See, e.g., discussion of echo simulators developed to train AI operators of an aircraft interception radar set in Garman (1942).

3. The use of computer simulations in the biological sciences—how it both draws from and differs from its uses in the physical sciences—is of particular interest to me, but I refer the reader to chapters 8 and 9 of Keller (2002) for discussion of this topic.