"Artificial Societies" and the Social Sciences

J. Stephen Lansing
Department of Anthropology
University of Arizona
Tucson, AZ 85721
jlansing@u.arizona.edu

Das Ganze ist das Unwahre.
        —Theodor Adorno

I

In 1961 the German Sociological Association invited Karl Popper to give a lecture on the logic of the social sciences, and invited Theodor Adorno to offer a critical response. It was to be a battle of the titans of postwar sociology, and afterwards became known as the "Positivist dispute" (*Positivismusstreit*) in German social theory. Yet from the beginning the debate went in an unexpected direction. Popper refused to accept the label of positivist, and explicitly attacked an inductivist and naturalist conception of science. Adorno was partially disarmed by this, and offered some reflections on Popper's theses rather than the expected attack on positivist social science. These were not well received; Popper apparently found Adorno's Hegelian language unintelligible, and complained that he was "simply talking trivialities in high-sounding language" [1]. The debate was carried on inconclusively for a decade by partisans of both schools.[1] Popper's original suggestions, and Adorno's response to them, were soon forgotten as the protagonists retreated to well-worn defensive positions.

It is interesting to speculate whether this debate might have gone differently if the development of methods for the study of "artificial societies" had begun fifty years earlier.[2] In retrospect it seems evident that both Popper and Adorno were struggling with the limitations of a social science methodology based on descriptive statistics. Yet they approached this problem from nearly opposite perspectives. Popper described his methodology as "criticist" rather than "positivist," and argued for a view of scientific method as consisting of "tentative attempts at solutions." To this Adorno made two responses. First, he argued that social theory must be able to conceive of an alternative to contemporary society: "only through what it is not will it disclose itself as it is . . . " [1]. This led him to a critique of descriptive statistics as the primary tool for social inquiry. He argued that "A social science that is both atomistic, and ascends through classification from the atoms to generalities, is the Medusan mirror to a society which is both atomized and organized according to abstract classificatory principles . . . " [1]. Adorno's point was that a purely descriptive, statistical analysis of society at a given historical moment is just "scientific mirroring" that ". . . remains a mere duplication. . . . As a corrective, it is not then sufficient simply to distinguish descriptively between the 'collective realm' and the 'individual realm' as Durkheim intended, but *rather the relationship between the two realms must be mediated and must itself be grounded theoretically*" [1] (emphasis mine).

---

1 Albrecht Wellmer responded with a critical analysis of Popper's views in 1967, and another of positivism and critical theory in 1969 [51]; Jurgen Habermas made several unsuccessful forays at clarifying Adorno's response to Popper (see his chapter in [1]) and later attempted to reformulate the entire debate [19].

2 For a historical overview, see Nigel Gilbert, "Modeling sociality: The view from Europe" in [18].

These arguments should sound familiar to students of "artificial societies." The limitations of descriptive methods, the concept of alternative pathways of historical development, and perhaps most crucially the relationship between the individual and the collective are all major themes in contemporary research. Compare, for example, Adorno's final point to the question with which Joshua Epstein and Robert Axtell introduce *Growing artificial societies* (1996):

> How does the heterogeneous micro-world of individual behaviors generate the global macroscopic regularities of the society? [13]

Purely descriptive methods, they go on to say, " 'filter out' all consequences of heterogeneity. Few social scientists would deny that these consequences can be crucially important, but there has been no natural methodology for systematically studying highly heterogeneous populations" [13]. That natural methodology, they suggest, is agent-based modeling. However, they observe, "it is only in the last decade that advances in computing have made large-scale agent-based modeling practical" [13].

But this emphasis on the new possibilities opened up by computers—a theme that is echoed by most authors writing about artificial societies—has some drawbacks. Scholars who are not modelers may see the field as technology-driven, often amnesiac, certainly insular, occasionally tinged with science fiction, and prone to extravagant claims. For example, Epstein and Axtell ask

> What constitutes an explanation of an observed social phenomenon? Perhaps one day people will interpret the question, "Can you explain it?" as asking "Can you *grow* it?" Artificial society modeling allows us to "grow" social structures *in silico* demonstrating that certain sets of microspecifications are *sufficient to generate* the macrophenomena of interest. [13]

Reading this, a mainstream scientist who is dubious about computer modeling may conclude that his skepticism was justified. One does not need to be a modeler to know that it is possible to "grow" nearly anything *in silico* without necessarily learning anything about the real world (one of the most frequent criticisms of this approach).[3] My point is not to criticize Epstein and Axtell for overshooting the mark in their enthusiasm for computer models, but rather to suggest that a broader historical perspective is needed. Large-scale social simulation models are indeed new, but the questions that prompt their creation are very old.

So before plunging into the new world of dynamical systems, attractors, and emergent properties, I propose to take a brief detour into the past. Suppose, then, we wind back the clock to the time of the *Positivismusstreit* and, with the assistance of Popper and a few of his colleagues, consider the condition of the social sciences before the age of computers. According to Popper, there are two types of social theory. One is essentially Newtonian, and is best represented by general equilibrium theories, for example in economics. Such theories take the form of systems of differential equations. Change occurs as a result of perturbations and leads from one equilibrium state to another. The second type of theory is statistical. If one cannot write the equations to define a dynamical system, it may yet be possible to observe regularities in social phenomena. Thus in 1897 Durkheim famously used statistics to prove the existence of what he termed "social forces" that "have an existence of their own; they are forces as

---

3 Thus John Maynard Smith writes in the *New York Review of Books* that he has "a general feeling of unease when contemplating complex systems dynamics. Its devotees are practicing fact-free science. A fact for them is, at best, the outcome of a computer simulation; it is rarely a fact about the world" [47].

real as cosmic forces." His argument was based on the observation that the statistical rate of suicides "is even more stable than that of general mortality":

> The individuals making up a society change from year to year, yet the number of suicides is the same so long as the society itself does not change. The population of Paris renews itself very rapidly, yet the share of Paris in the total of French suicides remains practically the same.... The causes which thus fix the contingent of voluntary deaths for a given society or one part of it must then be independent of individuals. [12]

Descriptive statistics offered a methodology for the analysis of "social forces," one that avoided the simplifying assumptions of equilibrium theories. As Durkheim observed, the significance of suicide is precisely that it is not commonplace; it is not the expected behavior of Quetelet's "average man." Suicides are a tiny fraction of the population; they are statistical outliers; what is interesting is that the magnitude of this fraction is remarkably constant over time. Is there then a "propensity" to suicide? Durkheim grappled with this question with the statistical ideas available to him in 1895, and Popper placed the question of "propensities" at the center of his history of scientific reasoning; his last book was entitled *A World of Propensities* [40]. If the reader will bear with me, the question of "propensity" bears directly on what I take to be the truly innovative possibilities offered by the study of "artificial societies."

Durkheim's reflections on propensities begin with his observations on the differences between what he termed "egoistic" suicide, with a constant rate, and "anomic" or "altruistic" suicide, which depended on the rate of breakdown of "traditional" social forms as a consequence of modernization. He termed anomic suicide the "ransom money of civilization," observing that it "springs not from the intrinsic nature of progress but from the special conditions under which it occurs in our day" [12]. The key question, for Durkheim, was the causal inferences that could be drawn from the statistics on anomic suicides, since the regularity of the patterns suggested a deterministic process: "Since each year has an equal number of suicides, the current does not strike simultaneously all those within its reach. The persons it will attack next year already exist; already, also, most of them are enmeshed in the collective life and therefore come under its influence. Why are they provisionally spared?...Why, to use the familiar expression, does society pay its bill only in installments?" [12]

In retrospect, the issue Durkheim was grappling with was a shift in the understanding of probability from chance to objective probability, or what Popper calls "propensity." The solution to this problem led to the reformulation of some physical laws in probabilistic terms. "The statistical method," wrote the physicist James Clerk Maxwell in 1890, "involves an abandonment of strict dynamical principles" [41]. This led to one of the most significant intellectual pirouettes in the history of science, according to Popper. With the notion of "objective probability," "the world is no longer a causal machine. It now can be seen as a world of propensities, as an unfolding process of realizing possibilities and of unfolding new possibilities" [40]. (It is an interesting footnote to the history of the sciences that this reformulation of the concept of chance occurred in the social sciences and was borrowed by the physicists. "Doubtless it would be too brave," writes Porter in *The Rise of Statistical Thinking* "to argue that statistical gas theory only became possible after social statistics accustomed scientific thinkers to the possibility of stable laws of mass phenomena with no dependence on predictability of individual events. Still, the actual history of the kinetic gas theory is fully consistent with such a claim" [41].)

In the natural sciences, it proved to be a reasonably straightforward matter to average over large numbers of heterogeneous gas molecules or, later, atoms, while retaining dif-

ferential equations. One of Popper's favored examples is the Franck-Hertz experiment, which measured the propensity of electrons to interact with gas atoms:

> The Franck-Hertz experiment, one of the classic experiments in quantum mechanics, studies the dependence of this interaction on an increasing voltage. As the voltage rises, the intensity of the current of electrons rises slowly and then, suddenly, falls; it rises again slowly to a still higher level and falls again suddenly. This is interpreted as the result of the single electrons reaching, step by step, the discrete excitation states of the gas atoms. [40]

Popper observes that for this type of experiment, we need a calculus of relative or conditional probabilities rather than absolute probabilities. He notes that a statement in the absolute calculus may be written

$$p(a) = r, \tag{1}$$

to be read: "The probability of the event a equals r." This contrasts with the relative or conditional probability statement

$$p(a, b) = r, \tag{2}$$

to be read: "The probability of the event a in the situation b (or given the conditions b) equals r" [40]. In general, "propensities in physics are properties of *the whole physical situation* ... and the same holds of the propensities in chemistry, in biochemistry and in biology" [40].

But how could this understanding of conditional probabilities, or propensities, be translated back into the social sciences? As Popper observed, "one important aspect of the Franck-Hertz experiment which it shares with many other quantum experiments is that, even though the conditions change, we can *measure* the propensities because there are so many electrons involved" [40]. But there are a limited number of quantum states for the particles in the Franck-Hertz experiment. In contrast, the social sciences must contend with agents whose heterogeneity seemingly knows no bounds, and on whom experiments may not be performed. And so we seem to be left with essentially the same tools that were available to Durkheim in 1895: either to sift for patterns in social statistics, or to construct dynamical equilibrium models of homogeneous actors.

It is here, I suggest, that the significance of artificial societies as a genuine methodological innovation becomes apparent. Formal models of actors and societies are classically represented by linear differential or difference equations. Heterogeneous actors and environments introduce nonlinear elements, which make these systems difficult or impossible to solve. But using computers it is possible to step through solutions by simple brute force. As Stephen Wolfram observes, "One way to find out how a system behaves in particular circumstances is always to simulate each step in its evolution explicitly" [52]. So computers offer a solution to the problem of incorporating heterogeneous actors and environments, and nonlinear relationships (or effects). Still, the worry is that the entire family of such solutions may be trivial, since an infinite number of such models could be constructed. If analytical solutions are forever out of reach, is there any point to performing the calculations?

The missing ingredient is a method for attaching objective probabilities to social outcomes. What we would really like to know, and what descriptive statistical methods alone usually cannot tell us, is the likelihood of a given outcome, a given social reality. "What if the evolutionary tape were run again?" as Stephen Jay Gould famously asked.

Here the physical sciences repay their debt to the social sciences with some new analytical tools. The most basic technique is dynamical systems theory, which provides methods to describe the global behavior of systems of differential equations. Systems of N variables are tracked in their passage through N-dimensional space, allowing the researcher to ask where the system spends most of its time, and thus to assess its characteristic patterns of behavior. In continuous deterministic dynamical systems, all possible time series make up the vector field that is represented by the system's phase portrait, an idea first proposed by Poincaré [39]; (see also [23]). Analogous questions can be asked about discrete dynamical systems, such as random Boolean networks and cellular automata (cf. [53, 54]). Where analytical solutions cannot be found, repeated simulations can be analyzed for their statistical properties. The initial exploration of these systems has been rewarded by the discovery of unsuspected patterns of simplicity and order, from Feigenbaum's number [15] to Langton's "edge of chaos" [30], Wolfram's Class 4 cellular automata [53], and Bak's self-organized criticality [6]. Closer to home (that is, to the social sciences), evolutionary game theory has been remarkably successful at predicting patterns of behavior in stickleback fish, naked mole rats, elephant seals—and, possibly, *Homo sapiens*.

The significance of these developments for the social sciences is nicely captured in an anecdote described in Waldrop's *Complexity* [49]. Waldrop is describing a meeting of physicists and economists at the Santa Fe Institute in November 1988. He quotes the economist Brian Arthur:

> We had just begun to realize that if you do economics this way—if there was this Santa Fe approach—then there might be no equilibrium in the economy at all. The economy would be like the biosphere: always evolving, always changing, always exploring new territory.
>
> "Now, what worried us is that it didn't seem possible to do economics in that case," says Arthur. "Because economics had come to mean the investigation of equilibria.... So Frankie Hahn said, 'If things are not repeating, if things are not in equilibrium, what can we, as economists, say? How could you predict anything? How could you have a science?' " [49]

The question was taken up by John Holland, who suggested a parallel with meteorology: "The weather never settles down, it's essentially unpredictable more than a week or so in advance. And yet we can comprehend and explain nearly everything that we see up there.... We can understand their dynamics." According to Waldrop [49], this answer came as something like a revelation to Brian Arthur:

> It left me almost gasping. I had been thinking for almost ten years that much of the economy would never be in equilibrium. But I couldn't see how to 'do' economics without equilibrium. John's comment cut through the knot for me. After that it seemed—straightforward. [49]

2

Society is a human product. Society is an objective reality. Man is a social product.

—Peter Berger and Thomas Luckmann [8]

Today the study of non-equilibrium economics is well under way, for example in the simulation of stock markets, both real and imaginary. The study of the global properties

of these economies has been accompanied by research on the behavior of economic actors [4]. The economist Samuel Bowles and his colleagues work with psychologists and anthropologists to investigate, as an empirical question, how social actors make decisions in games involving economic principles. Based on these results, they propose to supplement *Homo economicus* with a larger and more diverse "family of man" [9]. Research thus proceeds on two levels: the characteristics of individual social actors, and the global dynamics of economies or societies. This strategy is encouraged by the discovery that systems composed of heterogeneous agents with limited knowledge of their environments can exhibit emergent properties of order. The sources for this important idea include the physicist Philip Anderson ("Psychology is not applied biology, nor is biology applied chemistry" [2]), John Holland's work on "complex adaptive systems" [24], and Stuart Kauffman's studies of random Boolean networks [28]. The search for emergent patterns became the major theme of the initial program on economics at the Santa Fe Institute [3]. In principle, it offers a solution to Berger and Luckman's famous paradox (above): as time proceeds in a simulation, the "heterogeneous micro-world of individual behaviors" can generate global patterns, which in turn constrain the future behavior of individuals.

But economics is not the only discipline where the study of artificial societies is beginning to have an impact. One of the interesting features of this research is that ideas and methods move rapidly from field to field; there is a great deal of cross-disciplinary borrowing. At or near the center of these crisscrossing paths is the simulation of evolutionary processes, a field that is largely the creation of one man, John Holland. In the 1960s Holland developed methods to implement R. A. Fisher's mathematical model of natural selection in computer code, creating what Holland christened a *genetic algorithm*. These devices do not merely simulate the evolutionary process; they reproduce it *in silico*. Holland's work on genetic algorithms roughly coincided with important developments in the mathematical theory of evolution, and also in game theory. The convergence of these three fields has become perhaps the most fecund source of ideas in artificial societies. Here I briefly sketch this history.

In 1930 R. A. Fisher published his *fundamental theorem of natural selection*, and at about the same time V. Volterra described what is now known as the *predator-prey equation* [48]. In 1988, Hofbauer showed that Fisher's selection equation is a special case of Volterra's equation, and in 1983 Schuster and Sigmund showed that both equations could be subsumed by another, since "essentially one type of equation models the evolution (1) of allele frequencies in a gene pool, (2) of population densities in a habitat, (3) of concentrations of polynucleotides in a flow reactor, and (4) of probabilities of strategies for conflicts within one species" [44]. They proposed the term "replicator equation" for this differential equation, signifying its close relationship to the concept of replicators developed by Richard Dawkins. Stressing the wide applicability of this type of equation, Schuster and Sigmund "emphasize that this result has not been obtained through a mathematician's desire to apply a pretty equation to as many fields as possible: the modeling leading to Equations (1)–(4) proceeded independently in genetics, ecology, prebiotic chemistry, and sociobiology" [44]. (See the appendix for a description of the replicator equation.)

Holland's genetic algorithm implements this equation *in silico*, representing genotypes as computer code. As the eminent evolutionary biologist W. D. Hamilton observed, "It was a very brilliant step of John Holland to realize that he could probably do something in artificial intelligence by just importing the principles of life directly" (quoted in [33]). And indeed, the genetic algorithm has opened an entirely new approach in artificial intelligence, enabling researchers to "grow" or "evolve" solutions to problems. In the field of artificial societies, the genetic algorithm has been implemented directly, for example by Holland himself [25]. But perhaps more importantly,

it has helped to inspire a new approach to social simulations, emphasizing the evolution of strategies, agents, or societies. These models do not always employ the specific techniques of Holland's genetic algorithm, but use analogous procedures to mimic evolutionary processes. They represent an intersection of game theory, computer science, and evolutionary theory that has led to a proliferation of models and ideas. But again, a skeptic might ask, what is truly new about these implementations of evolutionary processes in the computer?

The answer is, emergence. Consider the replicator equation. It captures a very simple idea: increase in fitness under selection. But it does not include a term for the environment in which selection occurs. Instead it describes a process of selection based on a one-way influence between organism and environment. This assumption simplifies the analysis of selective processes, because it avoids consideration of feedback loops that could produce highly complex, or even chaotic, dynamics [16]. But experiments in evolutionary simulations have repeatedly found that when agents are embedded in an environment, the dynamical behavior of the system as a whole is altered. As with Popper's Equation 2, "propensities" turn out to be properties of the whole dynamical system. The most famous example is Axelrod's reformulation of the prisoner's dilemma, a two-player nonzero-sum game in which rational choice leads to ruin. By creating a population of simulated agents who play a series of games with each other—in other words, treating the game as a dynamical system evolving over time—Axelrod found that cooperation would emerge under a wide range of conditions. In further experiments he tested the ability of the genetic algorithm to evolve winning strategies, and found that evolved strategies typically did as well as or better than those created by game theorists. Axelrod concluded that the genetic algorithm proved to be "very good at what actual evolution does so well: developing highly specialized adaptations to specific environments" (quoted in [33]; see [5]). Subsequently, Kristian Lindgren embedded game-playing agents on a lattice, adding greater flexibility by making memory length an evolutionary variable. Over tens of thousands of generations, he observed the emergence of spatial patterns that resemble evolutionary processes and clarify preconditions for the emergence of cooperation and competition [35, 14]. Such simulation results have inspired behavioral ecologists to reexamine biological systems. For example Manfred Milinski has studied stickleback fish, which enjoy "a well-earned reputation for keeping abreast of the latest trends in animal behavior" [45]. According to Milinski, cooperation in predator inspection by the sticklebacks follows the dynamics of the iterated prisoner's dilemma. The results of these simulations have also been used to model problems in political science and economics.[4]

But cooperation is by no means the only emergent property investigated by social simulations. The philosopher Brian Skyrms has studied the "evolution of the social contract" by modeling it as a problem in the evolution of dynamical systems driven by replicator dynamics. His most ambitious models tackle such large questions as the evolution of justice, linguistic meaning, and logical inference. Skyrms finds that "the typical case is one in which there is not a unique preordained result, but rather a profusion of possible equilibrium outcomes. The theory predicts what anthropologists have always known—that many alternative styles of social life are possible" [46]. But this is a bit too modest. With respect to the evolution of meaning, for example, Skyrms shows that evolutionary processes provide a plausible answer to the fundamental question, "How do the arbitrary symbols of language become associated with the elements of reality they denote?" [46].[5]

---

4 There is a large literature on this topic. For a critique of the standard model for the evolution of altrusim based on kin selection, using agent-based modeling, see [38]; for a possible biological example, see [53].

5 For a review of other recent work on the computer modeling of the emergence and evolution of language, see [10].

While replicator dynamics is the dominant dynamical process in artificial societies models, there have also been experiments that illustrate the limitations of this approach. The idea that "survival of the fittest" is the principal driver of both social and biological evolution was criticized long ago by Karl Marx,[6] a criticism echoed today by Richard Lewontin. According to Lewontin, "the success of the program of physicalizing biology has encouraged the program, also inherited from the nineteenth century, of biologizing the psychic and the social" [34]. In particular, Lewontin takes a dim view of the "metaphor of adaptation," which he sees as "an impediment to a real understanding of the evolutionary process [which] needs to be replaced." Lewontin proposes an alternative metaphor, *niche construction*, to emphasize the active role of organisms in constructing their own environments.

To clarify the differences between simple replicator dynamics and niche construction, and the implications of these differences, consider James Lovelock's celebrated model Daisyworld [36]. This model illuminates the consequences of a shift in perspective from simple equations like the replicator to an entire dynamical system: from the world of Popper's Equation 1 to that of Equation 2.

Daisyworld is an imaginary planet orbiting a star like the sun and at the same orbital distance as the earth. The surface of Daisyworld is fertile earth sown uniformly with daisy seeds. The daisies vary in color, and daisies of similar color grow together in patches. As sunshine falls on Daisyworld, the model tracks changes in the growth rate of each variety of daisy and the resulting changes in the amount of the planet's surface covered by different-colored daisies. The simplest version of this model contains only two varieties of daisies, white and black. If the daisies cover a sufficiently large area of the surface of Daisyworld, their color affects not only their own microclimate but also the albedo of the planet as a whole.[7] Lovelock modeled the effect of gradually increasing solar radiation on the temperature of Daisyworld, similar to that thought to have occurred on earth as our sun became hotter. He showed that simple positive and negative feedback will occur in the proportion of colored daisies covering the planet, and that this results in stabilizing the average temperature on Daisyworld at approximately the optimum temperature for the flowers. The mathematician Peter Saunders comments:

> To see how different the situation on Daisyworld is from that envisaged in neo-Darwinian evolutionary theory, imagine that an inter-planetary expedition is sent from the earth to Daisyworld. The astronauts would report that they had found a planet with a mean temperature of 22.5°C and that all that seemed to grow on it were two species of daisies. Since there was no other life on the planet, they would have plenty of time to study the daisies in detail and they would find that both species grew fastest at a temperature of 22.5°C. And they would almost certainly conclude that here was a clear example of natural selection. Obviously the daisies had become adapted to the prevailing conditions on the planet.

---

6 In a letter to Engels, Marx wrote: "It is remarkable how Darwin recognizes among beasts and plants his English society with its division of labor, competition, opening of new markets, 'inventions,' and the Malthusian 'struggle for existence.' It is Hobbes's 'bellum omnia contra omnes,' and one is reminded of Hegel's *Phenomenology*, where civil society is described as a 'spiritual animal kingdom,' while in Darwin the animal kingdom figures as civil society." Cited in [43].

7 In the Daisyworld model, black daisies have an albedo (or reflectance) of 0.25 and so absorb more heat than bare earth, which has an albedo of 0.5, while white daisies have an albedo of 0.75. The model posits that as a result of these differences in albedo, clumps of same-colored daisies create a local microclimate for themselves, slightly warmer (if they are black) or cooler (if white) than the mean temperature of the planet. The mean temperature of the planet, in turn, is a function of the solar luminosity. Both black and white daisies grow fastest and at the same rate when their local effective temperature (the temperature within their microclimate) is 22.5°C, and they respond identically, with a decline in growth rate, as the temperature deviates from this ideal.

> That is not, of course, what actually happened. In particular, the number 22.5°C is a property of the daisies, not the planet. The daisies did not adapt to the environment. They adapted the planet to suit themselves. [42]

Temperature regulation is an emergent property of Daisyworld, one that improves the fitness of all life on the planet, not just those individuals that are temporarily favored by the replicator dynamics. It is not that the replicator is absent: the growth of the daisies is explicitly modeled by this equation. But the selective process is controlled at a higher level by the entire coevolving system including the star, the planet, and the flowers. We have moved from Popper's Equation 1 to Equation 2, where propensities are properties of the whole system. Perhaps for this reason (a shift in theoretical focus from competitive individuals to cooperative niche construction), Daisyworld has been quite controversial. Lovelock comments that his model "draws attention to the fallibility of the concept of adaptation. It is no longer sufficient to say that 'organisms better adapted than others are more likely to leave offspring.' It is necessary to add that the growth of an organism affects its physical and chemical environment" [37].

One way to think about the implications of models like Daisyworld for the social sciences is in the context of functionalism. The nature and origins of functional relationships in societies is one of the classical questions in the social sciences. But as Carl Hempel showed long ago, functionalist arguments are typically flawed by a vague assumption that societies are normally functionally integrated, ignoring the question of how this integration is achieved [22]. Still, as the historian Barry Barnes observes, "Functionalism is too deeply ingrained in the modes of thought of the social sciences for its abandonment to be a realistic possibility" [7]. Models like Daisyworld allow the investigator to explore the conditions in which functional integration may occur, and to estimate conditional probabilities. Feedback processes need not involve biology; they may be entirely social, as in recent studies of the emergence of social networks (cf. [50, 38]).

But to me as an anthropologist, an even more important result is the revelation that social institutions can emerge from the bottom up, as a result of feedback processes linking social actors to their environments (as Kohler and Gumerman observe in their recent volume [29]). Such institutions might look very different from those that social scientists normally study; they might even be invisible. Recently my colleagues and I have suggested that the water temple networks with which Balinese farmers manage their centuries-old irrigation systems and rice terraces are a real-world example of a complex adaptive system (*sensu* Holland), whose dynamics resemble those of Lovelock's Daisyworld [32]. For decades, both social scientists and engineers have marveled at the success of the Balinese in managing complex irrigation systems involving hundreds of villages. But the question of how this was achieved remained mysterious. Our agent-based computer simulations showed that even though local communities "do not consciously attempt to create an optimal pattern of staggered cropping schedules for entire watersheds ... the actual patterns [we] have observed in the field bear a very close resemblance to computer simulations of optimal solutions" [31]. We speculate that the water temples of Bali are probably not unique, and may be representative of a whole class of social institutions whose existence has so far gone unnoticed because until now we lacked the appropriate conceptual models.

The study of institutions like the water temple networks, or the emergence of cooperation in human communities, offers a new context in which to explore such fundamental issues as the relationship between structure and function. For example, Erica Jen has recently explored the difference between concepts of "robustness" and "stability" in dynamical systems, a contrast that (as she points out) "touches on essentially every aspect of what we instinctively find interesting about robustness in natural, en-

gineering, and social systems." "Robustness" has to do with the stability and resistance to perturbation of dynamical or complex systems as they evolve in state space: "A dynamical system is said to be *structurally stable* if small perturbations to the system result in a new dynamical system with the same qualitative dynamics. The concept of robustness raises additional questions that lie outside the purview of stability theory, such as the interplay of system organization and system dynamics" [27].

This brings me to my third and final point about models like Bali and Daisyworld: they provide a way to explore such questions as the robustness of coupled human and natural systems, from a cross-disciplinary perspective. This is becoming a matter of some urgency, since we are already performing the experiment (recently, ecologists have calculated the fraction of the earth's total biological productivity that is appropriated by *Homo sapiens* at nearly 40%, and rising fast) [17].

3

There have been several recent critiques of the field of artificial societies in addition to that of Maynard Smith (quoted earlier). The science writer John Horgan cautions that "as the philosopher Karl Popper pointed out, prediction is our best means of distinguishing science from pseudo-science.... The history of 20[th]-century science should also give complexologists pause. Complexity is simply the latest in a long line of highly mathematical 'theories of almost everything' that have gripped the imaginations of scientists in this century" [26]. (Here Horgan appears to be mostly concerned with the very general theories of emergence developed by Stuart Kauffman and Per Bak, among others.) In the social sciences, as we have seen, many artificial societies models are explicitly predictive, though the nature of the predictions is quite variable.

A more ambitious critique was recently published by an anthropologist, Stefan Helmreich, who offers an ethnographic account of the researchers working at the Santa Fe Institute in the mid-1990s. In *Silicon Second Nature* [20], Helmreich argues that artificial societies models reflect the unconscious cultural assumptions and social prejudices of their creators: "Because Artificial Life scientists tend to see themselves as masculine gods of their cyberspace creations, as digital Darwins exploring frontiers filled with primitive creatures, their programs reflect prevalent representations of gender, kinship, and race and repeat origin stories most familiar from mythical and religious narratives." For example, Helmreich describes Holland's genetic algorithms as reflecting a "heterosexual" bias: "There are a number of ways we might understand the exchange of bits between strings, but the metaphor of productive heterosex is gleefully emphasized by most authors" [20]. Thus for Helmreich, simulation models are like Rorschach tests, revealing the researcher's cultural background and psychological idiosyncrasies. All statements, especially theoretical pronouncements, are taken not as statements about the world, but as evidence about the author's beliefs and mode of thought. "That many Artificial Life practitioners are white men who grew up reading cowboy science fiction," observes Helmreich, "is not trivial" [20]. Simulation models may also be dangerous (as Helmreich suggests with reference to my own work), urging that "the use and abuse of computer simulations bears watching—especially in situations where there is a notable power differential between those putting together the simulation and those whose lives are the subjects and objects of these simulations" ([21]; for my response, see [31]).

Perhaps, when the science wars have quieted down a bit, a less tendentious critique will be forthcoming from the science studies quarter. In my view, such a critique is overdue. "The strain of simplistic scientism that characterized social theory from the beginning of the nineteenth century," observes Richard Lewontin, "continues to infect it today," perhaps nowhere more conspicuously than in some areas of artificial societies research [34]. The problem with Helmreich's critique is that it applies with equal force

to any and all models or theories, including those of the natural sciences. One wishes he had used a sharper instrument, one that could rid us, for example, of such fantasies as "memes." The field is in need of a more focused critique, and since Helmreich has not supplied one, consider Lewontin's suggestions as to where one might begin:

> The successes of the natural sciences in explaining the physical and biological world have affected not only the content of explanations of social phenomena but the image of how we are to go about investigating them. Studies of human societies become "social sciences" with an apparatus of investigation and statistical analysis that pretends that the process of investigation is not itself a social process. . . .
>
> I have considerable sympathy for the position in which sociologists find themselves. They are asking about the most complex and difficult phenomena in the most complex and recalcitrant organisms, without that liberty to manipulate their objects of study which is enjoyed by natural scientists. In comparison, the task of the molecular biologist is trivial. [33]

But are we then back where we began, merely ventriloquizing the arguments of Adorno versus Popper, *Vernunft* versus induction? Perhaps not quite. Certainly, Adorno's critique of the quantitative methods of his day remains relevant to the construction of artificial societies. If our "toy models" serve only to reify and naturalize the conventional social science wisdom, then they are indeed a Medusan mirror, freezing the victim by the monster's glance. But the success of the first generation of ALife simulations surely came about because these models surprised us and confounded our intuitions (one thinks of Robert May's discovery of chaos lurking unsuspected in the logistic equation, and the emergence of exotic ecologies in Tom Ray's Tierra). In his final work Adorno emphasized two themes: the need to reconceive society's relationship to nature, and to break the seal of reification on the existing social order by exploring that which does *not* exist. Some historians have argued that his most successful expositions of these ideas were articulated in his studies of classical music, where he pondered the relationship between the inner logic of musical abstractions and the social imagination. Perhaps, then, it is not too fanciful to imagine him intrigued by the transposition of these themes into the emergent patterns of computer simulations.

Acknowledgments

References

1. Adorno, T. W., Dahrendorf, R., Pilot, H., Albert, H., Habermas, J., & Popper, K. (1976). *The positivist dispute in German sociology*, transl. Glyn Adey and David Frisby (p. 296). London: Heinemann.

2. Anderson, P. W. (1972). More is different. *Science, 177*(4047, 4 August), 393.

3. Arthur, W. B. (1989, 1990). *Emergent Structures: A Newsletter of the Economic Research Program* (March 1989, August 1990). Santa Fe: The Santa Fe Institute.

4. Arthur, W. B., Durlauf, S. N., & Lane, D. A. (1997). *The economy as an evolving complex system II*. Addison-Wesley and the Santa Fe Institute.

5. Axelrod, R. (1997). *The complexity of cooperation: Agent-based models of cooperation and collaboration*. Princeton, NJ: Princeton University Press.

6. Bak, P., & Chen, K. (1991). Self-organized criticality. *Scientific American, 264*(1), 46–53.

7. Barnes, B. (1995). *The elements of social theory*. Princeton: Princeton University Press.

8. Berger, P. L., and Luckman, T. (1966). The social construction of reality. New York: Doubleday.

9. Bowles, S., & Gintis, H. (2002). Homo reciprocans. *Nature, 415*, 125–128.

10. Cangelosi, A. (2001). Evolution of communication and language using signals, symbols and words. *IEEE Transactions on Evolutionary Computation, 5*(2) (CogPrints).

11. Christiansen, F. B. (1988). Frequency dependence and competition. *Philosophical Transactions of the Royal Society of London. Series B, 319*, 587–600.

12. Durkheim, E. (1897). *Suicide: A study in sociology*, transl. by J. A. Spaulding and G. Simpson (pp. 303, 309). New York: The Free Press, 1979.

13. Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up* (p. 1). Washington, DC: The Brookings Institution Press and The MIT Press.

14. Eriksson, A., & Lindgren, K. (2001). *Evolution of strategies in repeated stochastic games* (Working paper 01-04-023). Santa Fe Institute.

15. Feigenbaum, M. J. (1980). Universal behavior in nonlinear systems. *Los Alamos Science, 1*, 4–27.

16. Ferriere, R., & Fox, G. A. (1995). Chaos and evolution, *Tree, 10*, 480–485.

17. Field, C. B. (2001). Sharing the garden. *Science, 294*, 2490–2491.

18. Gilbert, N. (2000). Modeling sociality: The view from Europe. In T. Kohler & G. Gumerman (Eds.), *Dynamics in human and primate societies: Agent-based modeling of social and spatial processes* (pp. 355–372). Oxford University Press and Santa Fe Institute.

19. Habermas, J. (1968, 1972). *Knowledge and human interests*. London: Heinemann.

20. Helmreich, S. (1998). *Silicon second nature: Culturing artificial life in a digital world*. University of California Press.

21. Helmreich, S. (1999). Digitizing "development": Balinese water temples, complexity, and the politics of simulation. *Critique of Anthropology, 19*(3).

22. Hempel, C. G. (1965). *Aspects of scientific explanation*. New York: Free Press.

23. Hofbauer, J., & Sigmund, K. (1988). *The theory of evolution and dynamical systems*. Cambridge, UK: Cambridge University Press.

24. Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor: University of Michigan Press.

25. Holland, J. H. (1995). *Hidden order: How adaptation builds complexity*. New York: Helix Books (Addison-Wesley).

26. Horgan, J. (1995). From complexity to perplexity. *Scientific American, 272*(6, June), 104–110.

27. Jen, E. Stable or robust? What's the difference? http://discuss.santafe.edu/robustness.

28. Kauffman, S. (1993). *The origins of order: Self-organization and selection in evolution*. New York: Oxford University Press.

29. Kohler, T., & Gumerman, G. J. (Eds.) (2000). *Dynamics in human and primate societies: Agent-based modeling of social and spatial processes*. New York: Oxford University Press and Santa Fe Institute.

30. Langton, C. G. (1990). Computation at the edge of chaos: Phase transitions and emergent computation. *Physica D*, 12–37.

31. Lansing, J. S. (2000). Foucault and the water temples: A reply to Helmreich. *Critique of Anthropology, 20*(3, September), 313.

32. Lansing, J. S., Kremer, J. N., & Smuts, B. B. (1998). System-dependent selection, ecological feedback and the emergence of functional structure in ecosystems, *Journal of Theoretical Biology, 192*, 377–391.

33. Levy, S. (1992). *Artificial life: A report from the frontier where computers meet biology* (p. 160). New York: Vintage Books.

34. Lewontin, R. (2000). *It ain't necessarily so: The dream of the human genome project and other illusions.* New York Review of Books.

35. Lindgren, K. (1994). Evolutionary dynamics of simple games. *Physica D, 75*, 292–309.

36. Lovelock, J. (1992). A numerical model for biodiversity. *Philosophical Transactions of the Royal Society of London Series B, 338*, 365–373.

37. Lovelock, J. (1988). *The ages of Gaia* (p. 63). New York: Bantam Books.

38. Nowak, M. A., & Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. *Nature, 393*, 573–577.

39. Poincaré, H. (1896). *Calcul des probabilités* (2nd ed., 1912). Paris: Gauthier-Villars.

40. Popper, K. (1990). *A world of propensities* (pp. 18–19). Bristol: Thoemmes.

41. Porter, T. M. (1986). *The rise of statistical thinking* (p. 202). Princeton University Press.

42. Saunders, P. T. (1994). Evolution without natural selection: Further implications of the Daisyworld parable. *Journal of Theoretical Biology, 166*, 370.

43. Schmidt, A. (1971). *The concept of nature in Marx.* London: New Left Books.

44. Schuster, P., & Sigmund, K. (1983). Replicator dynamics. *Journal of Theoretical Biology, 100*, 536.

45. Sigmund, K. (1993). *Games of life: Explorations in ecology, evolution, and behaviour* (p. 201). London: Penguin Books.

46. Skyrms, B. (1996). *Evolution of the social contract.* Cambridge, UK: Cambridge University Press.

47. Smith, J. M. (1995). Review of "The Origins of Order." *New York Review of Books* (2 March), 30.

48. Volterra, V. (1931). *Leçons sur la théorie mathématique de la lutte pour la vie.* Paris: Gauthier-Villars.

49. Waldrop, M. (1992). *Complexity: The emerging science at the edge of order and chaos.* (p. 255). New York: Touchstone.

50. Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of "small-world" networks. *Nature, 393*, 440–442.

51. Wellmer, A. (1967). *Methodologie als Erkenntnistheorie.* Frankfurt: Suhrkamp. *Critical theory of society.* New York: Herder and Herder (1969, 1971).

52. Wolfram, S. (1988). Complex systems theory. In *Emerging syntheses in science: Proceedings of the founding workshops of the Santa Fe Institute* (p. 182). Addison-Wesley.

53. Wolfram, S. (1994). *Cellular automata and complexity: Collected papers.* Addison-Wesley-Longman.

54. Wuensche, A., & Lesser, M. J. (1992). *The global dynamics of cellular automata: An atlas of basin of attraction fields of one-dimensional cellular automata,* Santa Fe Institute Studies in the Sciences of Complexity. Reading, MA: Addison-Wesley.

Appendix: The Replicator Equation

The replicator equation describes the behavior of a population divided into n phenotypes E_1 to E_n with relative frequencies x_1 to x_n specified by the vector $\mathbf{x}$. In a population with more than one phenotype, the fitness of phenotype E_i is conventionally denoted by $f_i(\mathbf{x})$. The rate of increase $\dot{x}_i/x_i$ under selection is equal to the difference between the fitness of E_i and the mean fitness of the population. Thus $\dot{x}_i/x_i =$ (fitness of E_i) − (mean fitness), and the process of selection can be succinctly described by the *replicator equation* [23]

$$\dot{x}_i = x_i[f_i(\mathbf{x}) - \bar{f}(\mathbf{x})], \qquad i = 1, \ldots, n. \tag{3}$$

If phenotypes differ in fitness, this equation will cause $\mathbf{x}$ to vary, with changes in $\mathbf{x}$ driven by these differences.[8] If the environment remains constant and the fitness of each phenotype does not change ("constant selection" [11]), then the highest attainable mean fitness is the fitness of the most fit phenotype in the initial population, and this phenotype will go to fixation.

---

8 Equation 3 defines a dynamical system on the unit simplex,

$$S_n = \left\{ \mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n : \sum_i x_i = 1, \quad x_i \geq 0 \quad \text{for} \quad i = 1, \ldots, n \right\}.$$

Hofbauer and Sigmund [23] showed that the n-dimensional Lotka-Volterra equation can be mapped onto orbits of Equation 3, as well as those of the time-continuous version of Fisher's selection equation and of the hypercycle equation.