COMPLEX ADAPTIVE SYSTEMS
John H. Holland, Christopher G. Langton, and Stewart W. Wilson, advisors

Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence, John H. Holland

Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life, edited by Francisco J. Varela and Paul Bourgine

Genetic Programming: On the Programming of Computers by Means of Natural Selection, John R. Koza

Genetic Programming: The Movie, John R. Koza and James P. Rice

From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior, edited by Jean-Arcady Meyer, Herbert L. Roitblat, and Stewart W. Wilson

Intelligent Behavior in Animals and Robots, David McFarland and Thomas Bösser

Advances in Genetic Programming, edited by Kenneth E. Kinnear, Jr.

Genetic Programming II: Automatic Discovery of Reusable Programs, John R. Koza

Genetic Programming II Video: The Next Generation, John R. Koza

Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds, Mitchel Resnick

From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior, edited by Dave Cliff, Philip Husbands, Jean-Arcady Meyer, and Stewart W. Wilson

Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems, edited by Rodney A. Brooks and Pattie Maes

Comparative Approaches to Cognitive Science, edited by Herbert L. Roitblat and Jean-Arcady Meyer

Artificial Life: An Overview, edited by Christopher G. Langton

Evolutionary Programming IV: Proceedings of the Fourth Annual Conference on Evolutionary Programming, edited by John R. McDonnell, Robert G. Reynolds, and David B. Fogel

An Introduction to Genetic Algorithms, Melanie Mitchell

Catching Ourselves in the Act: Situated Activity, Interactive Emergence, and Human Thought, Horst Hendriks-Jansen

Toward a Science of Consciousness: The First Tucson Discussions and Debates, edited by Stuart R. Hameroff, Alfred W. Kaszniak, and Alwyn C. Scott

Genetic Programming: Proceedings of the First Annual Conference, edited by John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo

Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming, edited by Lawrence J. Fogel, Peter J. Angeline, and Thomas Bäck

Elements of Artificial Neural Networks, Kishan Mehrotra, Chilukuri K. Mohan, and Sanjay Ranka

From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, edited by Pattie Maes, Maja J. Mataric, Jean-Arcady Meyer, and Stewart W. Wilson

Advances in Genetic Programming, Volume 2, edited by Peter J. Angeline and Kenneth E. Kinnear, Jr.

Growing Artificial Societies: Social Science from the Bottom Up, Joshua M. Epstein and Robert Axtell

Growing Artificial Societies

Social Science from the Bottom Up

Joshua M. Epstein
Robert Axtell

A product of the 2050 Project, a collaborative effort of the Brookings Institution, the Santa Fe Institute, and the World Resources Institute.

Brookings Institution Press
Washington, D.C.

The MIT Press
Cambridge, Massachusetts & London, England

CD-ROM Edition

A CD-ROM edition of *Growing Artificial Societies* is available. It operates on both Macintosh and Windows platforms and contains the complete text as well as the animations, which can be viewed in their entirety.

For ordering information, contact the Brookings Institution Press, 1775 Massachusetts Avenue N.W., Washington, D.C. 20036. Telephone: 202-797-6258 or 1-800-275-1447. Fax: 202-797-6004. E-mail: BIBOOKS@Brook.edu.

Introduction

Herbert Simon is fond of arguing that the social sciences are, in fact, the "hard" sciences.[1] For one, many crucially important social processes are *complex*. They are not neatly decomposable into separate subprocesses—economic, demographic, cultural, spatial—whose isolated analyses can be aggregated to give an adequate analysis of the social process as a whole. And yet, this is exactly how social science is organized, into more or less insular departments and journals of economics, demography, political science, and so forth. Of course, most social scientists would readily agree that these divisions are artificial. But, they would argue, there is no natural methodology for studying these processes together, as they coevolve.

The social sciences are also hard because certain kinds of controlled experimentation are hard. In particular, it is difficult to test hypotheses concerning the relationship of individual behaviors to macroscopic regularities, hypotheses of the form: If individuals behave in thus and such a way—that is, follow certain specific rules—then society as a whole will exhibit some particular property. How does the heterogeneous microworld of individual behaviors generate the global macroscopic regularities of the society?[2]

Another fundamental concern of most social scientists is that the rational actor—a perfectly informed individual with infinite computing capacity who maximizes a fixed (nonevolving) exogenous utility function—bears little relation to a human being.[3] Yet, there has been no natural methodology for relaxing these assumptions about the individual.

Relatedly, it is standard practice in the social sciences to suppress real-world agent heterogeneity in model-building. This is done either explicitly, as in representative agent models in macroeconomics,[4] or

1. Simon [1987]; the same point has been made by Krugman [1994: xi-xiii].
2. This is to be distinguished from the very different problem of determining what rules are actually employed by individual humans, a topic studied in experimental economics and other fields of behavioral science.
3. A recent statement of this basic concern is Aaron [1994].
4. Kirman [1992] makes this point forcefully.

implicitly, as when highly aggregate models are used to represent social processes. While such models can offer powerful insights, they "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.

Finally, it is fair to say that, by and large, social science, especially game theory and general equilibrium theory, has been preoccupied with static equilibria, and has essentially ignored time dynamics. Again, while granting the point, many social scientists would claim that there has been no natural methodology for studying nonequilibrium dynamics in social systems.

We believe that the methodology developed here can help to overcome these problems. This approach departs dramatically from the traditional disciplines, first in the *way* specific spheres of social behavior—such as combat, trade, and cultural transmission—are treated, and second in the way those spheres are *combined*.

"Artificial Society" Models

We apply *agent-based* computer modeling techniques to the study of human social phenomena, including trade, migration, group formation, combat, interaction with an environment, transmission of culture, propagation of disease, and population dynamics. Our broad aim is to begin the development of a computational approach that permits the study of these diverse spheres of human activity from an evolutionary perspective as a single social science, a *transdiscipline* subsuming such fields as economics and demography.

This modeling methodology has a long lineage. Beginning with von Neumann's work on self-reproducing automata [1966], it combines elements of many fields, including cybernetics (for example, Ashby [1956], Wiener [1961]), connectionist cognitive science (for example, Rumelhart and McClelland [1986]), distributed artificial intelligence (for example, Gasser and Huhns [1989]), cellular automata (for example, Wolfram [1994], Toffoli and Margolus [1987], Gutowitz [1991]), genetic algorithms (for example, Holland [1992]), genetic programming (Koza [1992, 1994]), artificial life (for example, Langton [1989, 1992, 1994], Langton *et al.* [1992], Brooks and Maes [1994]), and individual-based modeling in biology (for example, Haefner and Crist [1994] and Crist and Haefner [1994]). However, there have been very few attempts to bring these literatures to bear on social science.[5]

The first concerted attempts to apply, in effect, agent-based computer modeling to social science explicitly are Thomas Schelling's. In a classic series of papers—"Models of Segregation" [1969], "On the Ecology of Micromotives" [1971a], and "Dynamic Models of Segregation" [1971b]—and later in the book *Micromotives and Macrobehavior* [1978], Schelling anticipated many of the themes encountered in the contemporary literature on agent-based modeling, social complexity, and economic evolution. Among other things, Schelling devised a simple spatially distributed model of the composition of neighborhoods, in which agents prefer that at least some fraction of their neighbors be of their own "color." He found that even quite color-blind preferences produced quite segregated neighborhoods.[6]

But Schelling's efforts were constrained by the limited computational power available at that time. It is only in the last decade that advances in computing have made large-scale agent-based modeling practical. Recent efforts in the social sciences to take advantage of this new capability include the work of Albin and Foley [1990], Arifovic [1994], Arifovic and Eaton [1995], Arthur [1991, 1994], Arthur *et al.* [1994], Axelrod [1993, 1995], Carley [1991], Danielson [1992, 1996], Gilbert and Doran [1994], Gilbert and Conte [1995], Holland and Miller [1991], Kollman, Miller, and Page [1992, 1994], Marimon, McGrattan, and Sargent [1990], Marks [1992], Nagel and Rasmussen [1994], Tesfatsion [1995], and Vriend [1995]. Additionally, computer scientists interested in questions of distributed artificial intelligence (DAI), decentralized decisionmaking, and game theory have been actively researching multiagent systems. Important work here includes that of Huberman and coworkers (Huberman [1988], Huberman and Glance [1993, 1996], Glance and Huberman [1993, 1994a, 1994b], Huberman and Hogg [1995], Youssefmir and Huberman [1995]), Maes [1990], Miller and Drexler [1988], and Resnick [1994]. Biologists have even built models in which a population of agents representing humans exploits ecological resources (Bousquet, Cambier, and Morand [1994]).

In what follows we shall refer to agent-based models of social pro-

5. An important exception is Steinbruner's *The Cybernetic Theory of Decision* [1974].
6. Related work includes that of Vandell and Harrison [1978].

cesses as *artificial societies*.[7] In this approach fundamental social structures and group behaviors emerge from the interaction of individuals operating in artificial environments under rules that place only bounded demands on each agent's information and computational capacity. We view artificial societies as *laboratories*, where we attempt to "grow" certain social structures in the computer—or *in silico*—the aim being to discover fundamental local or micro mechanisms that are sufficient to *generate* the macroscopic social structures and collective behaviors of interest.[8]

In general, such computer experiments involve three basic ingredients: agents, an environment or space, and rules. A brief word on these may be in order before discussing the particular artificial society presented in this book.

Agents

Agents are the "people" of artificial societies. Each agent has internal states and behavioral rules. Some states are fixed for the agent's life, while others change through interaction with other agents or with the external environment. For example, in the model to be described below, an agent's sex, metabolic rate, and vision are fixed for life. However, individual economic preferences, wealth, cultural identity, and health can all change as agents move around and interact. These movements, interactions, changes of state all depend on rules of behavior for the agents and the space.

7. This term apparently originates with Builder and Banks [1991]; see also Bankes [1994].

8. So-called micro-simulation techniques, developed by social scientists at the dawn of the modern computer era, are philosophically similar to agent-based approaches insofar as both attempt to model social phenomena in a highly disaggregated way. An early pioneering work of this type is Orcutt *et al.* [1961], who wrote:

> Our socioeconomic system is a complicated structure containing millions of interacting units, such as individuals, households, and firms. It is these units which actually make decisions about spending and saving, investing and producing, marrying and having children. It seems reasonable to expect that our predictions would be more successful if they were based on knowledge about these elemental decision-making units—how they behave, how they respond to changes in their situations, and how they interact.

In comparison to agent-based modeling, micro-simulation has more of a "top-down" character since it models behavior via equations statistically estimated from aggregate data, not as resulting from simple local rules.

Environment

Life in an artificial society unfolds in an environment of some sort. This could be landscape, for example, a topography of renewable resource that agents eat and metabolize. Such a landscape is naturally modeled as a lattice of resource-bearing sites. However, the environment, the medium over which agents interact, can be a more abstract structure, such as a communication network whose very connection geometry may change over time. The point is that the "environment" is a medium separate from the agents, *on* which the agents operate and *with* which they interact.

Rules

Finally, there are rules of behavior for the agents and for sites of the environment. A simple movement rule for agents might be: Look around as far as you can, find the site richest in food, go there and eat the food. Such a rule couples the agents to their environment. One could think of this as an *agent-environment* rule. In turn, every site of the landscape could be coupled to its neighbors by cellular automata (see below) rules. For example, the rate of resource growth at a site could be a function of the resource levels at neighboring sites. This would be an *environment-environment* rule. Finally, there are rules governing *agent-agent* interactions—mating rules, combat rules, or trade rules, for example.

Object-Oriented Implementation

Contemporary object-oriented programming (OOP) languages are particularly natural ones for agent-based modeling. Objects are structures that hold both data and procedures. Both agents and environmental sites are naturally implemented as objects. The agent's data fields (its *instance variables*) represent its internal states (for example, sex, age, wealth). The agent's procedures (*methods*) are the agent's rules of behavior (for example, eating, trading). This *encapsulation* of internal states and rules is a defining characteristic of OOP and greatly facilitates the construction of agent-based models.[9]

9. For more on the software engineering aspects of artificial societies, see Appendix A.

Social Structures Emerge

Typically, we release an initial population of agent-objects into the simulated environment (a lattice of site-objects) and watch for organization into recognizable macroscopic social patterns. The formation of tribes or the emergence of certain stable wealth distributions would be examples. Indeed, the defining feature of an artificial society model is precisely that *fundamental social structures and group behaviors emerge from the interaction of individual agents operating on artificial environments under rules that place only bounded demands on each agent's information and computational capacity.* The shorthand for this is that we "grow" the collective structures "from the bottom up."

The Sugarscape Model

While the "bottom-up" approach to social science is quite general—as discussed at greater length in our concluding chapter—the primary focus of the present work is a particular instance of the artificial society concept that has come to be known as *The Sugarscape Model.* A brief summary of each chapter follows.

Life and Death on the Sugarscape

In Chapter II we introduce the sugarscape, a spatial distribution, or landscape, of generalized resource that agents like to "eat." The landscape consists of variously shaped regions, some rich in sugar, some relatively impoverished. Agents are born onto the sugarscape with a vision, a metabolism, and other genetic attributes. In Chapter II their movement is governed by a simple local rule. Paraphrasing, it amounts to the instruction: "Look around as far as your vision permits, find the spot with the most sugar, go there and eat the sugar." Every time an agent moves, it "burns" some sugar—an amount equal to its metabolic rate. Agents die if and when they burn up all their sugar.

A remarkable range of phenomena emerges from the interaction of these simple agents. The ecological principle of carrying capacity—that a given environment can support only some finite population—quickly becomes evident. When "seasons" are introduced, migration is observed. Migrators can be interpreted as environmental refugees, whose immigration boosts population density in the receiving zone, intensifying the competition for resources there—a dynamic with "national security" implications. Since agents are accumulating sugar at all times, there is always a distribution of wealth—measured in sugar—in the agent society. Does the wealth distribution mimic anything observed in human societies? Under a great variety of conditions the distribution of wealth on the sugarscape is highly skewed, with most agents having little wealth. Highly skewed distributions of income and wealth are also characteristic of actual human societies, a fact first described quantitatively by the nineteenth-century mathematical economist Vilfredo Pareto.[10] Thus we find the first instance of a qualitative similarity between extant human societies and artificial society on the sugarscape.

A CompuTerrarium

As a practical matter, if such highly skewed wealth distributions are immutable laws of nature, as some have claimed, then there is little hope of greater economic equity in society. A tool like Sugarscape can function as a kind of laboratory—a *CompuTerrarium*—where we alter agent behavioral rules, such as those governing trade or inheritance, in order to see how immutable this kind of distribution really is.

Agent Social Networks

Humans can be connected socially in various ways: genealogically, culturally, and economically, for example. Indeed, one of the things that makes humans complicated, conflicted, and interesting is that they can belong to many different communities, or social networks, at once. These networks change over time. And, most interestingly, group loyalties can come into profound conflict, as when brothers (members of a family group) fought each other (as members of competing political groups) in the American Civil War. One theme that runs through this entire book is *social connection*. In each chapter the local rules governing agent behavior permit us to define certain kinds of agent social networks. We represent such networks as graphs and track their evolution over time and space. In particular, Chapter II explores social networks of neighbors.

10. See Persky [1992] for an overview of the so-called Pareto law.

Up to this point, collective phenomena have emerged from interactions within a *single* population of agents. In Chapter III we "grow" distinct populations—cultural formations—of agents.

Sex, Culture, and Conflict: The Emergence of History

Indeed, the aim of Chapter III is to "grow" an entire history of an artificial civilization—a *proto-history*, as we call it. The storyline is as follows:

> *In the beginning, a small population of agents is randomly scattered about a landscape. Purposeful individual behavior leads the most capable or lucky agents to the most fertile zones of the landscape; these migrations produce spatially segregated agent pools. Though less fortunate agents die on the wayside, for the survivors life is good: food is plentiful, most live to ripe old ages, populations expand through sexual reproduction, and the transmission of cultural attributes eventually produces spatially distinct "tribes." But their splendid isolation proves unsustainable: populations grow beyond what local resources can support, forcing tribes to expand into previously uninhabited areas. There the tribes collide and interact perpetually, with penetrations, combat, and cultural assimilation producing complex social histories, with violent phases, peaceful epochs, and so on.*

This, then, is the social story we wish to "grow," from the bottom up. We will need a number of behavioral ingredients, each of which generates insights of its own.

The first ingredient of the proto-history is sexual reproduction. Like other rules that agents execute in the model, the "sex code" is completely *local* and very simple. Yet a rich variety of demographic trajectories is observed. For instance, populations—and population *densities* on the sugarscape—can fluctuate dramatically. Because mating is local, reproduction can cease and the population can crash if population becomes too sparse, or *thin*. Bottom up models such as Sugarscape suggest that certain cataclysmic events—like extinctions—can be brought on endogenously, without external shocks (like meteor impacts) through local interactions alone. Scientists have long been fascinated by the oscillations, intermittencies, and "punctuated equilibria" that are observed in real plant and animal populations. They have modeled these phenomena using "top-down" techniques of nonlinear dynamical systems, in which aggregate state variables are related through, say, differential equations. Yet we demonstrate that all these dynamics can be "grown" from the "bottom-up." And, when they are conjoined with the processes of combat, cultural exchange, and disease transmission, a vast panoply of "possible histories," including our proto-history, is realized on the sugarscape.

It is possible to *observe* evolutionary processes as they alter the genetic composition of our artificial society. For example, we expect that, over many generations, selection pressures will operate in favor of agents having relatively low metabolism and high vision. In fact, precisely this behavior emerges on the sugarscape. In Chapter III we assign a color to each agent according to its metabolism, then watch society change color as selection pressures "weed out" high metabolism individuals over time. Selection also operates on agent vision. There is a kind of genetic algorithm (GA) at work here, though we have not specified any "fitness function" beforehand. The topic of fitness, and the need to define it in coevolutionary terms, is addressed.

With a sexual reproduction rule in place, it is natural to study genealogy using social networks. It is very interesting to watch these "family trees" branch out across the sugarscape.

The next ingredient of the proto-history is tribe formation. How do tribes form? How does "social speciation" occur? To address these questions, we give agents cultural attributes and rules for their local transmission.[11] Cultural formations then "grow" from the bottom up. We represent cultural connections as lines between agents who have similar cultural attributes. These cultural connection networks expand, contract, and deform over time.

Finally, when agents of one cultural "tribe" encounter agents of a different tribe they may engage in a primitive kind of combat. That is, agents of opposite tribes may plunder one another for sugar. However, they are not so stupid as to attack agents who are capable of defeating them, or to attack an agent of a different tribe when there are others from that tribe in the vicinity who can retaliate successfully. Thus the combat rule results in agent movement patterns very different from the standard "eat all you can find" rule. We experiment with a variety of combat rules in Chapter III.

---

11. In Chapter IV we let economic preferences depend on these cultural attributes. Then, when cultural interchange and economic processes are both active, we have a model in which agent preferences change endogenously, in contrast to the assumption of fixed preferences standard in economic theory.

Sugar and Spice: Trade Comes to the Sugarscape

In Chapter IV a second commodity—"spice"—is added to the resource landscape, and each agent is given a corresponding metabolism for spice. The relative size of an agent's sugar and spice metabolisms determines its *preferences* for the two resources. The agents move around the landscape searching for those sites that best satisfy their preferences. Each agent must at all times possess positive quantities of *both* sugar and spice, or it dies.

Agents are then given the ability to trade sugar and spice. All trade is conducted in a decentralized fashion between neighboring agents, so-called bilateral exchange. Each pair of agents engaged in trade "bargains" to a local price and then exchanges goods only if it makes both agents better off. The main topics investigated in the chapter concern the relationship of local prices to the formation of a single "market-clearing" price and the welfare properties of these artificial markets.[12] These issues are investigated for two distinct classes of agents: the idealized economic agents found in economics textbooks and agents that are non-neoclassical insofar as they have finite lives and evolving preferences.

Markets of Neoclassical Agents

A crucial question is the following: *Under what conditions (for example, rules of agent behavior) will local prices converge to a market-clearing (general equilibrium) price?* We find that an equilibrium price is approached when our artificial society consists of a large number of infinitely lived agents having fixed preferences who trade for a long time. However, the resource allocations that obtain, although locally optimal, fail to be globally optimal. That is, there are additional gains from trade that our agents are unable to extract. The reason is that, while bilateral exchange is pushing the artificial economy toward a globally optimal configuration, production activities (resource gathering) are constantly modifying this configuration. These two competing processes—exchange and production—yield an economy that is perpetually out of equilibrium.

Because trade can simply be turned on or off in models of this type, we can study the effects of trade on other social variables. In particular, we find that the carrying capacity of the environment is *increased* when

12. On the notion of an artificial economy, see Lane [1993].

agents trade.[13] However, this salutary result does not come free, for under some circumstances trade increases societal inequality.

There are further implications for the welfare properties of markets. The "equilibrium" price that emerges under bilateral trade has a different character than the general equilibrium price of neoclassical theory; it is *statistical* in nature. One implication of statistical equilibrium is that agents having identical preferences and endowments can end up in very different welfare states through decentralized trade: they encounter different people, bargain to different prices, and trade different quantities, producing initially small differences in their respective welfare states, which may be amplified with time. This phenomenon is termed *horizontal inequality*.

Markets of Non-Neoclassical Agents

In neoclassical economic theory individual economic agents live forever and have fixed preferences. We give agents finite lives and the ability to reproduce sexually (as in Chapter III) and study the effects on economic behavior. The primary result of adding new agents to our artificial economy is to add variance to the distribution of trade prices in the sugar-spice market. This occurs because as new agents are born it takes time to have their internal valuations brought into line with those prevailing in the marketplace. The amount of price dispersion this effect produces increases as average agent lifetime decreases. Generally, increased variance in price corresponds to increased horizontal inequality, so the welfare properties of markets are further eroded by finite agent lives.

Preferences are permitted to evolve by coupling them to the cultural exchange process introduced in Chapter III. This yields several interesting economic phenomena. Agents whose preferences change from one period to the next find that their accumulated holdings—quite satisfactory in the previous period—may not satisfy their current wants, so they are more willing to trade than when preferences are fixed. Overall, we find that total trade volumes are larger with evolving preferences. Too, there is much more variation in prices under such circumstances, and the average price follows a kind of "random drift" process. Nothing like

13. In Chapter VI a set of Sugarscape model runs in which this phenomenon plays a crucial role is described. The evolutions of two societies, identical in all respects except that one engages in interagent trade while the other does not, are compared and contrasted. The nontraders end up extinct, while the traders are progenitors of a prosperous civilization.

the equilibrium of neoclassical theory emerges. Of course, the laissez-faire argument is precisely that markets, left to their own devices, allocate goods and services efficiently. The theoretical case for this is the so-called First Theorem of welfare economics. However, when markets fail to arrive at equilibrium, the First Welfare Theorem does not apply, and this case for laissez-faire is undermined.

Credit Networks

When agents are permitted to enter into credit relationships with one another very elaborate borrower-lender networks result. Agents borrow for purposes of having children. If an agent is of childbearing age but has insufficient wealth to produce offspring, then it will ask each of its neighbors in turn if they are willing to loan it the sugar it needs to become "fertile."[14] Prospective lenders assess the borrower's ability to repay a loan based on the borrower's past income. Once a loan has been consummated, it is repaid when due unless the borrower has insufficient accumulation, in which case it is renegotiated.

The credit connections that result from these rules are very dynamic. In order to study these relationships, graphs of creditor-debtor arrangements are shown in which each agent is a vertex and edges are drawn between borrowers and lenders. These graphs are updated each time period, thus showing the evolution of credit structures spatially. Unexpectedly, some agents turn out to be borrowers and lenders simultaneously, and this is most effectively displayed as a hierarchical graph, with agents who are only lenders placed at the top of the hierarchy and those who are only borrowers positioned at the bottom.

Social Computation

Yet another kind of social network is a trade partner graph, in which each vertex represents an agent and edges are drawn between agents who have traded with each other during a particular time period. Such graphs not only represent social relations but also depict the physical flow of commodities—sugar flows one way along an edge, while spice gets transferred in the opposite direction. These graphs—webs of economic intercourse—link agents who may be spatially quite distant, even though all trade is local, that is, between neighbors. These networks are

14. As we define this term in Chapter III, "fertility" includes an economic component.

ever-changing with time and are displayed in Chapter IV as animations.

During trade each agent acts to improve its welfare, that is, each participant optimizes its own utility function. A main question addressed in Chapter IV is: To what extent does individual (local) optimality result in overall social (global) optimality? Consider each agent to be an autonomous processing node in a computer, the agent society. *Individual* agents (nodes) compute only what is best for themselves, not for the society of agents *as a whole*. Over time, the trade partner network describes the evolution of connections between these computational nodes. Thus there is a sense in which *agent society acts as a massively parallel computer, its interconnections evolving through time.* This idea is fleshed out in Chapter IV.

Another important area where agent-based techniques apply very naturally is that of public health—epidemiology and immunology. We study this in Chapter V.

Disease Agents

Humans and infectious parasites have been coevolving for a long time. Certainly, it would be hard to overstate the impact of infectious diseases on human society. William McNeill [1976] has argued that infectious diseases played crucial roles in the spread of religions, political dominions, and social practices ranging from prohibitions on the consumption of pork to caste systems of the sort seen even today in India. In our own time, HIV has obviously had important sociopolitical impacts across a wide variety of groups on many continents. In light of all this, there is every reason to include epidemiology in social science. But there is equal reason to include social science in epidemiology! After all, the Black Death—*Pasteurella pestis*—could not have spread from China to Europe without human technological advances and commercial intercourse, notably in navigation and shipping. Needless to say, military conquest and migration have been equally efficient vehicles for the dissemination of infectious disease agents.

One aim of Chapter V, then, is to break down an artificial division between fields, presenting an adaptive agents model in which the spread of infectious diseases interacts with other social processes. We also hope to advance epidemiology proper, in several respects. First, our treatment of space differs fundamentally from that found in typical mathematical models. Also, mathematical epidemiology typically divides society into homogeneous subpopulations—compartments such as susceptibles and infectives within which there is no variation among individuals. In actuality, substantial variation exists; agents are heterogeneous precisely in that they have different immune systems. We endow every agent with its own adaptive immune system. Our immunology is, of course, very simple and highly idealized. Nonetheless, the explicit incorporation of an immune model into the epidemic model enriches and unifies the resulting picture. Important phenomena including immunological memory and the persistence of childhood diseases emerge very naturally. Moreover, since infected agents suffer a metabolic increase in our model, the epidemic dynamics affect (through the agents' metabolism-dependent utility functions) their movements and economic behavior.

A Society Is Born

Over the course of these chapters, the agents' behavioral repertoire grows to include movement, resource gathering, sexual reproduction, combat, cultural transmission, trade, inheritance, credit, pollution, immune learning, and disease propagation. In Chapter VI, we turn on all these dimensions and explore the complex, multidimensional artificial society that emerges. The book then concludes with a discussion of variations on, and extensions of, the current Sugarscape model.

Artificial Societies versus Traditional Models

Differences between our approach and certain other methodologies (for example, game theory) have already been noted. But additional ways in which artificial societies differ from traditional mathematical models and work in the field of artificial life (ALife) also merit review.

Heterogeneous Agent Populations

In a traditional ordinary differential equation model of an epidemic, the total population is divided into subpopulations of, say, susceptibles and infectives. These subgroups are *homogeneous*; nothing distinguishes one member from another. Similarly, in ecosystem models there are predators and prey, but homogeneity is assumed within each species. In macroeconomics the use of representative agents assumes away real-world heterogeneity.

By contrast, in agent-based models there is no such aggregation. The spatially distributed population is heterogeneous and consists of distinct agents, each with its own genetically and culturally transmitted traits (attributes and rules of behavior). Individual traits can change—adapt—in the course of each agent's life, as a result of interaction with other agents, with diseases, and with an environment. And, in evolutionary time (which can elapse quickly on computers), selection pressures operate to alter the distribution of traits in populations.

Space Distinct from the Agent Population

In *ordinary* differential equation models there is *no* spatial component at all. Susceptibles and infectives, predators and prey, interact in time but not in space.[15] In *partial* differential equation models there is a physical space **x**, but the state variables representing agent populations (such as the infection level) are continuous in **x**.

By contrast, in Sugarscape the agents live on a two-dimensional lattice, but are completely distinct from it. When diseases occur, they are passed from agent to agent, but the environment—and the agent's rules of interaction with it—affects the spatial distribution of agents, and hence the epidemic dynamics. Likewise, it affects the dynamics of trade, of combat, of population growth, of cultural transmission, and so on.

Agent-Environment and Agent-Agent Interactions according to Simple Local Rules

In the simplest form of our model agents are born with various genetic attributes, one of which is vision, and their rule of behavior is to look for the best unoccupied resource location. Their search is local; no agent has global information. Similarly, when we introduce trade there is no computation by any agent—or any "super agent" such as the Walrasian auctioneer—of a market-clearing price. Price formation takes place by a process of completely decentralized bilateral trade between neighbors. Under some conditions prices converge to a statistical equilibrium. This artificial economy stands in stark contrast to the neoclassical general equilibrium formalism, which relies on aggregate excess demand func-

15. These points apply with equal force to aggregate modeling of the system dynamics type (e.g., Stella, Dynamo).

tions—or some other form of *global* information—for the existence of and convergence to equilibrium.

Focus on Dynamics

One need not confine one's attention to equilibria, as is done in much of mathematical social science.[16] A social system's rest points, its equilibria, may be the most analytically tractable configurations, but it is by no means clear that they are either the most important or interesting configurations. Indeed, in much of what follows it will be the dynamic properties of the model, rather than the static equilibria, that are of most interest. In Chapter III, for instance, we study the dynamics of cultural transmission. Over thousands of time periods we see the sudden appearance of cultural "fads" and their irregular spatial propagation. These out-of-equilibrium dynamics seem far more interesting than the static cultural equilibrium into which the system is finally absorbed. With artificial societies built from the bottom up the transients are no more difficult to study than the equilibria.[17]

Beyond Methodological Individualism

Our point of departure in agent-based modeling is the individual: We give agents rules of behavior and then spin the system forward in time and see what macroscopic social structures emerge. This approach contrasts sharply with the highly aggregate perspective of macroeconomics, sociology, and certain subfields of political science, in which social aggregates like classes and states are posited *ab initio*. To that extent our work can be accurately characterized as "methodologically individualist." However, we part company with certain members of the individualist camp insofar as we believe that the collective structures, or "institutions," that emerge

16. Proofs of the *existence* of general economic equilibrium, *refinements* of equilibrium concepts in game theory (for example, Nash equilibrium), theories of equilibrium *selection* when multiple equilibria exist, and methods for evaluating the *stability* of equilibria are dominant themes in this literature.

17. When a model produces some interesting transient for which no explanation is immediately available, one can simply recreate the realization in question (by keeping track of seeds to the random number generators) and then glean data (noiselessly) from the agent population, data that will serve as the basis for analyses of the observed output. Or it may be useful to pause the model at some particular point in its execution and query particular agents for their state information.

can have feedback effects in the agent population, altering the behavior of individuals.[18] Agent-based modeling allows us to study the interactions between individuals and institutions.[19]

Collective Structures Emerge from the Bottom Up

A general equilibrium price, when obtained in our model, is an example of an emergent entity. In the usual general equilibrium story it is assumed that every agent "takes" a price issued from the top down, by the so-called Walrasian auctioneer. By contrast we "grow" an equilibrium price from the bottom up through local interactions alone, dispensing with the artifice of the auctioneer and the entire aggregate excess demand apparatus. Many other collective structures emerge in our artificial society: tribes of agents, stationary wealth distributions, and collective patterns of movement, for example.

Artificial Societies versus ALife

The Sugarscape synthesizes two "threads" from the ALife research tapestry. One is the field of cellular automata, or CA. A CA consists of a lattice of cells, or sites. At every time, each cell has a value, such as 0 or 1, black or white, "on" or "off," or a color selected from a set of colors, such as {red, blue, green}. These values are updated iteratively according to a fixed rule that specifies exactly how the "new" value of every site is computed from its own present value and the values of its immediate neighbors. Although, properly speaking, the pedigree of CAs extends at least as far back as von Neumann's work on self-replicating automata, the most familiar example is John Conway's game, "Life."[20] Cellular automata

18. Varying positions of methodological individualism are reviewed in Hausman [1992] and Arrow [1994].

19. The term "bottom up" can be somewhat misleading in that it suggests unidirectionality: everything that emerges is outside the agent. But in models with feedback from institutions to individuals there is emergence inside the agents as well.

20. The rules of "Life" are very simple:

1. A cell in state 0 switches to state 1 if three of its eight lattice neighbors are in state 1; otherwise, it stays in state 0.
2. A cell in state 1 stays in that state if two or three of its neighbors are in state 1; otherwise, it switches to state 0.
3. Each cell is updated once per time period.

have been created as models of fluid flow [Doolen *et al.* 1990], earthquakes [Bak and Tang, 1989], clouds [Nagel and Raschke, 1992], forest fires [Bak, Chen, and Tang, 1990], biological systems [Ermentrout and Edelstein-Keshet, 1993], and a vast array of other complex *spatial* processes. The sugarscape proper—as opposed to the agents—is modeled as a CA.

Another major line of work in the ALife field does not involve an explicit space, but rather concerns the interaction of agents in a "soup"—a (space-less) environment in which each agent may interact directly with every other agent. The agents—unlike the cells in "Life"—have many different attributes (e.g., internal states and rules of behavior) which change through social interaction; see, for instance, Arthur [1994].

Sugarscape agents are very simple by design. In particular, we specify the agents' behavioral rules and watch for the emergence of important macro-social structures, such as skewed wealth distributions. The agents' rules do vary, but only parametrically, not structurally. For instance, every agent has a utility function. Culturally varying parameters enter into these utility functions, but the *algebraic form* of the utility function remains fixed, as does the agent's practice of *maximizing* the function. We say, then, that the microrules governing economic behavior adapt parametrically, not structurally. Similarly, agent immune systems adapt parametrically to new disease strains. The game in this particular research has been to design the simplest possible agents and explore what happens when they interact. As we shall see, the analytical challenges are already formidable. However, this is not the only possible game.

Instead of giving all agents the same rule, one might begin with a population of agents, each with a different rule, and allow selection pressure to change the rule distribution over time. In other words, no individual agent adapts, but (as in evolutionary game theory) those who prosper replicate and those doing poorly eventually die out. Over time, the rule distribution evolves. Society "learns" though individuals do not.[21]

Another modeling avenue is essentially to move the evolutionary process inside the agent. Here, each individual entertains a number of behavioral rules. Successful rules are promoted, while failures are

Under these rules a random initial distribution of black and white sites gives rise to a spectacular world of blinkers, wiggly "snakes," self-replicating "gliders," and stable structures on the lattice. For more on Life, see Sigmund [1993: 10–15, 27–39].

21. Examples of this approach in the context of the iterated prisoner's dilemma include Axelrod [1987], Miller [1989], and Lindgren [1992].

demoted, so that evolutionary learning occurs "within" the agent; Arthur [1994] is an example.

Yet more complex are models in which agents, in effect, "invent" entirely novel behavioral rules. Classifier systems (Holland [1992]) and neural networks (Rumelhart and McClelland [1986] and McClelland and Rumelhart [1986]) have been used in such models; see, for example, Marimon, McGrattan, and Sargent [1990] and Vriend [1995].

Cellular Automata + Agents = Sugarscape

In any event, if the pure CA is a space with no agents living on it, and the pure adaptive agents model represents agent kinetics with no underlying space, then the Sugarscape model is a synthesis of these two research threads. There is an underlying space—a "sugarscape"—that is a CA. But, populations of agents live on the CA.[22] The agents interact with one another *and* they interact with the environment. Interagent dynamics affect environmental dynamics, which feed back into the agent dynamics, and so on. *The agent society and its spatial environment are coupled.*[23]

Toward Generative Social Science: Can You Grow It?

The broad aim of this research is *to begin the development of a more unified social science, one that embeds evolutionary processes in a computational environment that simulates demographics, the transmission of culture, conflict, economics, disease, the emergence of groups, and agent coadaptation with an environment, all from the bottom up.* Artificial society–type models may change the way we think about *explanation* in the social sciences.

22. Other models in which agents inhabit a landscape include Holland's *Echo* [1992: 186–198], Ackley and Littman [1992], Yeager's *PolyWorld* (Yeager [1994], Wolff and Yeager [1994: 170–171]), and *BioLand* of Werner and Dyer [1994].

23. Heuristically, one thinks of an artificial society as a discrete time dynamical system in which the vector **A** of all agent internal states and the vector **E** of all environmental states interact as a high-dimensional discrete dynamical system of the general form:

$$\mathbf{A}^{t+1} = \mathbf{f}(\mathbf{A}^t, \mathbf{E}^t)$$
$$\mathbf{E}^{t+1} = \mathbf{g}(\mathbf{A}^t, \mathbf{E}^t)$$

where the vector functions $\mathbf{f}(\bullet)$ and $\mathbf{g}(\bullet)$ map the space of all states at time t to the space at t+1.

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.[24] And that, after all, is a central aim. As social scientists, we are presented with "already emerged" collective phenomena, and we seek microrules that can generate them.[25] We can, of course, use statistics to test the match between the true, observed, structures and the ones we grow.[26] But the ability to grow them—greatly facilitated by modern object-oriented programming—is what is new. Indeed, it holds out the prospect of a new, *generative*, kind of social science.[27]

24. This usage of the term "sufficient" is similar to that of cognitive scientists Newell and Simon [1972: 13].

25. There may be many microspecifications that will do as well—the mapping from micro-rules to macrostructure could be *many to one*. In the social sciences, that would be an embarrassment of riches; in many areas, *any to one* would be an advance.

26. Issues of agent-based model validation—objectives, methods, and software tools—are discussed in Axtell and Epstein [1994].

27. On artificial societies and generative social science, see Epstein and Axtell [1996]. Further discussion of generative social science appears in Chapter VI.

II

Life and Death on the Sugarscape

In this chapter the simplest version of our artificial world is described. A single population of agents gathers a renewable resource from its environment. We investigate the distribution of wealth that arises among the agents and find that it is highly skewed. It is argued that such distributions are *emergent* structures. Other emergent phenomena associated with mass agent migrations are then studied. Social networks among neighboring agents are illustrated and their significance is discussed. Finally, it is argued that artificial societies can serve as laboratories for social science research.

In the Beginning . . . There Was Sugar

Events unfold on a "sugarscape." This is simply a spatial distribution, or topography, of "sugar," a generalized resource that agents must eat to survive. The space is a two-dimensional coordinate grid or lattice. At every point (x, y) on the lattice, there is both a sugar level and a sugar capacity, the capacity being the maximum value the sugar level can take at that point. Some points might have no sugar (a level of zero) and low capacity, others might have no sugar but large capacity—as when agents have just harvested all the sugar—while other sites might be rich in sugar and near capacity.

The Sugarscape software system (that is, the computer program proper) permits one to specify a variety of spatial distributions of levels and capacities. But let us begin with the particular sugarscape shown in figure II-1, which consists of 2500 locations arranged on a 50 x 50 lattice with the sugar level at every site initially at its capacity value.

The sugar score is highest at the peaks in the northeast and southwest quadrants of the grid—where the color is most yellow—and falls off in a series of terraces.[1] The sugar scores range from some maximum—

1. Terms like "peak" or "mountain" are not used to suggest physical elevation, but to denote regions of high sugar level.