AGENT-BASED MODELS

Nigel Gilbert
University of Surrey, Guildford, UK

1. THE IDEA OF AGENT-BASED MODELING

Agent-based modeling is a new analytical method for the social sciences, but one that is quickly becoming popular. This short book explains what agent-based modeling is. It also warns of some dangers and describes typical ways of doing agent-based modeling. Finally, it offers a range of examples from many of the social sciences.

This first chapter begins with a brief overview of agent-based modeling before contrasting it with other, perhaps more familiar forms of modeling and describing several examples of current agent-based modeling research. Chapter 2 goes into more detail, considering a range of methodological and theoretical issues and explaining what "agents" are. Chapter 3 dives into the specifics of building agent-based models, reviewing available software platforms and showing step by step how one can build an agent-based model using one of these. Chapter 4 provides some practical advice about designing agent-based models, using them in social science research, and publishing articles based on agent-based modeling. Finally, Chapter 5 discusses the future of agent-based modeling research and where advances are likely to be made. The book concludes with a list of resources useful to agent-based modelers on the Web and in print.

Agent-based simulation has become increasingly popular as a modeling approach in the social sciences because it enables one to build models where individual entities and their interactions are directly represented. In comparison with variable-based approaches using structural equations, or system-based approaches using differential equations, agent-based simulation offers the possibility of modeling individual heterogeneity, representing explicitly agents' decision rules, and situating agents in a geographical or another type of space. It allows modelers to represent in a natural way multiple scales of analysis, the emergence of structures at the macro or societal level from individual action, and various kinds of adaptation and learning, none of which is easy to do with other modeling approaches.

1.1   Agent-Based Modeling

Formally, agent-based modeling is a *computational* method that enables a researcher to create, analyze, and *experiment* with *models* composed of *agents* that interact within an *environment*. Let us consider each of the italicized terms in this definition.

1.1.1   A Computational Method

First, agent-based modeling is a form of *computational social science*. That is, it involves building models that are computer programs. The idea of modeling is familiar in most of the social sciences: One creates some kind of simplified representation of "social reality" that serves to express as clearly as possible the way in which one believes that reality operates. For example, if one has a dependent variable and one or more independent variables, a regression equation serves as a model of the relationship between the variables. A network of nodes and edges can model a set of friendships. Even an ordinary language description of a relationship, such as that between the strength of protection of intellectual property rights and the degree of innovation in a country, can be considered a model, albeit a simple and rather unformalized one.

Computational models are formulated as computer programs in which there are some inputs (somewhat like independent variables) and some outputs (like dependent variables). The program itself represents the processes that are thought to exist in the social world (Macy & Willer, 2002). For example, we might have a theory about how friends influence the purchasing choices that consumers make. As we shall see, we can create a program in which there are individuals ("agents") that buy according to their preferences. The outcome is interesting because what one agent buys will influence the purchasing of a friend, and what the friend buys will influence the first agent. This kind of mutual reinforcement is relatively easy to model using agent-based modeling.

One of the advantages of computational modeling is that it forces one to be precise: Unlike theories and models expressed in natural language, a computer program has to be completely and exactly specified if it is to run. Another advantage is that it is often relatively easy to model theories about processes, for programs are all about making things within the computer change. If the idea of constructing computational models reminds you of computer games, especially the kind where the player has a virtual world to build, such as The Sims (http://thesims.ea.com/), that is no accident. Such games can be very close to computational modeling, although in order to make them fun, they often have fancier graphics and less social theory in them than do agent-based models.

1.1.2 Experiments

Whereas in physics and chemistry and some parts of biology, experimentation is the standard method of doing science, in most of the social sciences, conducting experiments is impossible or undesirable. An experiment consists of applying some *treatment* to an isolated system and observing what happens. The treated system is compared with another otherwise equivalent system that receives no treatment (the control). The great advantage of experiments is that they allow one to be sure that it is the treatment that is causing the observed effects, because it is only the treatment that differs between the treated and the control systems and the systems are isolated from other potential causes of change. However, with social systems, isolation is generally impossible, and treating one system while not treating the control is often ethically undesirable. Therefore, it is not surprising that experiments, despite the potential clarity of their results, are rarely used by social scientists.

A major advantage of agent-based modeling is that the difficulties in ensuring isolation of the human system and the ethical problems of experimentation are not present when one does experiments on virtual or computational systems. An experiment can be set up and repeated many times, using a range of parameters or allowing some factors to vary randomly. Of course, carrying out experiments with a computational model of some social phenomenon will yield interesting results only if the model behaves in the same way as the human system or, in other words, if the model is a good one, and one may not know whether that is the case. So, experimentation on models is not a panacea.

The idea of experimenting on models rather on the real system is not novel. For example, when architects put a model tower block in a wind tunnel to investigate its behavior in high winds, they are experimenting on the model for just the same reasons as social scientists might want to experiment on their models: The cost of experimenting on a real tower block is too high. Another reason for experimenting with models is that this may be the only way to obtain results. Deriving the behavior of a model analytically is usually best because it provides information about how the model will behave given a range of inputs, but often an analytical solution is not possible. In these cases, it is necessary to experiment with different inputs to see how the model behaves. The model is used to *simulate* the real world as it might be in a variety of circumstances.

1.1.3 Models

Computational social science is based on the idea of constructing models and then using them to understand the social world (Sawyer, 2004). Models have a long history in the social sciences—much longer than the

use of computers—but came to the fore when statistical methods began to be used to analyze large amounts of quantitative data in economics and demography. A model is intended to represent or simulate some real, existing phenomenon, and this is called the *target* of the model. The two main advantages of a model are that it succinctly expresses the relationships between features of the target, and it allows one to discover things about the target by investigating the model (Carley, 1999).

One of the earliest well-known social science models is the Phillips (1950) hydraulic model of the economy in which water flowing through interconnected glass pipes and vessels is used to represent the circulation of money. This model can still be admired at the Science Museum, London (http://en.wikipedia.org/wiki/MONIAC_Computer). The effect of changing parameters such as the interest rate can be investigated by changing the rate of flow of water through the pipes.

Models come in several flavors, and it is worth listing some of these to clarify the differences:

• *Scale models* are smaller versions of the target. Together with the reduction in size is a systematic reduction in the level of detail or complexity of the model. So, for example, a scale model of an airplane will be the same shape as its target, but probably would not show the electronic control systems or possibly even the engines of the real plane. Similarly, a scale model of a city will be much smaller than the real city and may model only two dimensions (the distances between buildings, but not the heights of buildings, for instance). When drawing conclusions about the target by studying the model, one needs to bear in mind that the results from the model will need to be scaled back up to the target's dimensions, and that it is possible that some of the features not modeled may affect the validity of the conclusions.

• An *ideal-type* model is one in which some characteristics of the target are exaggerated in order to simplify the model. For example, an idealized model of a stock market may assume that information flows from one trader to another instantaneously, and an idealized model of traffic may assume that drivers never get lost. The idealization has the effect of removing one or more complicating factors from the model, and if these have negligible effects on how the model works, the model will remain useful for drawing conclusions about the target.

• *Analogical models* are based on drawing an analogy between some better understood phenomenon and the target. The most famous example is the billiard ball model of atoms, but there are also social science examples such as the computer model of the mind (Boden, 1988) and the garbage can model of organizations (Cohen, March, & Olsen, 1972). Such models are

useful because well-established results from the analogy can be carried over and applied to the target, but of course the validity of these depends on the adequacy of the analogy.

These are not mutually exclusive categories; it is possible, and indeed common, for a model to be a scale model and an analogy (for example, the hydraulic model of the economy mentioned above is such a combination).

Some models fall into a fourth category that is somewhat different, but also commonly encountered in the social sciences; these are often called *mathematical* or *equation-based* models. Examples are the structural equation models of quantitative sociology and the macroeconomic models of neoclassical economics. These models specify relationships between variables, but unlike models in the other three categories, they do not imply any kind of analogy or resemblance between the model and the target. Usually, the success of a mathematical model is indicated by the degree to which some data fit the equation, but the form of the equation itself is of little interest or consequence. For example, the Cobb-Douglas "production function" is a mathematical model of how manufactured outputs are related to inputs (Cobb & Douglas, 1928):

$$Y = AL^{\alpha}K^{\beta},$$

where $Y =$ output, $L =$ labor input, $K =$ capital input, and A, α, and β are constants determined by technology. The form of this equation was derived from statistical evidence, not by theorizing about the behavior of firms. Although mathematical models have been very successful in some parts of the social sciences in clarifying the relationships between variables, they are often not very useful in helping to understand *why* one variable is related to another, or in other words, in expressing ideas about process and mechanism, where the other types of models are generally more helpful.

1.1.4   Agents

Agent-based models consist of agents that interact within an environment. Agents are either separate computer programs or, more commonly, distinct parts of a program that are used to represent social actors—individual people, organizations such as firms, or bodies such as nation-states. They are programmed to react to the computational environment in which they are located, where this environment is a model of the real environment in which the social actors operate.

As will be seen later, a crucial feature of agent-based models is that the agents can interact, that is, they can pass informational *messages* to each other and act on the basis of what they learn from these messages. The messages may represent spoken dialogue between people or more indirect means of information flow, such as the observation of another agent or

the detection of the effects of another agent's actions. The possibility of modeling such agent-to-agent interactions is the main way in which agent-based modeling differs from other types of computational models.

1.1.5   The Environment

The environment is the virtual world in which the agents act. It may be an entirely neutral medium with little or no effect on the agents, or in other models, the environment may be as carefully crafted as the agents themselves. Commonly, environments represent geographical spaces, for example, in models concerning residential segregation, where the environment simulates some of the physical features of a city, and in models of international relations, where the environment maps states and nations (Cederman, 1997). Models in which the environment represents a geographical space are called *spatially explicit*. In other models, the environment could be a space, but one that represents not geography but some other feature. For example, scientists can be modeled in "knowledge space" (Gilbert, Pyka, & Ahrweiler, 2001). In these spatial models, the agents have coordinates to indicate their location. Another option is to have no spatial representation at all but to link agents together into a network in which the only indication of an agent's relationship to other agents is the list of the agents to which it is connected by network links (Scott, 2000).

To make these definitions somewhat more concrete, in the next section we shall introduce some examples of agent-based models in terms of these concepts.

1.2   Some Examples

Agent-based models are of value in most branches of social science. The models that are briefly described in the rest of this section have been chosen to illustrate the diversity of the problem areas where they have been used productively.

1.2.1   Urban Models

In 1971, Thomas Schelling (1971, 1978; see also Sakoda, 1971) proposed a model to explain observed racial segregation in American cities. The model is a very abstract one as originally conceived, but it has been influential in recent work on understanding the persistence of segregation not only in the United States but also in other urban centers. The model is based on a regular square grid of cells representing an urban area on which agents, representing households, are placed at random. The agents are of two kinds (let us call them "reds" and "greens"). Each cell can hold only one

household agent at a time, and many cells are empty. At each time step, each household surveys its immediate neighbors (the eight cells surrounding it) and counts the fraction of households that are of the other color. If the fraction is greater than some constant threshold "tolerance" value, that is, there are more than a fixed proportion of reds surrounding a green, or greens surrounding a red, that household considers itself to be "unhappy" and decides to relocate. It does so by moving to some vacant cell on the grid.

At the next time step, the newly positioned household may tip the balance of "tolerance" of its neighbors, causing some of them to become "unhappy," and this can result in a cascade of relocations. For levels of the tolerance threshold at or above about 0.3, an initially random distribution of households segregates into patches of red and green, with households of each color clustering together (Figure 1.1). The clustering occurs even when households "tolerate" living adjacent to a majority of neighbors of the other color, which Schelling interpreted as indicating that even quite low degrees of racial prejudice could yield the strongly segregated patterns typical of U.S. cities in the 1970s.

The Schelling model has been influential for several reasons (Allen, 1997). First, the outcome—clusters of households of the same color—is surprising and not easily predictable just from considering the individual agents' movement rule. Second, the model is very simple and has only one parameter, the "tolerance" threshold. It is therefore easy to understand. Third, the emergent clustering behavior is rather robust. The same outcomes are obtained for a wide range of tolerance values, for a variety of movement rules (e.g., the household agent could select a new cell at
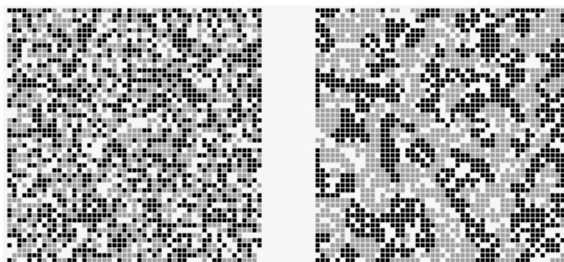
Figure 1.1   The Schelling Model at the Start (Left) and After Equilibrium Has Been Reached (Right), With a Uniform Tolerance of 0.3

SOURCE: Wilensky, U. (1998). NetLogo Segregation model. http://ccl.northwestern.edu/netlogo/models/Segregation. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

random, or use a utility function to select the most preferred cell, or take into account affordability if cells are priced, and so on) and for different rules about which neighbors to consider (e.g., those in the eight surrounding cells; the four cells to the north, east, south, and west; or a larger ring two or more cells away) (Gilbert, 2002). Fourth, the model immediately suggests how it could be tested with empirical data (Benenson, Omer, & Hatna, 2002; Bruch & Mare, 2006; Clark, 1991; Pancs & Vriend, 2004; Zhang, 2004), although in practice it has proved quite difficult to obtain reliable and extensive data on household location preferences to calculate ratings of ''unhappiness.'' The advantages of the Schelling model over others that had been previously proposed, which were based on equations relating migration flows and the relative values of residential properties (e.g., Fotheringham & O'Kelly, 1989), are that the number of parameters to be estimated is lower and that the model is simple to simulate and analyze. Current work has focused on making the model more concrete, replacing the abstract square grid with actual urban geographies and adding further factors, such as the affordability of the locations to which households want to move.

1.2.2 Opinion Dynamics

Another interesting group of models with potentially important implications is concerned with understanding the development of political opinions, for example, with explaining the spread of extremist opinions within a population. We shall review just one such study, although there are a number that explore the consequences of different assumptions and opinion transmission mechanisms (Deffuant, 2006; Deffuant, Amblard, & Weisbuch, 2002; Hegselmann & Krause, 2002; Lorenz, 2006; McKeown & Sheehy, 2006; Stauffer, Sousa, & Schulze, 2004). Deffuant et al. (2002) ask,

> How can opinions, which are initially considered as extreme and marginal, manage to become the norm in large parts of a population? Several examples in world history show that large communities can more or less suddenly switch globally to one extreme opinion, because of the influence of an initially small minority. Germany in the thirties is a particularly dramatic example of such a process. In the last decades, an initial minority of radical Islamists managed to convince large populations in Middle East countries. But one can think of less dramatic processes, like fashion for instance, where the behavior of minority groups, once considered as extremist, becomes the norm in a large part of the population (it is the case of some gay way of dressing for instance). On the other hand, one can also find many examples where a very strong bipolarization of the population takes place, for instance the Cold War period in Europe. In these cases, the whole population becomes extremist (choosing one side or the other).

In Deffuant et al.'s model, agents have an opinion (a real number between −1.0 and +1.0) and a degree of doubt about their opinion, called *uncertainty* (a positive real number). An agent's *opinion segment* is defined as the band centered on the agent's opinion, spreading to the right and left by the agent's value for uncertainty. Agents interact randomly. When they meet, one agent may influence the other if their opinion segments overlap. If the opinion segments do not overlap, the agents are assumed to be so different in their opinions that they have no chance of influencing each other. If an agent does influence another, the opinion of one agent (agent j) is affected by the opinion of another agent (agent i) by an amount proportional to the difference between their opinions, multiplied by the amount of overlap divided by agent i's uncertainty, minus one. The effect of this formula is that very uncertain agents influence other agents less than those that are certain (for full details, see Deffuant et al., 2002, Equations 1 to 6).

Every agent starts with an opinion taken from a uniform random distribution and with a common level of uncertainty, with the exception of a few extremists, those who have the most positive or negative opinions. The latter are given a low level of uncertainty, that is, the extremists are assumed to be rather certain about their extreme opinions. Under these conditions, extremism spreads, and eventually the simulation reaches a steady state with all agents joining the extremists at one or the other end of the opinion continuum. Restarting the simulation without the politically certain extremists, the population converges instead so that all agents share a middle view. Thus, the model shows that a few extremists with opinions that are not open to influence from other agents can have a dramatic effect on the opinions of the majority. This work has some implications for the development of terrorist movements, where a few extremists have been able to recruit support from substantial proportions of the wider population.

1.2.3 Consumer Behavior

Businesses are naturally keen to understand what influences their customers to buy their products. Although the intrinsic qualities of the product are usually important, so are the influence of friends and families, advertising, fashion, and a range of other "social" factors. To examine the often complex interactions between these, some researchers have started to use agent-based models in which the agents represent consumers. Among the first to report such a model were Janssen and Jager (1999), who explored the processes leading to "lock-in" in consumer markets. Lock-in occurs when one among several competing products achieves dominance so that it is difficult for individual consumers to switch to rival products. Commonly cited examples are VHS videotapes (dominating Betamax), the QWERTY

keyboard (dominating other arrangements of the keys), Microsoft operating systems (dominating Apple and Linux), and so on. Janssen and Jager focus on the behavioral processes that lead to lock-in, and therefore they give their agents, which they call "consumats," decision rules that are psychologically plausible and carefully justified in terms of behavioral theories of, for example, social comparison and imitation.

Another example of modeling consumers is a study by Izquierdo and Izquierdo (2006) in which the authors consider markets such as the second-hand car market, where there is quality variability (different quality for different items) and quality uncertainty (it is difficult to know the quality of an item before buying it and using it). The study explores how quality variability can damage a market and affect consumer confidence. There are two agent roles: buyer and seller. Sellers sell products by calculating the minimum price they will accept, and buyers buy products by offering a price based on the expected quality of the product. The expected quality is based on experience, either just of the agent or accumulated from the agent's peers over its social network. There are a finite number of products in the system, buyers and sellers perform one deal per round, and the market is cleared every round as these deals are done. The social network is created by connecting pairs of agents at random, with a parameter used to adjust the number of connections, from completely connected to completely unconnected.

The authors found that without a social network, consumer confidence fell to the point where the market was no longer viable, whereas with a social network, the aggregation of the agent's own experience and the more positive collective experience of others (which is not so volatile) helped to maintain the market's stability. This shows how social information can aggregate group experience to a more accurate level and so reduce the importance of a single individual's bad experiences.

1.2.4 Industrial Networks

Most economic theory ignores the significance of links between firms, but there are many examples of industrial sectors where networks are of obvious importance. A well-known instance is the "industrial districts" of northern Italy, such as the textile production district, Prato. Industrial districts are characterized by large numbers of small firms clustered together in a small region, all manufacturing the same type of product, with strong, but varying, links between them. The links may be those of a supplier/customer, a collaboration to share techniques, a financial link, or just a friendship or familial relationship (Albino, Carbonara, & Giannoccaro, 2003; Boero, Castellani, & Squazzoni, 2004; Borrelli, Ponsiglione, Iandoli, & Zollo, 2005; Brenner, 2001; Fioretti, 2001; Squazzoni & Boero, 2002).

Another example is the "innovation networks" that are pervasive in knowledge-intensive sectors such as biotechnology and information technology. The firms in these sectors are not always geographically clustered (although there tend to be concentrations in certain locations), but they do have strong networking relationships with other, similar firms, sharing knowledge, skills, and technology with them.

For example, Gilbert et al. (2001) developed a model of innovation networks in which agents have "kenes" that symbolize their stock of knowledge and expertise. The kenes are used to develop new products that are marketed to other firms in the model. However, a product can be produced only if its components are available to be purchased from other firms, and if some firm wants to buy it. Thus, at one level, the model is one of an industrial market with firms trading between each other. In addition, firms can improve their kenes either through internal research and development or through incorporating knowledge obtained from other firms by collaborative arrangements. The improved knowledge can be used to produce products that may sell better, or require fewer or more cheaply available components. At this level, the model resembles a population that can learn through a type of natural selection (see Section 5.2.2) in which firms that cannot find a customer cease trading, whereas the fittest firms survive, collaborate with other firms, and produce spin-offs that incorporate the best aspects of their "parents." For another example, see Pajares, Hernández-Iglesias, and López-Paredes (2004).

1.2.5 Supply Chain Management

Manufacturers normally buy components from other organizations and sell their products to distributors, who then sell to retailers. Eventually, the product reaches the user, who may not realize the complex interorganizational relationships that have had to be coordinated to deliver the product. Maximizing the efficiency of the supply chain linking businesses is increasingly important and increasingly difficult as products involve more parts, drawn more widely from around the world, and as managers attempt to reduce inventory and increase the availability of goods. Modeling supply chains is a good way of studying order fulfillment processes and investigating the effectiveness of management policies, and multiagent models are increasingly being used for this purpose.

A multiagent model fits well with the task of simulating supply chains because the businesses involved can be modeled as agents, each with its own inventory rules. It is also easy to model the flow of products down the chain and the flow of information, such as order volumes and lead times, from one organization to another. This was the approach taken by Strader,

Lin, and Shaw (1998), who described a model they built to study the impact of information sharing in divergent assembly supply chains. Divergent assembly supply chains are typical of the electronics and computer industries and are those in which a small number of suppliers provide materials and subcomponents (e.g., electronic devices) that are used to assemble a range of generic products (e.g., hard disk drives) that are then used to build customized products at distribution sites (e.g., personal computers). Strader et al. compared three order fulfillment policies: make-to-order, when production is triggered by customer orders; assembly-to-order, when components are built and held in stock, and only the final assembly is triggered by an order; and make-to-stock, when production is driven by the stock level falling below a threshold. They also experimented with different amounts of information sharing between organizations, and found that in the divergent assembly supply chains that they modeled, an assembly-to-order strategy, coupled with the sharing of both supply and demand information between organizations along the supply chain, was the most efficient. They also pointed out that their results reinforce the general point that information can substitute for inventory: If one has good information, uncertainty about demand can be reduced, and the required inventory level to satisfy orders can also be reduced as a consequence.

1.2.6   Electricity Markets

In many developed countries, in recent years, the electricity supply has been privatized. It is now common for there to be two or three electricity utilities that sell power to a number of distributors that in turn sell the electricity to commercial and domestic users. The change from a monopoly state-owned or state-regulated supplier to one in which there are a number of supply firms bidding into a market has inspired a range of agent-based models that aim to anticipate the effect of market regulations; changes to the number and type of suppliers and purchasers; and policy changes intended, for example, to reduce the chance of blackouts or decrease the environmental impact of generation (Bagnall & Smith, 2005; Batten & Grozev, 2006; Bunn & Oliveira, 2001; Koesrindartoto, Sun, & Tesfatsion, 2005; North, 2001).

In these models, the agents are the supply companies that make offers to the simulated market to provide a certain quantity of electricity at a certain price for a period, such as a day or an hour. This is also how the real electricity markets work: Companies make offers to supply and the best offer is accepted (different markets have different rules about what is meant by the "best" offer). Usually, the demand varies continuously, so supply companies have a difficult job setting a price for the electricity that maximizes

their profit. A further complication is that the cost of generation can be very nonlinear: Matching peak demand may mean starting up a power station that is used for only a few hours.

By running the simulation, one can study the conditions under which the market price comes down to near the marginal cost of generation; the effect of mergers that reduce the number of supply companies; and the consequences of having different types of market "design," such as allowing futures trading. Most of the current models allow the agents to "learn" trading strategies using a technique known as *reinforcement learning* (see Section 5.2.1). A supply agent starts by making a bid using a pricing strategy selected at random from a set common to all the suppliers. If the bid is accepted and profitable, the value of this strategy is reinforced and the probability of using it again in similar circumstances is increased, or if it is unsuccessful, the chance of using it again is decreased (Roth & Erev, 1995).

1.2.7 Participative and Companion Modeling

Agent-based models have been used with success in rural areas in Third World countries to help with the management of scarce natural resources such as water for irrigation. This surprising use of agent-based models is due to their fit with "participative" (or participatory) research methods. As well as being used for research, multiagent models have been used as a support for negotiation and decision making and for training with, for example, the farmers in Senegal (D'Aquino, Le Page, Bousquet, & Bah, 2003), foresters and farmers in the Central Massif in France (Etienne, 2003; Etienne, Le Page, & Cohen, 2003), and the inhabitants of an atoll in Kiribati in the South Pacific (Dray et al., 2006).

The approach, also called *companion modeling* (Barreteau, 2003; Barreteau, Bousquet, & Attonaty, 2001; Barreteau, Le Page, & D'Aquino, 2003) involves building a multiagent system in close association with informants selected from the people on the ground. As a preliminary, the informants may be interviewed about their understanding of the situation, and they then engage in a role-playing game. Eventually, when sufficient knowledge has been gained, a computer model is created and used with the participants as a training aid or as a negotiation support, allowing the answering of "what-if" questions about possible decisions.

For example, Barreteau et al. (2001) describe the use of participative modeling in order to understand why an irrigation scheme in the Senegal River Valley had produced disappointing results. They developed both a role-playing game (RPG) and a multiagent system called SHADOC to represent the interactions between the various stakeholders involved in

making decisions about the allocation of water to arable plots in the irrigated area. In this instance, the multiagent model was developed first and then its main elements converted to an RPG (in which the players were the equivalent of the agents in the multiagent model), partly to validate the agent-based model, and partly because the RPG is easier to use in a rural environment. The authors sum up the value of this approach as "enhancing discussion among game participants" and enabling "the problems encountered in the field and known by each individual separately [to be] turned into common knowledge."

1.3 The Features of Agent-Based Modeling

These examples, chosen to illustrate the spectrum of agent-based modeling now being undertaken, also provide examples of some characteristic features of agent-based modeling (Fagiolo, Windrum, & Moneta, 2006).

1.3.1 Ontological Correspondence

There can be a direct correspondence between the computational agents in the model and real-world actors, which makes it easier to design the model and interpret its outcome than would be the case with, for example, an equation-based model. For instance, a model of a commercial organization can include agents representing the employees, the customers, the suppliers, and any other significant actors. In each case, the model might include an agent standing for the whole class (e.g., "employees"), or it might have a separate agent for each employee, depending on how important the differences between employees are. The models of electricity markets described above have agents for each of the main players in the market.

1.3.2 Heterogeneous Agents

Generally speaking, theories in economics and organization science make the simplifying assumption that all actors are identical or similar in most important respects. They deal, for example, with the "typical firm," or the economically rational decision maker. Actors may differ in their preferences, but it is unusual to have agents that follow different rules of behavior, and when this is allowed, there may be only a small number of sets of such actors, each with its own behavior. This is for the good reason that unless agents are homogeneous, analytical solutions are very difficult or impossible to find. A computational model avoids this limitation: Each

agent can operate according to its own preferences or even according to its own rules of action. An example is found in the models of supply chains, in which each business can have its own strategy for controlling inventory.

1.3.3   Representation of the Environment

It is possible to represent the ''environment'' in which actors are acting directly in an agent-based model. This may include physical aspects (e.g., physical barriers and geographical hurdles that agents have to overcome), the effects of other agents in the surrounding locality, and the influence of factors such as crowding and resource depletion. For example, Gimblett (2002) and colleagues have modeled the movement of backpackers in the Sierra Nevada Mountains in California to examine the effect of management policies in helping to maintain this area of wilderness. The agents simulated trekking in a landscape linked to a *geographical information system* that modeled the topology of the area. The environment also plays an important role in the models of industrial districts mentioned in the previous section.

1.3.4   Agent Interactions

An important benefit of agent-based modeling is that interactions between agents can be simulated. At the simplest, these interactions can consist of the transfer of data from one agent to another, typically another agent located close by in the simulated environment. Where appropriate, the interaction can be much more complicated, involving the passing of messages composed in some language, with one agent constructing an ''utterance'' and the other interpreting it (and not necessarily deriving the same meaning from the utterance as the speaker intended). The opinion dynamics models (Section 1.2.2) are a good example of the importance of interactions in agent-based models.

1.3.5   Bounded Rationality

Many models implicitly assume that the individuals whom they model are rational, that is, that they act according to some reasonable set of rules to optimize their utility or welfare. (The alternative is to model agents as acting randomly or irrationally, in a way that will not optimize their welfare. Both have a place in some models.) Some economists, especially those using rational choice theory, have been accused of assuming that individuals are ''hyperrational,'' that is, that people engage in long chains of complex reasoning in order to select optimal courses of action, or even that people are capable of following chains of logic that extend indefinitely. Herbert Simon (1957), among others, criticized this as unrealistic and

proposed that people should be modeled as *boundedly rational*, that is, as limited in their cognitive abilities and thus in the degree to which they are able to optimize their utility (Kahneman, 2003). Agent-based modeling makes it easy to create boundedly rational agents. In fact, the challenge is usually not to limit the rationality of agents but to extend their intelligence to the point where they could make decisions of the same sophistication as is commonplace among people.[1] Models of stock markets and the segregation model introduced in Section 1.2.1 are examples where agents have been designed with strictly limited rationality.

1.3.6   Learning

Agent-based models are able to simulate learning at both the individual and population levels. For example, the firms in the model of innovation networks described above (Section 1.2.4) are able to learn how to produce a more salable and more profitable product, and the sector as a whole (that is, all the agents in the model as a collection) learn over time which products will form a compatible set so that products from one firm will provide the components bought by another firm. Learning can be modeling in any or all of three ways: as individual learning in which agents learn from their own experience; as evolutionary learning, in which the population of agents learns because some agents "die" and are replaced by better agents, leading to improvements in the population average; and social learning, in which some agents imitate or are taught by other agents, leading to the sharing of experience gathered individually but distributed over the whole population (Gilbert et al., 2006). The model of innovation networks summarized above is an example of a model incorporating learning: The individual innovating firms learn how to make better products, and because poorly performing firms become bankrupt to be replaced by better start-ups, the sector as a whole can learn to improve its performance.

Some techniques for designing models that incorporate learning will be discussed in Section 5.2.

1.4   Other Related Modeling Approaches

The previous section has reviewed some areas where agent-based models have been useful. However, agent-based models are not appropriate for every modeling task. Before starting a new project, it is worth considering the alternatives. This section introduces two styles of modeling used in the social sciences that stand comparison with agent-based modeling: microsimulation and system dynamics.

1.4.1 Microsimulation

Microsimulation starts with a large database describing a sample of individuals, households, or organizations and then uses rules to update the sample members as though time was advancing. For example, the database might be derived from a representative national survey of households and include data on variables such as household members' age, sex, education level, income, employment, and pension arrangements. These data would relate to the specific time period when the survey was carried out. Microsimulation allows one to ask what the sample would be like in the future. For example, one might want to know how many in the sample would be retired in 5 years' time and how this would affect the distribution of income. If we have some rules about the likely changes in individual circumstances during the course of a year, these rules can be applied to every person in the sample to find what might have changed by the end of the first year after the survey. Then the same rules can be reapplied to yield the state of the sample after 2 years, and so on. After this aging process has been carried out, aggregate statistics can be calculated for the sample as a whole (for example, the mean and variance of the distribution of income, which can be compared with the distribution at the time of the survey) and inferences made about what changes are to be expected in the population from which the sample was drawn (Gupta & Kapur, 2000; Harding, 1996; Mitton, Sutherland, & Weeks, 2000; Orcutt, Merz, & Quinke, 1986; Redmond, Sutherland, & Wilson, 1998).

Microsimulation has been used to assess the distributional implications of changes in social security, personal tax, and pensions. For example, it can be helpful in evaluating the effects of changing the income threshold below which state benefits become payable (Brown & Harding, 2002). Experimental prototypes have also been developed in which there are several databases, describing not only individuals but also firms, and in which the aging process is affected not only by individual characteristics but also by macroeconomic variables such as inflation and the growth in gross domestic product (GDP); see the bibliography at http://www.microsimulation.org/.

An advantage of microsimulation models is that they start not from some hypothetical or randomly created set of agents but from a real sample, as described by a sample survey. Hence, it is relatively easy, in comparison with agent-based models, to read back from the results of the microsimulation to make predictions about the future state of a real population. There are two main disadvantages. First, the aging process requires very detailed transition matrices that specify the probability that an agent currently in some state will change to some other state in the following year. For example, one

needs to know the probability that someone currently in employment will become unemployed 1 year later. Moreover, because this transition probability will almost certainly differ between men and women, women with and without children, young and old people, and so on, one needs not a single probability value but a matrix of conditional probabilities, one for each combination of individual circumstances. Obtaining reliable estimates of such transition matrices can be very difficult, requiring estimation from very large amounts of data. Second, each agent is aged individually and treated as though it is isolated in the world. Microsimulation does not allow for any interaction between agents and typically has no notion of space or geography. So, for instance, it is hard to take account of the finding that the chances of getting a job if one is unemployed are lower if one lives in an area where the unemployment rate is high.

1.4.2 System Dynamics

In the system dynamics approach to modeling, one creates a model that expresses the temporal cause-and-effect relationships between variables, but agents are not represented directly. One of the earliest and best known examples is Forrester's model of the world, which was used to make predictions about future population levels, growing pollution, and rates of consumption of natural resources (Forrester, 1971). System dynamics, as its name implies, models systems of interacting variables and is able to handle direct causal links, such as a growth in population leading to increased depletion of resources, and feedback loops, as when population growth depends on the food supply, but food supply depends on the level of the population (Sterman, 2000).

It is often convenient to represent a system dynamics model with a diagram in which arrows represent the causal links between variables. Figure 1.2 shows a typical, although simple, model of an ecosystem in which sheep breed in proportion to their population, wolves eat the sheep, but if there are too few sheep, the wolves starve. The rectangular boxes represent the stocks of sheep and wolves, the tap-like symbols are flows into and out of the stocks, and the diamond shapes are variables that control the rate of flow. The population of sheep increases as sheep are born, and the rate at which this happens is determined by the constant sheep-birth-rate. The diagram shows that sheep die at a rate that is a function of the number of sheep living (the curved arrow from the stock of sheep to the flow control labeled sheep-deaths), the probability that a wolf will catch a sheep (the arrow from the predation-rate variable), and the number of wolves (the arrow from the stock of wolves). Although this illustrative model is concerned with somewhat imaginary wolves and sheep, similar models can be constructed for
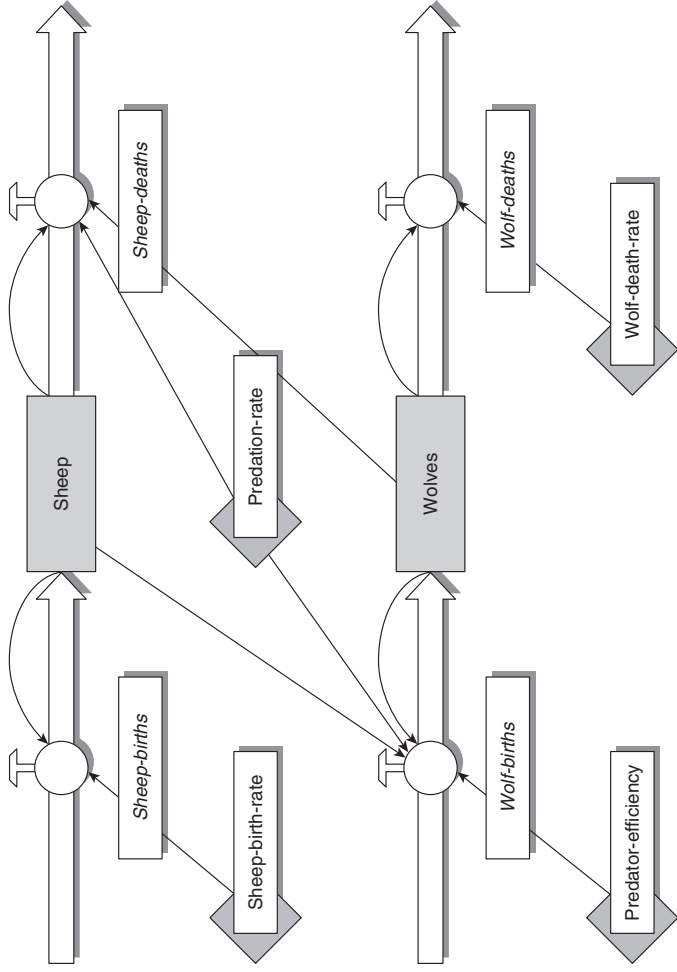
Figure 1.2   A System Dynamics Model of a Simple Ecosystem, With Wolves Eating Sheep According to the Lotka-Volterra Equations

topics of sociological interest, such as the number of illegal drug users and enforcement agents, and public health epidemics (Homer & Hirsch, 2006; Jacobsen, Bronson, & Operations Research Society of America, 1985).

System dynamics is based on the evaluation of sets of simultaneous differential or difference equations, each of which calculates the value of a variable at the next time step given the values of other, causal variables at the current time step. Software such as Stella (http://www.iseesystems .com/) and NetLogo (http://ccl.northwestern.edu/netlogo/, described in more detail in Chapter 3) can help with drawing the diagrams and also execute the simulation by computing these equations.

In comparison with agent-based modeling, the system dynamics approach deals with an aggregate, rather than with individual agents. For example, in the wolves and sheep model, the simulation will compute the total population of sheep at each time step, but each individual sheep is not represented. This makes it hard to model heterogeneity among the agents; although one could, in principle, have a distinct stock for each different type of agent (e.g., a stock of white sheep, a stock of black sheep, a stock of mottled sheep, and so on), in practice this becomes extremely cumbersome with more than a few different types. It is also hard to represent agent behaviors that depend on the agent's past experience, memory, or learning in a system dynamics model. On the other hand, because they deal with aggregates, the system dynamics approach is good for topics where there are large populations of behaviorally similar agents. Thus, system dynamics was an appropriate method for Forrester's models of the global economy because individual action was unimportant and the focus was on the state of the world as a whole.

Note

1. Nevertheless, it has been found in several contexts, such as in modeling stock markets, that the aggregate behavior of agents with very little rationality (or ''zero intelligence'') matches the observed behavior at the macro level surprisingly well (Chan, LeBaron, Lo, & Poggio, 1999; Farmer, Patelli, & Zovko, 2005).