

Chapter 1

Simulation and social science

Using computer simulation in the social sciences is a rather new idea – although the first examples date from the 1960s, simulation only began to be used widely in the 1990s – but one that has enormous potential. This is because simulation is an excellent way of modelling and understanding social processes.

This book has been written for social scientists interested in building simulations. All research should be theoretically informed, methodologically sophisticated and creative. These qualities are especially necessary when doing simulations because the field is only about 20 years old, so there are no well-established traditions to rely on, and there are a wide variety of approaches to simulation from which to choose. One additional skill needed by the researcher wanting to use simulation is some facility in using computers (all simulations nowadays are run on computers). It helps to know how to write simple programs, although the first half of this book does not demand any programming knowledge at all, and the second half needs only a beginner's level of skill.

Simulation introduces the possibility of a new way of thinking about social and economic processes, based on ideas about the emergence of complex behaviour from relatively simple activities (Simon 1996). These ideas, which are gaining currency not only in the social sciences but also in physics and biology, go under the name of complexity theory (see, by way of introduction, Waldrop 1992). However, we do not consider the theoretical implications of simulation in any depth in this book although there are frequent references to the theoretical foundations. Instead, the book

2 *Simulation and social science*

focuses on the practical and methodological issues of how to do simulation, covering matters such as the approach to adopt, the stages one should expect to go through and the traps and difficulties to avoid. In this first chapter, we discuss the types of problem and purposes for which simulation is best suited, present a few examples of simulation as it is used in social science and develop a classification of the types of simulation that will be described later in the book.

What is simulation?

Simulation is a particular type of modelling. Building a model is a well-recognized way of understanding the world: something we do all the time, but which science and social science has refined and formalized. A model is a simplification – smaller, less detailed, less complex, or all of these together – of some other structure or system. A model aeroplane is recognizably an aeroplane, even if it is much smaller than a real aeroplane and has none of its complex control systems. More relevant to social science are statistical models which are used to predict the values of dependent variables. Chapter 2 describes the idea of modelling and the differences between statistical models and simulation models in detail.

Like statistical models, simulations have ‘inputs’ entered by the researcher and ‘outputs’ which are observed as the simulation runs. Often, the inputs are the attributes needed to make the model match up with some specific social setting and the outputs are the behaviours of the model through time. An example – based loosely on the work of Todd (1997) – may make this clearer. Suppose that we are interested in how people choose a marriage partner. Do you (perhaps, did you?) keep looking and dating until you found someone who meets all your romantic ideals, or do you stop as soon as you find someone ‘good enough’? Do people use a sufficiently rigorous search procedure or, as Frey and Eichenberger (1996) suggest, should they search longer, possibly reducing the divorce rate as a result?

Asking people about their searching behaviour is unlikely to be very helpful: they may not be following any conscious strategy and may not reveal it even if they do have one. Instead, we might set up a model (in this case, a computer program) which embodies some plausible assumptions and see what happens, comparing the behaviour of the program with the observed patterns of searching for a partner.

This example is typical in several ways of how simulations can be used.

- When we have a theory of how people choose mates, we can express it in the form of a procedure and ultimately in the form of a computer program. The program will be much more precise than the textual form of the procedure and is therefore helpful in refining one's theory. Simulation can thus be used as a *method of theory development*.
- Once the theory is formalized into a program and we have made some assumptions, the program can be run and the behaviour of the simulation observed. Let us assume that we have a population of simulated potential suitors, each with a 'suitability' score chosen at random. Suppose further that the simulated person looking for a partner (the 'agent') can date potential suitors, selected at random, one after the other. At the end of every date, the agent has to choose whether to settle down with that person or break up and go on to date another suitor. This decision has to be made without knowing about the suitability of others whom the agent has not yet met and without the possibility of ever going back to a rejected suitor.

Figure 1.1: The mate searching game

55	116	149	217	117	81	308	193	78	239
85	15	294	110	219	275	151	310	191	75
110	21	23	132	259	264	194	59	273	239
166	254	136	100	172	30	172	288	128	276
94	169	38	208	145	73	147	13	256	280
312	187	158	124	203	264	142	241	192	54
27	216	316	301	0	183	250	112	30	19
189	273	29	111	259	97	256	249	130	13
53	253	15	273	148	6	97	295	22	238
98	141	88	60	279	211	35	160	304	10

Instructions: Cover up the rows of numbers with a piece of paper and gradually reveal them, starting from the top left corner, working downwards row by row. Wait for a couple of seconds between revealing each new number (this represents the time you spend dating your potential partner!). Decide for yourself when you want to stop. The last number you revealed is the suitability score of the person you would 'marry'. What is the best strategy to maximize the score, while minimizing the number of partners you have to date? (Try not to cheat by looking before you start either at the overall distribution of numbers or how many numbers there are in all.)

- To get the feel for this, cover up the array of numbers in Figure 1.1

4 *Simulation and social science*

with a piece of paper and then, moving the paper from left to right, row by row, gradually reveal more and more numbers. These numbers represent the suitability of successive dates. Stop whenever you feel that you have seen enough scores, remembering that if you spend too long dating you will have missed many years of married bliss!

- The suitability score of the selected partner is the ‘output’ for one run of the simulation. We can repeat the simulation many times. Since in the simulation the suitors are given random scores and the agent picks them in random order, the result may be different for each run, but the average score over a large number of runs will be useful. We can thus see that simulation allows the researcher to conduct experiments in a way that is normally impossible in social science.

Todd (1997) explores a number of possible strategies, including those that have been proved analytically to be optimal in terms of finding the best partner, but which require unrealistic amounts of search, and some other strategies that are much simpler and have better results when one takes into account that search is expensive in time and effort. He also begins to investigate the implications for search strategies when there is a possibility that you might want to settle down with a partner but the partner may still be wanting to continue to search for someone else. Even in this much more complex situation, simple strategies seem to suffice.

The uses of simulation

The example of strategies for searching for a partner illustrates one purpose of simulation: to obtain a better *understanding* of some features of the social world. We can observe dating behaviour going on all the time but the underlying strategies that people use are hard to discover directly, so simulation can be useful. However, this is not the only value of simulation (Axelrod 1997a).

Another classic use of simulation is for *prediction*. If we can develop a model that faithfully reproduces the dynamics of some behaviour, we can then simulate the passing of time and thus use the model to ‘look into the future’. A relatively well-known example is the use of simulation in demographic research, where one wants to know how the size and age structure of a country’s population will change over the next few years or decades. A model incorporating age-specific fertility and mortality rates can be used to predict population changes a decade into the future with fair accuracy.

Another example is the use of simulations for business forecasting.

A third use of simulation is to develop new tools to *substitute* for human capabilities. For example, expert systems (Hayes-Roth *et al.* 1983) have been constructed to simulate the expertise of professionals such as geologists, chemists and doctors. These systems can be used by non-experts to carry out diagnoses which would otherwise require human experts.

These and other simulations have been used for *training*. For example, an expert system that classifies rocks according to the likelihood that valuable minerals will be found in them can be used to train novice geologists. Flight simulators can be used to train pilots. And simulations of national economies can be used to train economists (see, for example, the simulation of the British economy available on the World Wide Web at <http://www.bized.ac.uk/virtual/economy/>).

A related use of simulation is for *entertainment*. Flight simulators are used not only for training pilots, but also for fun on home personal computers. Some simulations sold as games are very close to being social simulations of the type described in this book. For example, in Maxis' SimCity, the user plays the part of a city mayor and can alter property tax rates and other parameters to build a simulated city.

The major reason for social scientists becoming increasingly interested in computer simulation, however, is its potential to assist in *discovery* and *formalization*. Social scientists can build very simple models that focus on some small aspect of the social world and discover the consequences of their theories in the 'artificial society' that they have built. In order to do this, they need to take theories that have conventionally been expressed in textual form and formalize them into a specification which can be programmed into a computer. The process of formalization, which involves being precise about what the theory means and making sure that it is complete and coherent, is a very valuable discipline in its own right. In this respect, computer simulation has a similar role in the social sciences to that of mathematics in the physical sciences.

Mathematics has sometimes been used as a means of formalization in the social sciences, but has never become widespread except, perhaps, in some parts of econometrics. There are several reasons why simulation is more appropriate for formalizing social science theories than mathematics (Taber and Timpone 1996). First, programming languages are more expressive and less abstract than most mathematical techniques, at least those accessible to non-specialists. Second, programs deal more easily with parallel processes and processes without a well-defined order of actions than systems of mathematical equations. Third, programs are (or can easily be made to

be) modular, so that major changes can be made in one part without the need to change other parts of the program. Mathematical systems often lack this modularity. Finally, it is easy to build simulation systems that include heterogeneous agents – for example, to simulate people with different perspectives on their social worlds, different stocks of knowledge, different capabilities and so on – while this is usually relatively difficult using mathematics. Examples in which we compare mathematical and simulation treatments of a problem can be found in Chapters 3 and 6.

It is the use of simulation for experiment, proof and discovery in the social sciences which is the major concern of this book.


The history of social science simulation

Computer simulation in the social sciences had a difficult birth (Troitzsch 1997). Although there are isolated earlier examples, the first developments in computer simulation in the social sciences coincided with the first use of computers in university research in the early 1960s (Figure 1.2). They mainly consisted of discrete event simulations or simulations based on system dynamics. The former approach, described in Chapter 5, models the passage of units through queues and processes in order to predict typical throughput – for example, the waiting time of customers in a queue or the time a city’s police cars take to reach an emergency (Kolesar and Walker 1975). The system dynamics approach makes use of large systems of difference equations to plot the trajectories of variables over time – for example, the Club of Rome studies of the future of the world economy (Meadows *et al.* 1974; 1992). System dynamics and world models are described further in Chapter 3. The Club of Rome simulations that predicted global environmental catastrophe made a major impact but also gave simulation an undeservedly poor reputation as it became clear that the results depended very heavily on the specific quantitative assumptions made about the model’s parameters. Many of these assumptions were backed by rather little evidence.

This early work also suffered in another respect: it was focused on prediction, while social scientists tend to be more concerned with understanding and explanation. This is due to scepticism about the possibility of making social predictions, based on both the inherent difficulty of doing so and also the possibility, peculiar to social and economic forecasting, that the forecast itself will affect the outcome.

One approach that did blossom for some years became known as ‘Simulmatics’ (Sola Pool and Abelson 1962). The Simulmatics project was

Figure 1.2: The development of contemporary approaches to simulation in the social sciences (after Troitzsch 1997)



Legend: grey shaded area: equation based models; white area: object, event or agent based models; 'sCA' means cellular automata used for social science simulation; the other names of tools are explained in the respective chapters

originally designed to advise John F. Kennedy's presidential campaign. It tried to predict the reactions of voters to the measures taken by Kennedy and his campaign team, and was also used to understand voters' behaviours in the referendum campaigns about the fluoridation of drinking water, which were frequent in the United States in the early 1960s (Abelson and Bernstein 1963). The latter project was very similar to present-day multi-agent simulation (the term was only coined some 20 years later). Fifty simulated individuals were exposed to information about the topic of the referendum from several different channels and additionally exchanged information among themselves. How much information they absorbed and how much of this led to attitude change depended on their simulated communication habits,

but above all on their current attitudes (for example, the more extreme their current attitude was, the less susceptible they were to new information). The whole model included 51 rules of this kind, of which 22 refer to communication channels and other sources of information and 27 concern information exchange among the simulated individuals (the remaining two determine the ballot cast at the end of the simulated campaign).

Another approach that has thrived for more than two decades, impelled by policy concerns, is rather misleadingly called ‘microsimulation’ (Orcutt *et al.* 1986; Harding 1990). This is a very specific technique, yet until recently was the only form of simulation that had widespread recognition within the social sciences. Microsimulation, described in Chapter 4, is based on a large random sample of a population of individuals, households or firms. Each unit is ‘aged’ using a set of transition probabilities, which determine the chance that the unit will undergo some change during the passage of a year (for example, the probability that a woman within a certain age range will give birth to a child). After every unit has been aged by one year, the process is repeated for the next year, thus advancing the sample through simulated time. Aggregate statistics can be calculated and used as estimates of the future characteristics of the population. Microsimulation has become well established in some parts of the world (particularly in Germany, Australia and Canada) where its results have been influential in devising policies for state pensions, graduate taxes and so on.

Microsimulation has some characteristics that are instructive when compared with other approaches to simulation. First, it has no pretensions to explanation: it is simply a means of predicting future fiscal distributions. Second, it treats each unit (person, household or firm) individually: there is no attempt to model interactions between units. Third, the motivations or intentions of the units are disregarded: each unit develops from year to year only in response to the throw of the dice represented by a random number generator.

Apart from microsimulation, little was heard about simulation during the 1980s, in marked contrast to the situation in the natural sciences where simulation is now a basic methodological tool. However, in the early 1990s the situation changed radically, mainly as a result of the development of multi-agent models which offered the promise of simulating autonomous individuals and the interactions between them. These opportunities came from techniques imported from the study of nonlinear dynamics and from artificial intelligence research.

Physicists and mathematicians had been trying to understand the properties of large aggregates of matter and had devised models called cellular

automata to do so. These models have been applied to explain the properties of magnetic materials, turbulent flow in liquids, crystal growth, soil erosion and in many other areas of science (Toffoli and Margolus 1987). In all these cases, the properties of the material as a whole can be modelled by simulating the interactions between the component units (molecules, soil particles or whatever). Cellular automata consist of a large grid of cells in a regular arrangement. Each cell can be in one of a small number of states and changes between these states occur according to rules which depend only on the states of the cell's immediate neighbours. Cellular automata form a useful framework for some models of social interaction, for example the spread of gossip between people and the formation of ethnically segregated neighbourhoods. They are described in more detail in Chapter 7.

Another approach that has been influenced by ideas from physics is multilevel modelling (Chapter 6) which has taken its inspiration from the theory of synergetics, originally developed for application to condensed matter physics.

Artificial intelligence is an area of computer science concerned with the development of simulations of human intelligence and with building tools which exhibit some of the characteristics of intelligent behaviour. Until recently, artificial intelligence had only been involved with modelling individual cognition, but in the 1980s there was increasing interest in distributed artificial intelligence, a field which examines the properties of interacting artificial intelligence programs. With the growth of the Internet and the World Wide Web, many artificial intelligence researchers became interested in software 'agents', programs that can receive or collect information from other computers, assess it in the light of their past experience and decide what action to take (Doran 1997a). Both distributed artificial intelligence and the agent technology strands of research developed models which, because they involved interacting autonomous agents, could be applied to the simulation of human societies. Distributed artificial intelligence and multi-agent systems are discussed in Chapter 8. Chapter 9 considers strategies and techniques for designing multi-agent models.

Artificial intelligence researchers have also devoted a great deal of attention over the last decade to techniques of 'machine learning' (Michalski *et al.* 1983), which allow computer programs to increase their knowledge and their procedural skills by learning from experience. Models with the ability to learn are very useful both for simulating the cognitive processes of individuals and for modelling whole societies which adapt over time to new circumstances. Chapter 10 discusses some approaches to modelling learning and their application to social simulation.

Simulating human societies

This brief history of social science simulation research indicates that several of the approaches used in contemporary social simulation were originally developed in fields such as physics and artificial intelligence. Although the subject matter of the social sciences differs from that of the natural sciences and different issues are important in modelling societies compared with modelling, for example, aggregates of physical particles, these science and engineering techniques are proving to be very useful. On the other hand, some issues are specific to the social sciences and the relevance of computer simulation to understanding human societies therefore needs to be considered carefully.

One of the themes of social simulation research is that even when agents are programmed with very simple rules, the behaviour of the agents considered together can turn out to be extremely complex. Conventional statistical methods for analyzing social systems are almost all based on the assumption of a linear relationship between variables. That is, the effect on the dependent variable is proportional to a sum of a set of independent variables. But this is a very restrictive assumption. A new interdisciplinary field called complexity theory (Waldrop 1992; Kauffman 1995; Sole and Goodwin 2002) is developing general results about nonlinear systems. An example: consider pouring a steady stream of sand out of a pipe so that it mounts up into a pyramid. As you pour on more sand, there will be little landslides down the side of the pile. While the pyramidal shape of the pile and, in particular, the angle of the side are predictable, depending on the properties of the average sand grain, the timing, location and scale of the landslides are unpredictable because the slippage is nonlinear. Once a grain of sand starts sliding, it pulls others along with it and there is positive feedback leading to a mass of sand slipping (Bak 1996). Similar nonlinearities are thought to cause stock market crashes.

From the point of view of the scientist or mathematician, nonlinear systems are difficult to study because most cannot be understood analytically. There is often no set of equations that can be solved to predict the characteristics of the system. The only generally effective way of exploring nonlinear behaviour is to simulate it by building a model and then running the simulation (see Chapter 6). Even when one can get some understanding of how nonlinear systems work, they remain unpredictable. However much one studies stock markets or the properties of sand, it will still be impossible (in principle) to predict the timing of a crash or a landslide.

This does have some lessons for explanation in the social sciences. For

instance, conventional philosophy of social science has often made too ready a connection between explanation and prediction. It tends to assume that the test of a theory is that it will predict successfully. This is not a criterion that is appropriate for nonlinear theories, at least at the micro scale. Complexity theory shows that even if we were to have a complete understanding of the factors affecting individual action, this would still not be sufficient to predict group or institutional behaviour. The message is even stronger if we make the plausible assumption that it is not only social action that is complex in this sense, but also individual cognition (Conte and Castelfranchi 1995).

Emergence

A formal notion of emergence is one of the most important ideas to come from complexity theory. Emergence occurs when interactions among objects at one level give rise to different types of objects at another level. More precisely, a phenomenon is emergent if it requires new categories to describe it which are not required to describe the behaviour of the underlying components. For example, temperature is an emergent property of the motion of atoms. An individual atom has no temperature, but a collection of them does.

That the idea of emergence in the social sciences is not obvious is attested by the considerable debate among sociologists, starting with Durkheim (1895), about the relationship between individual characteristics and social phenomena. Durkheim, in his less cautious moments, alleged that social phenomena are external to individuals, while methodological individualists argued that there is no such thing as society (for example, Watkins 1955). Both sides of this debate were confused because they did not fully understand the idea of emergence. Recent social theorists (Kontopoulos 1993; Archer 1995; Sawyer 2001, forthcoming) are now beginning to refine the idea and work through the implications. Simulations can provide a powerful metaphor for such theoretical investigations.

There is one important caveat in applying complexity theory to social phenomena. It appears to leave human organizations and institutions as little different in principle from animal societies such as ants' nests (Drogoul and Ferber 1994) or even piles of sand. They can all be said to emerge from the actions of individuals. The difference is that while we assume that, for instance, ants have no ability to reason – they just follow instinct and in doing so construct a nest – people do have the ability to recognize, reason about and react to human institutions, that is, to emergent features. The institutions that result from behaviour that takes into account such emergent features

are characteristic of human societies (for example, governments, churches and business organizations). The emergence of such reflexive institutions is called ‘second-order emergence’ and might be one of the defining characteristics of human societies, distinguishing them from animal societies (Gilbert 1995). It is what makes sociology different from ethology. Not only can we as social scientists distinguish patterns of collective action, but the agents themselves can also do so and therefore their actions can be affected by the existence of these patterns.

A theoretical approach that was originally developed within biology, but which is becoming increasingly influential because it takes this reflexive character of human interaction seriously, is known as autopoietic or self-organization theory (Varela *et al.* 1991; Maturana and Varela 1992). Autopoietic theory focuses on organisms or units that are ‘self-producing’ and self-maintaining. An autopoietic system is one that consists of a network of processes that create components that through their interactions continuously regenerate the network of processes that produced them. Social institutions and cognitive systems have both been analyzed in these terms by Maturana and Varela (see also Winograd and Flores 1986). The emphasis on process and on the relations between components, both of which can be examined by means of simulation, accounts for the developing link between this theoretical perspective and simulation research.

Simulation can also usefully be applied to theories involving spatial location and rationality, two topics that have often been neglected in social science, but which are increasingly recognized to have profound implications. Geographical effects can be modelled by locating agents on a simulated landscape, faithfully reproducing an actual terrain – see, for example, Lansing’s (1991) simulation of the irrigation system in Bali – or on the regular grid of cells used with a cellular automata model. Rationality (Elster 1986) can be modelled using the artificial intelligence techniques described in Chapters 8 and 9, but often the main concern is not to model general cognitive capability, but to investigate the consequences of bounded rationality. For example, some theories about markets assume that traders have perfect information about all other traders and all transactions and are able to maximize their own profits by calculating their optimum strategy on the basis of all this information. In large markets, this is obviously unrealistic. What are the consequences for markets reaching equilibrium if the traders have limited information and limited capacity to process that information? Epstein and Axtell (1996: Chapter 4) describe a model they constructed to study the effect of decentralized markets where traders possess only local information and bounded rationality.

Conclusion

In the following chapters, we shall consider in turn the main techniques available for building simulations. These techniques each have their own specific characteristics and areas of application. In Table 1.1, the ‘number of levels’ refers to whether the techniques can model not just one level (the individual or the society) but the interaction between levels. A technique capable of modelling two or more levels is required to investigate emergent phenomena. Some techniques allow the modelling of communication (for example, the passing of messages) between agents and so are appropriate for modelling language and interaction; others do not. The techniques based on artificial intelligence (distributed artificial intelligence and learning models) are able to accommodate sophisticated agent designs; others derive some of their benefit from constraining the researcher to very simple agents. Finally, most techniques are able to handle the large number of agents that one would expect to find in social simulation, although the first to be considered here, system dynamics, is oriented to the development of models of a whole system, where the system itself is the one and only agent simulated.

Table 1.1: A comparison of social science simulation techniques

<i>Chapter</i>	<i>Number of levels</i>	<i>Communication between agents</i>	<i>Complexity of agents</i>	<i>Number of agents</i>
3 System dynamics	1	No	Low	1
4 Microsimulation	2	No	High	Many
5 Queuing models	1	No	Low	Many
6 Multilevel simulation	2+	Maybe	Low	Many
7 Cellular automata	2	Yes	Low	Many
8 Multi-agent models	2+	Yes	High	Few
9 Learning models	2+	Maybe	High	Many

We have suggested in this chapter that simulation has a number of valuable features for social science research. One of the clearest is that it is well adapted to developing and exploring theories concerned with social processes. In comparison with some other methods of analysis, computer simulations are well able to represent dynamic aspects of change. A second important feature of simulation is that it can help with understanding the relationship between the attributes and behaviour of individuals (the ‘micro’ level) and the global (‘macro’) properties of social groups. That is, it is possible to use simulation to investigate emergence.

Simulation is akin to an experimental methodology. One can set up a simulation model and then execute it many times, varying the conditions in which it runs and thus exploring the effects of different parameters. Experimental research is almost unknown in most areas of the social sciences, yet it has very clear advantages when one needs to clarify causal relationships and interdependencies. However, while simulation has similarities with experimentation, it is not the same. The major difference is that while in an experiment one is controlling the actual object of interest (for example, in a chemistry experiment, the chemicals under investigation), in a simulation one is experimenting with a model rather than the phenomenon itself.

We shall develop this idea further in the next chapter, which is concerned with the methodology of simulation research.