

1. Why simulate innovation?

This book seeks to innovate in the tools we use for thinking about innovation. Computer simulation models can clarify our thoughts and explore their implications. Over the last two decades there have appeared descriptions of computer simulation models that address some of the issues surrounding innovative ideas, practices and technology, including how innovations can be generated, how they diffuse among people and organizations, and the impact innovations have on people's and organizations' other ideas, practices and technologies. This book will provide a critical survey of some of these tools for thinking, while also introducing a few tools of our own.

In this chapter we explain why one might want to be thinking about innovation, how it involves complex adaptive systems and how these can be studied, and hence why one might want to add computer simulation models to the tools one uses for innovation studies. The chapter concludes with an outline of the rest of the book.

WHY STUDY INNOVATION TODAY?

The Trouble with Financial Innovation

Innovation is currently held responsible for a lot. During the work for this book (2009–12) countries around the world have been suffering the after-effects of a wave of innovation in the financial world. The tale, as told by *Financial Times* journalist and trained social anthropologist Gillian Tett (2009), tells of brilliant minds being hired by investment banks and, full of excitement for their work, putting in long hours to generate innovative ways of making money (see also MacKenzie, 2009, 2011a). They began with the idea of extending the centuries-old concept of derivatives, a form of insurance, to a new application, that of insuring against the risk of a borrower defaulting on their debts. The tale is woven around a diverse collection of novel financial concepts and products, each requiring a new name or phrase: from CDS to CDO, CDO of ABS, mortgage-backed CDO, slice-and-dice, tranches, CDO-squared, Gaussian

copula, sub-prime, super-senior, SIV and ABCP. Each one represented a new combination of pre-existing components or a new application for an existing idea. Once invented, the innovations were offered to new markets, scaled up to new levels, and sold off in unprecedented quantities to a variety of customers. These customers included not only traditional, but also some new financial players, most of whom had little knowledge about how these financial products worked. Understanding of the risks and value behind these novel products was also scarce among their producers, their owners, government regulators and insurers, but few saw any incentive for asking questions. When the underlying assets, mostly mortgages on houses, began to lose their market value, a house of cards was set to tumble, freezing markets and taking down banks and brokerages, insurers, investment and hedge funds, government finances and even governments themselves, and finally reduced the power of several nations. All this followed a frenzy of innovation.

At the same time, innovation is at the heart of proposed solutions to the crisis. Governments should introduce new regulations for the banking sector. R&D spending should be increased in other industries, especially manufacturing, in order to generate new growth to compensate for the losses. The gap between universities and businesses must be narrowed, with exchanges of knowledge between them, and more patents generated and spin-off companies set up based on academic research. Geographic clusters of firms must be seeded and protected, where interactions between the firms will generate the next big ideas in technologies. So innovation has been hero, then villain, and is now our best hope for salvation. It seems an apt moment to be writing about innovation.

The Trouble with Economics

Given the economic causes and effects of innovations such as those in the financial world, it might be thought that the topic of innovation would best be studied by economists. The primary focus of mainstream economics is efficient resource allocation, for which mathematical models have been developed based on the idea of a system in equilibrium. Solow (1956) provided a mathematical treatment to add resource growth to modelling as part of a dynamical equilibrium theory, but these models assume both population growth and technological change are givens, exogenous to the model. By this light, technological innovation is just an unexplained leftover when one has subtracted other factors behind resource stocks. Endogenous growth theory (e.g. Romer, 1986, 1990) considers some of the factors thought to be behind technological change, chiefly those that increase human capital, knowledge and innovation, such as R&D

spending, the level of government regulation and a culture of openness to change. A key difference from previous economic theories is the idea that investing in R&D can produce *increasing returns* to scale: acquired knowledge enables improvements in future knowledge production.

These attempts to study innovation endorse most of the common assumptions of mainstream economics, such as rational agents forming systems at equilibrium, and largely consist in developing equation-based models that will reproduce statistical patterns observed in data, in this case, data by country on GDP and growth, population size and R&D spending, among other measures. In so far as correlations are found between these variables, how the correlations come to be there is poorly understood. Representing the generating mechanisms means representing human behaviour, including representing its diversity, mathematically in such a way that it can be aggregated easily.

While mainstream economics remains attached to its assumptions and mathematical techniques, it continues to treat the topic of innovation poorly. This can be seen by the continued neglect in mainstream economics textbooks of fields that deal primarily with innovation: evolutionary economics and behavioural economics.

The pioneer of *evolutionary economics*, Joseph Schumpeter (1939, 1943), writing in the middle of the twentieth century, identified innovations and the entrepreneurs who develop innovative ideas into marketable products as vital to economic growth (Heilbroner, 2000, Chapter 10). The theories of neoclassical economics focus on markets at equilibrium. But according to these theories, at the equilibrium point, competition between firms has reduced profit to zero. In this case, why remain in the market? This seeming puzzle could be solved, according to Schumpeter, by reference to innovation. When companies bring new products to market, or develop improved methods of production resulting in lower costs, they enjoy an advantage over their competitors and may charge prices that include a premium, thus yielding non-zero profits. Their new offerings may also enhance the value of other goods and services, and undermine the market appeal of yet others, a process Schumpeter dubbed ‘creative destruction’ (Schumpeter, 1943, p.83). The advantage is only temporary, however, since competitors may imitate the innovator. For this reason, some of the innovator’s profit should be invested in the R&D that could generate future innovations and maintain some competitive advantage. Alongside efficient allocation of resources, forcing firms to innovate is the second major justification for markets. But uncertainty exists about how much to invest in this R&D, how best to go about seeking innovations, how much one innovation depends on knowledge of another and how long it will take to generate the next one. Different companies may adopt different

strategies for this, with some investing heavily in R&D and others hoping to be able to imitate quickly and cheaply when the investments by others have generated results. At some times there may be a flurry of new products, at other times the diffusion of recently introduced products, and at yet other times there may be a period of relative quiet, perhaps resembling a market equilibrium state. Thus, while undergraduate economics courses teach students to focus on the ideas of equilibrium being reached by a market of identical competing firms, the vision developed from Schumpeter's work is that of *heterogeneous* (diverse) firms in a *dynamic* market.

Another field trying to attract more attention within economics is *behavioural economics*. When reasoning about the decisions made by suppliers and customers, neoclassical economics assumes that decision makers know all the available options, the probabilities and monetary values of all consequences of these options, and will choose between the options so as to maximise their expected monetary gain. This view of human decision makers as rational optimisers with perfect information, or *homo economicus*, was criticised by the political scientist, Herbert Simon, beginning in the 1940s and continuing in the decades since (Simon, 1948, 1955a, 1957, 1991). In its place, Simon and collaborators proposed that human decision makers had limited information on options, probabilities and values, and limited ability to process the information they had in a short enough time for it to be useful. Instead, of being infinitely capable rational optimisers, 'bounded rational' humans employed relatively quick and easy rules of thumb, called heuristics, to search for solutions that were, if not the best possible, usually sufficiently good for survival (Simon, 1955a; Simon and Newell, 1958). Nelson and Winter (1982) combined this view of bounded rational agency with evolutionary economics. Laboratory experiments by psychologists Kahneman, Slovic and Tversky (1982) confirmed that how human beings actually performed decision making resembled the use of heuristics more than it did mathematical optimisation. Both Simon and Kahneman have since been rewarded with Nobel Memorial Prizes in Economic Sciences (in 1978 and 2002, respectively). In the 1990s, support for research into actual economic behaviour continued to build (Akerlof and Shiller, 2009; Kahneman, 2011; Klein, 1998). More recently, interest has grown in the study of what it is that decision makers seek to improve, in particular, happiness (Frey, 2008; Layard, 2011), instead of money. Despite this, an informal survey of the undergraduate-level textbooks in the economics sections of bookshops and libraries reveals that most still lack chapters devoted to either evolutionary or behavioural economics.

Following the financial crisis, however, confidence in mainstream economics has been shaken (Blanchard, 2012; Frydman and Goldberg, 2011;

Keen, 2011; Turner, 2012). There is an opportunity for rethinking the subject's core material, that is, what is taught to students, and also what is funded, what research is published in the most widely read journals, who gets employed by the most prestigious academic institutions and who will go on to influence the next generation of society's leaders. Time and effort is being devoted to innovative approaches, be these either the invention of new methods, or the importing of ideas from other fields, including psychology, sociology, neuroscience, cognitive science, biology and the various fields which study complex adaptive systems.

New Sources on Innovation

The information age has brought new data sources to help the change in focus. There is more emphasis on attempts to count innovations. In technology there are data on patents, including who patents what, who they patent it with and which patents refer to which others (Fleming, Mingo and Chen, 2007; Fleming and Sorenson, 2001; Sorenson, Rivkin and Fleming, 2006; Trajtenberg, 1990). Similarly, data on academic publications, their co-authors and their citations, give insights into innovation production within universities and other research institutes (Boerner, Maru and Goldstone, 2004; Goldenberg et al., 2010; Price, 1965; Small, 1973). Electronic records of individuals' interactions, such as email communications, the Internet Movie Database (www.imdb.com) or geographical tracking devices can provide impressions of the social networks within which information about innovations flows and ideas are combined to generate new innovations.

In addition to these quantitative sources of data, qualitative sources, especially ethnographic studies over the last 30–40 years, have caused a revision of views of innovation generation and adoption. Seen close up, the supposed events of invention and adoption of new ideas, practices or products become more complex and less identifiable (Akrich, Callon and Latour, 2002; Akrich et al., 2002; Bijker, 1995; Bijker and Law, 1992). Since the 1990s, developments in artificial intelligence, robotics and cognitive science (Clark, 1997; Hutchins, 1995a, b) have promoted a view of the human decision maker as being *embedded*, *embodied* and *social*, with decisions dependent on a historical context, on interaction with a material environment and on collective effort.

It remains to be seen whether analyses of these quantitative and qualitative datasets will lead to better policies on innovation. Some uses of the datasets, such as policies that attempt to base continuance of funding on past production of patents or publications, could cause innovators to adapt their behaviour from that which helped generate the past data.

Unlike, say, astrophysicists, social scientists have the potential to disturb the systems they study. However, where policy and behaviour has yet to reflect the results of analyses, the datasets may help us to understand retrospectively how innovations were generated, how they interrelate, how they diffuse and what their impact may be.

Both quantitative and qualitative studies can inform the creation and revision of theories about innovation, which in turn can inform policy making. Theorising, however, can be hard to perform in unambiguous, coherent detail, with its implications spelled out. The time is ripe for a technique that allows theorising to capture some of the complex networks of interdependencies, and the dynamic behaviour that results. In recent decades a new type of tool has emerged for improving the rigour of theories and exploring their coherence and consequences, generating new hypotheses for empirical studies (Davis, Eisenhardt and Bingham, 2007). These are computer simulation models, and this book applies them to the study of innovation. In this we draw upon papers and books by others that have appeared over the last 20 or so years. These works apply simulation models to the diffusion of innovations through social networks, to collective learning in organisations, to the structure of academic science publications, to the adoption and adaptation of technologies in complex contexts and to technological evolution and the formation of innovation networks, to name the major topics of our chapters. Given that innovation remains as important an area as ever, and given the numbers of these tools, it seems a good time to highlight some of the models, including their features, assumptions and purposes, and identify some recommendations for future models.

WHAT IS MEANT BY ‘INNOVATION’?

A Few Common Distinctions

There are many uses for the word ‘innovation’, and uses in this book will reflect several different bodies of literature, although the authors of models can be quite vague about the types of innovation they intend to apply them to. A few common distinctions may be made, however.

Two ideas seem essential to the concept of innovation. The more obvious idea is that it involves *newness*, or novelty. For example, there may be a new item or service brought to market (*product innovation*), or a new method for producing a product more cheaply than before (*process innovation*). The second idea is that the new thing will be of some value to someone, that is, it will be an improvement, reaching a new level of

quality, useful, relevant or appropriate to the concerns of some person or persons, be they customers in a market, academics in a scientific field or workers in a firm. Combining, for example, a banana with a spreadsheet may be an original or novel idea, but if it has no use or makes no sense to anybody, it seems pointless to treat it as an innovation. This does not mean bananas and spreadsheets will always remain apart, however. We cannot rule out the possibility that someone, somewhere, will one day find some meaning in that particular combination. Indeed, it may be that someone already has, but that the news of this has yet to reach us. It would seem, then, that both ideas implicit in innovation, novelty and appropriateness, are relative to some particular audience. Like beauty, innovation is in the eye of the beholder.

We will write about ‘innovation’ as a process and about ‘innovations’, the objects of the process. In fact, there are two types of process going by the name of innovation. The first is the *generation* of an innovation, for example, the combining of two ideas to make a third idea. The second is the introduction of an innovation to a group or a market, and its spread thereafter, a topic usually called the *diffusion of innovation*. However, both generation and diffusion will be represented in most of our models.

There is a formal model of innovation taught in business schools called the *linear model*, which focuses attention on a division into separate stages, innovation generation followed by innovation diffusion (Godin, 2006). Sometimes the first phase is divided further, between the results of basic research, such as a prototype or a discovery, and the results of development of the research results into an application or marketable product (e.g. Cooper, 1990). In recent decades, these divisions have received much criticism (Balconi, Brusoni and Orsenigo, 2010). Some of the models in this book will represent the distinction, others will reflect the critiques.

Part of the grounds for dispute is that an innovation can be both a physical construction, and also something less tangible: an idea, practice or meaning. The former might be, for example, an object produced by the particular combination of components which have never before been combined in this way (*recombinant innovation*). This novel object, however, may then be used for an old, familiar purpose. Conversely, an existing, familiar object may be used in a novel way, and given a new interpretation or value (*transfer innovation*). If the object is then modified physically in order to improve its ability to enact the new application, is this an additional innovation? Likewise, if our behaviour adapts to a new object because the old purpose is only imperfectly served by it, should we count this an additional innovation? Again, the models in this book will vary in where they identify innovations.

Innovations, whether products or processes, have relations to each other. Some new things can replicate, partially or wholly, the functions of others. These can then serve as *substitutes* for each other. Other new things enhance the functionality and value of existing products and processes. These are *complementary*. In some cases, the effect of one item on the ability to produce or use another may be quite strong. If product B does not work without the use of product A, product A having no substitutes, then A is *necessary* for B. If B is an automatic consequence of A being used, with no other products necessary, A is *sufficient* for B. *Dependency relations* between things will play an important role in innovation studies.

Several typologies found in the innovation literature are worth mentioning. First, a distinction is sometimes drawn between *qualitative* and *quantitative* innovations. The invention of the aeroplane, combining the internal combustion engine from automobiles with the wings-and-tail airframes of gliders, is a qualitative innovation. It produced capabilities very different to those of the car or the glider, including airborne reconnaissance, bombers and fighter planes. In comparison with ships and trains, however, it offers primarily a quantitative improvement: faster transport of people. Whether one sees an innovation as quantitative or qualitative depends on which dimensions one focuses on and which comparison technologies one selects.

Other distinctions to be drawn are those between *radical*, *incremental*, *modular* and *architectural* innovations (Henderson and Clark, 1990). Consider an electric air fan, mounted on the ceiling of a room. Improvements to the blades or the motor would be incremental. Replacing it with another technology, such as air conditioning – based on a very different principle and physical phenomenon but aimed at similar effects – would be radical. The fan blades and motor could be reassembled in another way, as a desk-mounted portable fan. Such rearrangement of component parts would be architectural innovation. Keeping the configuration of components, but replacing one of them with a new technology – such as a new type of motor to replace the electric one – would be a modular innovation. Incremental innovation maintains both the core design concepts of a technology and the linkages between its concepts and components. Modular innovation involves a change in the core concepts. Architectural innovation involves a change in the linkages. Radical innovation involves both. Henderson and Clark, who introduced this framework, admit that ‘the distinctions between radical, incremental and architectural innovations are matters of degree’ (Henderson and Clark, 1990, p. 13). But they invoke the framework to explain the relative degrees of disruption that technological innovations can cause. If a producer firm

is well-established, incremental changes build upon its core competences and have little effect on its strategy and organisation. Radical innovations render its expertise obsolete and can be devastating to the way it runs itself. Architectural innovation, however, mostly preserves the usefulness of the firm's knowledge of the components, while demanding that the firm rethinks how it uses the components. The combination of giving up some areas of expertise while preserving others may prove difficult for an established firm – in contrast, perhaps, to a young start-up firm with no emotional attachments to the expertise of particular staff members, and no financial investments in particular production machinery. A modular firm, that mirrors the structure of its product technology by its organisational structure, may find it easier to handle modular and architectural technological innovations as changes respectively within and between organisational units. Thus technological and social organisational structures can interact.

Creativity Myths and Some Insights into Innovation

As human beings we love to tell stories, including stories about inventions, discoveries and how some important component of our present lives came into being. But what makes for a good story may not reflect the real processes of innovation. Sawyer (2012) notes that Western societies employ a number of 'creativity myths' concerning how the fruits of creativity, or innovations, come about. Inventions are given mysterious origins: the flash of genius, the bolt from the blue or divine inspiration. We also like to think of creation being the act of a lone individual: the misunderstood artist starving in his garret, producing masterpieces that will not sell until he is dead; or the mad scientist, living in obscurity, and making the discoveries that will go unrecognised for 50 years. Even our legal systems reflect a focus on individuality, with patented inventions giving special rights to the holders, who may or may not have been the first to have the patented idea, and who may or may not be capable of developing a commercially viable product from it.

The reality may be more complicated. Sawyer lists a number of insights into business creativity (Sawyer, 2012, p. 285), which he illustrates with the case of the development of the graphical user interface (GUI) of windows, icons, menus and pointers (WIMP) by Xerox, Apple and Microsoft. We repeat the insights here (in italics below), but illustrated with our earlier case of credit derivatives (Tett, 2009).

Each innovation builds incrementally on a long history of prior innovations. One might wish there to be some invention or decision event, to which one could point saying, 'Ah! That was when credit derivatives were

created which would destroy the financial world!' As mentioned already, the real history does not oblige this wish, with developments going back decades and even centuries. Concerning the GUI, the elements of the WIMP paradigm came together from several sources and several stages, including academic research in the 1960s, demonstrations of the personal computer concept by Xerox in the 1970s and computers intended for mass production launched in the mid-1980s. The market dominance of Microsoft's *Windows*, today the best selling operating system, obscures the fact that it took major revisions before a PC-compatible GUI achieved commercial success with *Windows 3.0* and *3.1* in the early 1990s and another major revision, *Windows 95*, was needed to attain the GUI standards set by mid-1980s computer platforms.

Innovations emerge from collaborative teams. Tett describes a collection of brilliant individuals with international backgrounds and some diversity in education and career paths. Among them were people with strengths in banking itself, mathematical modelling and customer relations. Beyond this team there were lawyers and technical services, representatives of the customers, and regulators, all of whom had some input into the development of the new types of deals. If the development of an innovation involves multiple insights, each of these probably comes from a different team member. But a single person usually comes to be associated with the innovation, a Thomas-Edison type, who then gains the most in reputation from it. In the case of credit derivatives, one member of the J. P. Morgan team became a spokesperson for the new field, with many media appearances, and reportedly has received hate mail since the financial crisis emerged. However, as Tett points out, the team at J. P. Morgan in no way intended or anticipated how their idea would be reapplied at other banks, and at some stages were alarmed at the growth in scale of the new markets.

Multiple discovery is common. Tett focuses on one team at J. P. Morgan, who have some claim to having been the first to arrange a credit default swap deal. But rival firms were quick to copy this type of deal, partly because they were already familiar with the concepts involved, and competition during the growth years of the credit derivatives markets was intense. In the case of the multiple GUI developers, the extent to which they acted independently rather than imitated was a matter for legal disputes.

There is frequent interaction between the teams. In the case of the GUI, both Apple's Steve Jobs and Microsoft's Bill Gates visited Xerox PARC and saw some of their developments. In the case of banking services, staff are often poached by rival firms, sometimes taking whole groups of colleagues with them. In addition, mergers between firms led to rival teams being brought together.

A product's success depends on broad contextual factors. In the case of the GUI, the cost of a personal computer employing the new software designs was prohibitive until the late 1980s. It took the spread of the idea of desktop publishing, a so-called 'killer app', to boost sales of Apple's Macintosh computer. The growth of the credit derivatives markets was helped by a number of contextual factors. To fuel it, there was a plentiful inflow of money from pension funds, oil sheiks, a policy of cheap lending from the US Federal Reserve, and a growing trade deficit with China. From the 1980s on there was a celebration, especially on the part of politicians, of the entrepreneurial work performed in financial markets, including the invention of new product types, the adoption of new ways to trade, such as computer-based trading, and acceptance of new levels of risk taking in order to seek out the best returns for one's clients and shareholders. These cultural factors led to the removal of old regulations, the removal of restrictions on transactions, especially cross-border ones, and a relaxed, *laissez-faire* approach on the part of governments to the introduction of regulations to deal with the new financial products. The culture also infected people outside the financial industry, such as local government financial operators, who sought riskier places to invest money on the grounds that it would maximise returns – not previously thought of as an aim of local government. Even when concerns were expressed about the new markets, especially their rapid growth and sheer scale and the lack of accurate assessments of these, the growing practice of paying lobbyists to represent the interests of businesses to politicians meant that attempts to regulate the financial sector more closely were headed off or watered down. Furthermore, the novelty of the products obscured what was familiar. While some might have identified the boom as an old-fashioned market bubble, and plenty have since noted the parallels with the Wall Street Crash and the resultant Great Depression, at the time the novelty of the products meant that, as with the Dot.Com boom in the 1990s, people could argue that 'this time is different'. So, rather than one person, one team or one bank being responsible for the products behind the financial crisis, we find whole societies collectively creating the situation they then find themselves in.

Innovations, then, despite the creativity myths, seem to be the result of collective labours from many participants, in many locations, taking many steps, involving many components with many connections between them. Innovations emerge from the interactions of a complex system of social and technological parts. If we are to think about innovation then we need tools for dealing with complex systems and the emergence of novelties from these.

INNOVATION AND COMPLEX SYSTEMS

The Growing Interest in Studies of ‘Complexity’

Having suggested some reasons for studying innovation and what that might involve, we turn to the second theme of this introductory chapter, and of the book itself: complexity, emergence and complex adaptive systems.

An innovative academic field has emerged over the last few decades with complex systems as its focus. It is best encapsulated by the Santa Fe Institute (SFI), founded in the 1980s and still serving as an inspiration for researchers (www.santafe.edu/about/history). Part of this influence stemmed from the publication of engagingly written early accounts of its work by popular science writers (Lewin, 1993; Waldrop, 1993), but from early on, the subject matter was engaging as well. Unlike most funded research projects of the time it did not engage in providing confirmation of small-scale, pre-specified beliefs within established academic disciplines, but instead aimed at bigger, more fundamental questions, and transcended disciplinary boundaries. SFI brought together established researchers from multiple disciplines, beginning with physicists from nearby Los Alamos, and adding economists when Citicorp, keen to develop alternatives to mainstream economic thinking, offered to fund some economists known for their more maverick interests in technology and evolution. Biologists, mathematicians, anthropologists and computer scientists followed, as well as cognitive scientists and psychologists in the 1990s, and the occasional artist, especially writers. The biological sciences made particularly important contributions, including theoretical biologists debating evolution and adaptation, the origin of life as self-organising systems, the emergence of cooperation, and the population dynamics of ecosystems, where complex networks of who-eats-whom relations made different species interdependent for their evolutionary success in complicated, impossible-to-predict ways.

A common thread in SFI work throughout its history was the development of new computer tools for studying complex systems, including agent-based simulation models (which we will return to in a later section), statistical data analysis and pattern-recognition tools, and problem-solving heuristic search methods, such as genetic algorithms. Many SFI researchers believed that they each faced analogous problems which might be tackled using ideas and techniques inspired by each other’s disciplines. Once SFI acquired buildings, visiting scholars from all over the world and from multiple disciplines were able to share offices or meet in corridors and the canteen, and discuss ideas for each other’s problems.

A common feature of these problems was that they involved systems composed of multiple interdependent parts, and the behaviour of the system could not be equated to a linear sum of the behaviours of the parts. Unlike Newtonian mechanics, where vector sums are made to calculate the overall behaviour of some combination of forces, mainstream econometrics, where market behaviour was assumed to be an aggregation of individual customers and suppliers, or social statistics employing linear regression models, in these complex systems behaviour was assumed to be *nonlinear*. The scale of the effects was not proportional to the scale of the causes. For instance, doubling the person-months invested in a project might not double the output or halve the project duration. More people mean more interdependencies between them, and more time means more opportunity for adapting to each other and the task. The overall effects of additional resources may difficult to anticipate. Experience of software development taught that adding manpower to a late project makes it later (Brooks, 1975).

Also common to these problems was that the systems involved were dynamic, or changing, with component parts continually adapting to each other, in contrast to mainstream economics' focus on systems at rest, or equilibrium. Among the phenomena that could be found in nonlinear, complex adaptive systems' behaviours were phase transitions, sudden shifts in the nature of the system, in response to perhaps only small changes in a single parameter. Phase transitions were familiar to physicists from the study of matter, for example where gradual increments in temperature can cause ice to melt quite suddenly to a liquid around 0°C, and boil quite suddenly around 100°C. Another key concept was emergence, the surprising appearance of some kind of order or pattern where previously there had been only disorder. The best known phase transitions were those between order and disorder, or 'chaos' as it was popularly known, though no real connection was ever established from complexity to chaos theory, made famous slightly earlier (Gleick, 1987). Systems frozen in ordered states had no interesting consequences; neither did those in random flux. Interesting phenomena in nature and social systems could be found in-between, 'at the edge of chaos' (Lewin, 1993; Waldrop, 1993). This evocative phrase quickly spread to, among other places, the business literature (Cohen, 1997; Conner, 1998; Pascale, Millemann and Gioja, 2000) where the new 'science' was mined for metaphors for how organisations should be run: not planned and controlled by management at the top – businesses and economies as complex adaptive systems were too unpredictable for that – but instead directed from the bottom up, with workers and other system components given sufficient freedom for the business to self-organise and its new policies to emerge.

After a wave of publicity, SFI encountered scepticism about whether a general theory of complex systems was a sensible goal – complex social systems might have very different laws to complex biological systems, for example – and whether any contributions recognisable to other scientists would emerge amid all the dreams and hype (Horgan, 1995). In response, SFI researchers began to rely less on toy models of abstract complex systems, and more on empirical studies. In this they were helped by work from elsewhere that raised awareness of the importance of network structures in biological and social systems (Barabási, 2002; Buchanan, 2002). New technologies meant that large-scale datasets were being generated in biological and social sciences that could be analysed to test hypotheses about how the real systems functioned.

The Multiple Discoveries of Complex Systems

It should be noted that the researchers meeting at the SFI were neither the first nor the only ones to think about complex systems. Bronk (2009), for example, argues it is possible to identify many of the complexity ideas relevant to economists in the works of the Romantic philosophers of the early nineteenth century. Even if attention is restricted to the second half of the twentieth century, there are still plenty of examples, of which we can cite a few.

The work of Herbert Simon and his collaborators from the 1940s onwards at what is today Carnegie-Mellon University has already been mentioned. Simon's paper on 'The architecture of complexity' (Simon, 1962) was an early contribution, arguably decades before its time. His former co-workers pioneered the computer simulation of organisations, including how individual workers could collaborate to learn more than they could in isolation (Cyert, March and Clarkson, 1964; March, 1991). Chapter 4 will draw upon some of these ideas. This work was combined with themes from evolutionary economics (Nelson and Winter, 1982).

Also stemming from the 1940s and '50s, cybernetics considered the dynamics of systems, including business organisations, for the interdependencies between their parts, especially *feedback* and *feedforward loops* that could regulate a system or send it spiralling out of control (Ashby, 1956; Wiener, 1948). One result from cybernetics was Ashby's law of requisite variety, which holds that for a subsystem to be controlled (for example, by a manager) the variety in its behavioural states must be matched by the variety in the states of the controlling system (Ashby, 1956, 1958). The General Systems Theory of Bertalanffy (1971) attempted to apply the concept of a biological system, the organism, to the phenomena of other disciplines. Both cybernetics and systems theory owed their appeal to a

belief that studies of systems in general could inform attempts to manage human systems (Beer, 1959). Systems thinking prompted the development of system dynamics modelling (Forrester, 1961; Sterman, 2000), a computer simulation of stocks and flows of materials, people or other quantifiable things, in which the behaviour of any one stock level could depend on the levels of stocks and flows elsewhere in the system. Another outcome of systems thinking was a focus on how firms adapt over time in response to experience of an environment, in the *learning organisation* (Senge, 1992).

Those schooled in the mathematical techniques of operational research (OR) or management science became increasingly concerned during the 1970s that the real-world problems to which they tried to apply their skills did not resemble textbook exercises or idealised conditions (Rosenhead and Mingers, 2001). Instead, organisations presented them with ‘messes’ (Ackoff, 1981), ‘wicked problems’ (Rittel and Webber, 1973) and ‘swamp conditions’ rather than the high ground (Schön, 1987). Indeed, often the biggest challenge was in identifying what problems were faced within the organisation, rather than in solving the problems. Since then expertise has been developed in so-called ‘Soft OR’, applicable when systems include hard-to-quantify phenomena, such as opinions, norms, politics and emotions. Techniques for facilitating group discussions and involving stakeholders are employed in order to improve members’ understanding of their own organisation’s situation, seek consensus as to what problems should be dealt with, and generate more buy-in for candidate solutions (Rosenhead and Mingers, 2001). Examples include soft systems methodology (Checkland, 1998; Wilson, 2001) and causal mapping/cognitive mapping (Ackermann and Eden, 2011; Eden, 1988). Outside of business organisations, those working in the field of human–environment relations, studying social-ecological systems and their sustainable development, have adopted ideas from complexity science and a belief in the value of stakeholder participation (Berkes, Colding and Folke, 2003; Voinov and Bousquet, 2010).

Working primarily in the sociology of science (to which we will return in Chapter 5), Robert Merton identified several phenomena relating to how a complex society produces and uses innovations, including ‘multiple independent discoveries’ (Merton, 1973, Chapters 16–17; Zuckerman, 1979), ‘unintended consequences’ (Merton, 1968b, Chapter 15), ‘self-fulfilling prophecies’ (Merton, 1968b, Chapter 16) and serendipity (Merton and Barber, 2004). Another sociologist, Luhmann, drew upon Maturana and Varela’s work in biology to pursue a theory of communication and self-reference, based on their concept of auto-poietic (‘self-producing’) or self-organising systems (Luhmann, 1990; Maturana and Varela, 1980; Mingers, 1995).

Anthropologists had long been aware of the complexity of the social systems they studied, but the application of ethnological methods to study how scientists worked and how technologies were developed and used transformed understanding in Science and Technology Studies in the 1970s and 1980s (Bijker, Hughes and Pinch, 1987; Bijker and Law, 1992). (We will return to this in Chapter 6.) In the new picture, human agents, technologies and practices become intertwined in a complex network of social, economic, political and physical relations.

Interest in social networks began with empirical explorations of network-related social phenomena, such as the small-world effect (Killworth and Bernard, 1978; Milgram, 1967) and the strength of weak ties when seeking information on job opportunities (Granovetter, 1973), and with the development of statistical metrics for network structures in social network analysis (Boorman and White, 1976; Lorrain and White, 1971; Wasserman and Faust, 1994; White, Boorman and Breiger, 1976).

A number of physicists became interested in the phenomena of dissipative thermodynamic systems, that is, systems which take in free energy from the outside (such as from the sun) but partially constrain its dissipation, so that on a local scale (e.g. on planet Earth) ordered structures can build up and entropy decrease, in contrast to the second law of thermodynamics, which concerns closed systems and the increase of entropy. In this respect, the physicist Prigogine's work on self-organisation and irreversible systems is particularly notable, including his collaboration with the philosopher Stengers (Prigogine and Stengers, 1984). From the 1970s on, other physicists had begun to apply the techniques of statistical mechanics to social interactions (Galam, 2004), though it would take until the 1990s for this work to become well known, for example Bak's popularisation of his theory of self-organised criticality (Bak, 1997). In particular, social physicists have turned their attention to social networks, including relating network structural properties to processes of network growth and change (Albert and Barabási, 2002; Barabási, 2002; Newman, 2003, 2010; Newman, Barabási and Watts, 2006).

So, many of the ideas promoted by the SFI as part of a 'Complexity Science' or 'Complexity Theory', such as the importance in social systems of complex networks of interdependencies, were available from alternative sources. Many past researchers have been interested in transferring biological metaphors and formal or mathematical models of interdependent systems to other fields, especially those involving people. However, SFI has been a powerful inspiration for academics working in this area. Today journals exist dedicated to a complexity science approach. In the United States and UK several major universities have set up research centres in 'Complexity Science' and are now producing MSc and PhD students

specialised in its techniques and topics. Papers drawing upon complexity science have appeared in mainstream journals in a variety of more traditional subject areas, with biology and statistical physics the greatest beneficiaries. However, as this book will show, social sciences can use its concepts as well.

The Variety of ‘Complexities’

Both SFI-style complexity science and the fields invoking similar concepts have used the terms ‘complex’ and ‘complexity’ in diverse ways. SFI’s Seth Lloyd published a list of dozens of different definitions and measures of ‘complexity’ (Lloyd, 2001), but he identified three emerging themes: (1) ‘difficulty of description’ (typically measured in bits), (2) ‘difficulty of creation’ (measured in time, money, energy, etc.) and (3) ‘degree of organisation’, subdivided into (3a) ‘difficulty of describing organisational structure, whether corporate, chemical, cellular, etc.’ (‘effective complexity’), and (3b) ‘amount of information shared between the parts of a system as the result of this organisational structure’ (‘mutual information’). We shall not attempt to rival this list, but instead pick out some of the main complexity concepts relevant to a book on simulating innovation.

The first point has been mentioned already: complex systems consist of multiple parts or agents. For example, many people may be involved in the production of some innovation, which then diffuses among many others. Innovations themselves may be composed of multiple components. Because of the multiple parts, recording the state of this system may take many bits of information – part of Lloyd’s ‘difficulty of description’ theme. Each part may have multiple attributes, and each attribute may take multiple values. For example, if there are n agents in some population, and each agent has F cultural features or dimensions within which they can differ from each other, and q cultural traits within each dimension, then there are q^n different states of the system. If the population n increases linearly, the number of system states goes up exponentially. For quite modest values of population size, number of attributes and number of attribute values this can produce dauntingly large numbers. This is *combinatorial complexity*, the number of different ways of combining things.

A lot of problem solving can be thought of as involving a search within a vast space of possible combinations for the optimal solution, or at least a satisfactory one (combinatorial optimisation). In some cases (though not all), following a rule of thumb during this search – Simon’s heuristic search methods – may be an efficient way to get from one solution to a much better one in a relatively short number of steps, which is much more appealing than trying out every possible combination in turn. If it

is necessary to try nearly every combination, this is a sign of high *computational complexity*, and introduces Lloyd's 'difficulty of creation' theme. Our collective ability to solve combinatorial problems during our limited lifetimes, such as innovating in our individual attributes in order to minimise conflicts with each other and maximise mutual benefits, depends upon the difficulty of this task and the particular methods in use for searching through the space of possibilities.

Mainstream economics has tended to focus on collections of *homogeneous* agents, identical in their attributes and behaviour and therefore easier to represent in mathematical calculations, perhaps as some idealised Mr Average. This neglects the fact that real people and firms are *heterogeneous*, varied or diverse, in their attributes, not least in their spatial locations. In addition, the attribution of a particular attribute value to an agent may be a simplification as well. People can vary over time, often for reasons not perceivable to them or us. Random variation over time due to chance, or *stochasticity*, introduces more diversity into a system. Sometimes it does so in well-behaved ways, ways that are regular enough to show an identifiable pattern if we collect data, plot frequency distributions and analyse statistics. These then become knowable unknowns, predictable uncertainties, manageable risks. But when is it safe to assume we have now identified the correct pattern? And how do we know the pattern observed in the past will continue in the future? Even when economists recognise the possibility of random variation, there is a temptation to assume manageability of risk, because then they can employ mathematical techniques to deliver explanations for what has happened and advocate new policies for the future. It can feel reassuring to hear someone claim the world's complexity can be tamed, but they can still be wrong.

What makes the behaviour of a complex system so difficult to describe and predict is the interdependencies between the component parts. The systems are adaptive; parts change their states or attribute values. Reasons for state changes include manipulation by some agent, such as a human designer or manager, chance mutation and cross-over, such as when genes are reproduced, and random fluctuation or noise, such as happens due to heat. When changes occur, however, the new state is not purely a question of chance. The probabilities of different states being adopted by one part are determined by the states of other parts. Once a part has changed, its state may then affect the probabilities of other parts changing. In chemistry, the presence of molecules of one compound can catalyse, or raise the chances of reactions between other molecules. In technology the capabilities of one component are enhanced or inhibited by the other components, and also by how people use them. In ecosystems a population belonging to one species can affect the species it preys on and the predator species

that prey on it, as well as competitor species. Evolutionary changes in one species' behaviour and population can affect the fitness value of variations in other species, in a complex process of co-evolution.

In social groups the presence of one person can inspire or constrain the behaviour of other members, including affecting whether they continue to remain in the group. In a person's life, one behavioural practice can affect the performance of other practices, such as by learning skills that can then be transferred to other applications, and by reducing the amount of time and resources one has available for other activities. People, their beliefs and ideas, their practices, their props and tools, the places that serve as venues for them, and their social networks or knowledge of other people can all be intertwined in networks of interdependencies, catalysis and constraints. The number of such interdependency relations gives us another aspect of complexity. However, what happens in networks is not just the result of the number of links, but also other features of its structure (a point we will demonstrate in several models, and especially in Chapter 3). One important feature is the presence of loops, whereby a change in one component can initiate a chain of changes in other components, eventually affecting the original component itself. Such loops are known variously as self-referring, self-reflexivity (Popper, 1957; Soros, 1988), virtuous and vicious circles, self-fulfilling and self-refuting prophecies (Merton, 1968b), and positive and negative feedback loops (Arthur, 1994).

For example, suppose a person whose opinion is sought – a government finance minister, perhaps, or a technology consultant – makes a prediction that some tradable item, for example a company stock, will increase in price over the next year. Those who believe this person will decide they can make money by buying this stock now and selling in a year's time. Others may realise that one effect of this first group buying the stock will be to increase its price in response to the increase in demand. So even if they did not trust the first person predicting its rise, they may still believe the price will go up, and therefore that they too can make money by buying it now and selling later. This additional increase in demand again increases the price. Thus the first person appears to have made a self-fulfilling prophecy. But what if the original opinion was that the price would peak in a year's time? Those acting on such a belief will want to sell before the stock goes into a decline. They may therefore seek to sell before or at the predicted time of the peak. When they do so, this drop in demand will have an impact on the price, sending it downwards. Thus the prophecy of a peak is self-fulfilled. However, if people sell before the predicted time of the peak, it will go down early. Thus the prediction of when the peak would occur is self-refuted. Real-world trading can be more complicated than this. For instance, a firm's share price can affect its ability to raise new capital and

its ability to grow and thus justify an increased share price, a form of virtuous circle that billionaire George Soros attributes his success to (Soros, 1988, Introduction). Financial derivatives, futures and options allow hedge funds to bet on, and try to make money from, stocks going down in value. The presence of so many self-reinforcing and self-refuting actions, or positive and negative feedback loops, is what makes stock markets and economies so hard to predict.

It is because of these interdependency relations that the behaviour of complex systems is not a sum of the behaviour of their parts. Adding one extra person to the group may increase or decrease its total productivity. Adding a catalyst to some chemicals may greatly alter the relative levels of different compounds, and also alter whether they are solid, liquid or gas. Introducing a foreign species to an island's eco-system may have a devastating impact on other species' numbers, or have no discernable impact at all. Similarly, introducing an innovative technology may initiate Schumpeter's 'gales of change', rendering multiple technologies and practices obsolete while creating market opportunities for others, or it may have little or no economic impact.

The uncertainty about such waves of changes may resemble that about certain natural phenomena, part of the study of *self-organised criticality* (Bak, 1997). For example, earthquakes: there is no such thing as an average earthquake. The frequency distribution for earthquakes is 'scale-free'; it tends towards a straight line when plotted with logarithmic scales on both axes, and mathematically takes the form of a power law: $y = A x^b$. Power laws and scale-free distributions are among the signatures of complex adaptive systems.

But studies of complex adaptive systems are not always about cascades of changes. Stable patterns among the component and system states may *emerge*. For example, if one combination of component states is superior to all variations, it may come to dominate. A genome that is fitter than its rivals, or better at replicating itself in the current environment, will tend to increase its relative population size. A technology that is more useful and valuable to consumers than its competing substitutes will tend to increase market share. A solution to an organisation's problems, a strategy or combination of actions that seems more profitable to the organisation's members than its alternatives will tend to be adopted by the members, as a consensus view emerges. Such emergent patterns and order from heterogeneous, adapting parts are the other well-known system-level behaviour of complex systems.

Exactly which pattern or combination of component states emerges may be unpredictable, the result of building on and reinforcing chance, micro-level events. There may be several emergent patterns possible from

a particular complex system, several ‘attractors’ in its system state space, several ‘peaks’ in its ‘landscape’. But that some sort of pattern or order will emerge may be predictable, as may the expected number of possible emergent patterns. Chaos theory (Gleick, 1987) showed how simple, deterministic functions could have complicated behaviour, impossible to summarise or predict without performing all the calculations. Complex adaptive systems, of simple or of complicated parts, can have relatively simple or relatively complicated behaviour – often a relatively simple change to the parts’ behaviour shifts the system between the two.

Clearly, if humans live in complex adaptive systems, there are implications for decision makers and planners. Organisations are both complex systems and operate within them. The study of complexity and emergence is commonly associated with *bottom-up explanations*: that is explaining the system-level behaviour by reference to its parts’ behaviour. This is not to be confused with the ‘laissez-faire’ approach to management, whereby leaders give their employees or citizens the freedom to act without interference from upper management, and layers of corporate hierarchies are dispensed with, in the belief that the people at the bottom will self-organise and good collective behaviour will emerge from their activities (Goldstein, Hazy and Lichtenstein, 2010, p. 4). Leaders are also components in the complex system that is the organisation; a bottom-up explanation must include them. The self-organising system includes people at all levels having to adapt to each other. If complexity science has a message for management it is more likely to be about how uncertain are the consequences of stripping away a level of organisational structure or altering the behaviour of a key component. There is no guarantee that such attempts at managerial innovation will be beneficial in their consequences. The emergent system state may be one with no organisation at all.

To sum up what is meant by complexity, there are four components: diversity, interdependencies, adaptation and emergent order. In the case of innovation, diversity means heterogeneous agents, variety of parts and stochastic variability. Interdependencies mean networks of influence, catalysts and constraints, and positive and negative feedback loops. Adaptation includes trial-and-error experimentation, organisational problem solving and social learning, natural selection, learning by doing and market pressure, all of which can lead to some changes in states being more likely than others. The existence of stable system states is determined by the presence of interdependencies between the combinatorially complex system states, a mathematical property of the complex systems that Kauffman has called ‘order for free’ (Kauffman, 1996). But if they are there, they can emerge, searched for by the system itself, even though none of its parts knows of them. Social–economic–technical systems collectively

self-organise and search for more stable states, via relatively simple adaptive processes. Innovations emerge from the dynamics of such complex adaptive systems. This is why the study of complexity is relevant to the study of innovation. But how best to study complexity?

Research Approaches and Complex Systems

Studying innovation means studying complex adaptive social systems. Which research approaches from the social sciences will be appropriate for this? There are two familiar types of social science. The *quantitative* approach primarily involves statistical analyses of quantitative data obtained by questionnaire surveys or other numerical counting and measuring processes. A qualitative approach primarily involves constructing interpretations by reference to written and oral accounts by interviewed participants in the system and/or accounts by researchers as observers of the system, especially those observers who have participated themselves in the systems. We shall consider the quantitative approach first.

Given quantitative data one can try to propose a model, a mathematical description of some set of relations between different attributes of the items in the dataset, that *explains* the data, that is, explains in terms of some attributes why certain other attributes have the values they have. Fitting involves finding the parameters of a model that minimise the extent to which the variables to be explained deviate from what is predicted by the model's transformation of the explaining variables. The most familiar type of model is linear regression. As the name implies, this is unsuitable for systems that are nonlinear, where one component's behaviour is definitely not expected to be a linear sum of the behaviours of the other components. Regression techniques are possible with nonlinear models, but these are not widely used. There are indefinitely many mathematical functions that could be tried out, and why choose one model over another? Given enough parameters, a model can be constructed that exactly fits a dataset, but may be useless in telling us why the parameters are as they are, and unreliable in telling us how the system will behave in the future, beyond the current dataset's coverage. These are the problems of *overfitting* the data. Regression models only tell us that some variables are correlated. They do not tell us how the association has come to be. Hence, they provide no warning if the association is about to disappear (raising the philosopher's problem of induction: how can one know that a past pattern, such as the sun rising in the morning, will be repeated in the future). For that, information is needed about the *mechanisms* generating the observed pattern. As studies of chaotic functions showed earlier, the initial behaviour of some kinds of system may give little warning of its

later behaviour. Small differences in initial conditions may quickly lead to large differences in outcomes – the famous butterfly effect.

In addition to these problems, many datasets cover not the entire system of interest, but only sample some of its components, or sample a dynamic system at only a few points in time. If the sampling attempts are independent of each other and concern a component whose behaviour is distributed identically each time, then statistics textbooks tell of various techniques that can be used to infer things about the system underlying the sampled data. If, on the other hand, it is plausible that the items in the sample are not independent and identically distributed, then most textbooks are less forthcoming. Also, some techniques, for example analysis of variance (ANOVA), assume that the data are being sampled from normally distributed processes. Normal distributions may apply for some phenomena that have a characteristic scale (for example, heights and IQ measures); learning the average value is of interest. But normal distributions are rarely found in empirical economic data or in innovation studies (for example, GDP, income, age and innovation generation). As mentioned already, the behaviour of some complex adaptive systems is associated with *scale-free* distributions. For these, the average value is of no use, and the variance and standard error, used for statistical tests, may be undefined.

Even when the data sources are suitable for statistical testing, researchers often draw their conclusions from tests of *statistical significance*. This is not equivalent to measures of *importance* (McCloskey, 1985; McCloskey and Ziliak, 1996). By basing decisions on statistical significance, researchers can miss important causal relations due to noisy data, while at the same time publishing claims about causal relations which if true may still be very weak. Epidemiology is another field that suffers from this problem (McCloskey and Ziliak, 2009). Given how complex the human body is, and how varied the things we do with it, one should be particularly careful in making statements about the causal effects of one lifestyle factor or about the factors behind some disease.

When studying the emergence of innovations most quantitative methods may not be appropriate anyway. Mostly they work with averages from data, but a novelty in the data is more likely to appear as an extreme case, an outlier, far from what has been typical behaviour so far. Also, when collecting data one chooses metrics and designs questions for surveys on the basis of what has been considered useful to ask in the past. Novel, emergent order, not resembling previously observed patterns, may therefore be neglected by these methods. If we have not seen it before, we might not know to ask about it.

Qualitative research methods may be able to overcome this problem. If the researcher is a part of the system, or talks to those who are, when an

innovative event occurs, for example when a novel problem is encountered or a new idea is developed, the researcher may be able to trace its genesis. Human natural language is much richer than the subset employed in a quantitative survey. Even when no word exists for what is developing, metaphors and analogies may serve to construct its meaning. However, one's ability to recognise the importance of the new event depends on knowledge of other cases to compare and contrast it with. Indeed, this relatively high dependence on the subjective background of the researcher can lead many people, quantitative researchers especially, to disregard qualitative research. Also, compared to questionnaire surveys, the richer experience obtained by qualitative research can only be obtained at a slower rate and comparisons between cases take longer and are disputed more often. Extrapolating from some observed cases to a yet-to-be-unobserved case may invoke controversy. Human lives are so complex; trying to abstract and generalise from particular cases seems to threaten their complexity and autonomy. So scaling up qualitative studies is rarely performed unchallenged.

But without large-scale studies, how can we be sure that an innovation, the emergence of the next big technology, say, will be captured in detail by a researcher? With too few qualitative researchers, the chance of a researcher being in the right place at the right time seems slight. A researcher can perform a retrospective study, examining documents from the time when what we now know to have been an important technology was developed. But some documents may be missing – why keep them if you do not know how important this technology is going to be? And like war stories, tales of technological development tend to be histories of the victors.

So quantitative research, such as questionnaire surveys, can achieve large scale and the models commonly employed are easy to reapply to new cases. But the models are relatively simple and data collection is insensitive to emergence and life's complexity. Good models – models that fit the data – demonstrate correlations between variables. They do not demonstrate causal relations, or offer guidance as to what mechanisms generated the patterns, and hence provide no indication about whether the patterns of behaviour are set to continue. Qualitative research, such as analysis and critical reflection on oral and written accounts obtained from interviews and participant observation, can match the richness of social experience with the richness of natural language. But aggregating qualitative case findings and extrapolating to new ones is controversial and frequently disputed. Is there a middle way or a third option? One approach would be to try to combine qualitative and quantitative data collection. Advocates of *mixed methods* propose to do just that (Creswell and Plano Clark, 2007).

We will not explore that option in this book. Instead we will turn to the subject matter of the next section, computer simulation models.

Social simulation is the paradigm method in analytical sociology (Hedström and Bearman, 2009), whereby macro-level social facts, such as the emergence and diffusion of an innovation, are explained in terms of the mechanisms, mostly human actions and interactions, by which the social facts were generated. Pattern-oriented modelling (Grimm et al., 2005), developed in ecology, follows a similar approach. The intended contrast is with explanations of social facts in terms of their relations to other social facts, whether by statistical association or by logical deduction from grand theoretical assumptions. Analytical sociology thus bridges the gap between micro-level processes and macro-level phenomena. Economists also offered a micro-macro bridge, using mathematical integration over individuals' decisions. But it was based on homogeneous individual agents with unrealistic cognitive abilities and information, and the assumption that collectively these agents would form a system at equilibrium, so one had only to start with that state. A more responsible approach to studying complex adaptive social systems must respect our knowledge about the limits to human rationality and information sources, the diversity in human attributes and environments, and the dynamics of collective behaviour.

SIMULATING INNOVATION IN COMPLEX SYSTEMS

Why Use Computer Simulation?

The intended purpose of simulations of real-world systems is to give us something useful that we could not – for a variety of reasons – obtain from the system itself (Ahrweiler and Gilbert, 2005). An aircraft simulator, for example, gives would-be pilots an experience analogous to that of flying a plane, without the risks and costs of practising on the real thing, and with simulated, hypothetical situations that might not occur very often in real life, such as engine failures in bad weather. A key role for simulations is to answer what-if type questions, that is, simulations are not limited to the representation of the real world. Simulations of social phenomena can also save costs and avoid risks, not least the risks to professional ethics implied by experimenting on real people. Policy makers wishing to think through the consequences of their actions before they make them may appreciate being able to experience a simulation of their implementation first.

It might be wondered, however, whether computer-based representations of people can provide an adequate analogy to the real thing. Human

beings can certainly be very complicated in their behaviour, and often quite mystifying. This does not stop other human beings trying to predict their actions. Indeed, humans seem to be particularly good at interpreting other humans. These attempts often fail, of course, but the failures have not been so great that we have chosen to cease the effort of making them. Part of the skill in interpreting others lies in focusing on some attributes of the person while neglecting others. Likewise, the art of modelling people requires that we leave something out. Constructing a model of a person will not make you that person, though it may lead to you performing similar actions or making similar judgments to that person and it may inspire in you the same response as you would take to the real person.

Our ability to work with mental models of people and their interactions actually goes some way to explaining how computer models of the same complex systems come to be useful. Ashby's principle that variety must match variety (Ashby, 1958) might seem to pose a problem when trying to model a complex system. To understand how simple computer models can be adequate to the task of modelling complex minds or social systems, it must be remembered that the computer models in this book are not interacting with the real-world system automatically, but only as part of modelling projects, designed, run and interpreted by human modellers. It is the combination of modeller plus model that has to meet the complexity of the real-world system (Pidd, 1999), not models in isolation, and human beings, as noted already, are particularly adapted to responding to other human beings and interpreting social situations.

Agent-based Simulation Modelling

There are various approaches to simulation, and in the next chapter we will illustrate some of them. *System dynamics* uses difference equations to represent stocks and flows (Sterman, 2000). *Discrete-event simulation* processes lists of events, with inter-event time periods determined by random sampling from particular probability distributions (Law, 2006; Robinson, 2004). In this book, however, nearly every model is an *agent-based simulation*, also known as individual-based simulation and multi-agent simulation (Axelrod, 1997a; Gilbert, 2008; Gilbert and Troitzsch, 2005).

The agent-based simulation approach explicitly represents individuals with particular attributes engaging in interactions with each other and with a shared environment. The agents are most often intended to represent people, but agents representing animals, inanimate objects, firms and countries are found within the modelling literature, and often a model will include more than one type of agent. Time steps are also explicitly represented. The attributes of the agents at one time step are determined

by their attributes at previous time steps. At a time step, agent interactions occur according to some relatively simple rules of behaviour, usually represented as a few lines of computer code, for example, the rules of thumb, or heuristics, that Simon claimed were employed by human decision makers (Simon, 1955a, 1991). Agents may differ from each other in their current attributes as well as their rules of behaviour. In particular, there may be constraints about which other agents a given agent is capable of interacting with, that is, their social network. Agent behaviour may also vary according to some random elements, or stochastic processes. As a result of interactions the participants' attributes may change.

When attributes are represented visually, for example, as x and y coordinates, or as colours, a simulation user can look for on-screen visual patterns among the population of interacting agents, such as crowd formation. There will be a number of parameters to the model, controlling such features as the number of agents, the agents' initial attributes, their behaviour and their environment, and a human user may be able to learn about the model's behaviour by altering these parameters during a simulation run or between runs, using on-screen controls. Statistics can also be collected to summarise the population at a particular step of the simulation, across multiple steps during a simulation run, or across multiple runs. Such data can then be turned into charts or used in statistical tests. Thus, besides agents, agent-based modelling offers us rules of behaviour, heterogeneous agent attributes, networks, random variability, visualisation and user interaction, emergent patterns and the ability to experiment with an abstract, model system. No other simulation approach offers all this in so convenient a form. Each of the agent-based models in the following chapters will employ some or all of these features.

By comparing the output from simulation runs with varied parameter settings, the user can perform a computer simulation experiment. Given the focus of this book, we shall shorten 'computer simulation experiment' to 'experiment'. This is not an experiment on the real-world system that the model is intended to represent, but rather an exploration, made with scientific rigour, of an abstract system, albeit one which may provide some analogy to the real-world. Since experimenting with real people is often impractical, for reasons of cost, danger, ethics or lack of participants, the simulation model may be the only option we have for producing useful answers to what-if questions about the real world (Ahrweiler and Gilbert, 2005).

In other research fields and the business world, common applications for computer simulation models are fitting past data and forecasting future or hypothetical events, the most common uses for statistical models. Simulations may also be used for finding practices or quantities

that produce supposedly optimal outcomes from some system, in effect a form of prediction. As will become clear, some of the social simulations described in this book have a rather different purpose. Indeed, we give demonstrations of why forecasting innovation diffusion is unlikely to succeed, because of both random variability and the complexity of the would-be adopters' contexts. Instead, our aim is often the facilitation of understanding. Often the biggest problem for the members of an organisation is not solving some problems, but rather knowing what problem they face (Rosenhead and Mingers, 2001, Introduction). Once problems have been collectively identified, structured and agreed upon, the methods for solving the problems may be straightforward. Methods for problem structuring can still be rigorous and grounded in scientific research, such as social psychology and cognitive science, and simulation modelling can be included among these methods (Robinson, 2001).

Because of this different purpose, when discussing a social simulation model, there is less emphasis on validating the model by fitting it to some historical dataset, and greater reflection on conceptual modelling, that is, the question of which concepts from the real-world system should be included in the model and which left out (Robinson, 2008a, b). Indeed, sometimes the model might not even need to be completed, that is, debugged and run, in order for participants in a modelling project to feel that their understanding of the real-world system has improved.

Compared to statisticians' models, agent-based models are likely to have many more parameters. For a statistician, this would seem to reduce their power to explain anything, since with regression models the more parameters one has when fitting data, the easier fitting becomes, and therefore the less informative it is. However, the many options and parameters in an agent-based model help make more explicit the assumptions behind the model, and allow users to focus on and experiment with the model features they think most important, rather than the features the modeller thought important at the time of programming. By including alternative functionality and optional features, the modeller aims to avoid excluding any participants in a discussion of what one can learn from the model.

Research with Simulation

The idea of using computers to simulate business processes was being written about by the early 1960s (Cyert et al., 1964; Tocher, 1963), although some of the pioneering examples of cellular automata, a form of simulated system of multiple interacting individuals, Conway's game of 'life' (Gardner, 1970) and Schelling's segregation model (Schelling, 1969, 1971) were developed initially using such non-electronic technologies as

sketches in the margins of a newspaper, nickels and dimes on checkers boards and floor tiles. The growing availability of personal computers during the 1980s, and the ability to program them oneself, together with exponential rates of improvement in computer speed and storage capacity and more detailed and friendly GUIs, meant that sophisticated computer simulations could be made available to all. By the early 1990s there began to appear computer-programming languages with built-in functions developed specifically for the simulation of interacting individuals. Also during the 1990s two academic journals were launched specialising in the simulation of social phenomena, namely *Computational and Mathematical Organization Theory* in 1995 and the *Journal of Artificial Societies and Social Simulation* in 1997, since when both journals have continued to flourish (Meyer, Lorscheid and Troitzsch, 2009; Meyer, Zaggel and Carley, 2011) and more journals have been launched. Papers based on simulation modelling have also been published in mainstream journals in various fields, including sociology, psychology, environmental studies, geography, economics and business studies.

AN OUTLINE OF THE BOOK

The best-known use for models in innovation studies is that of modelling the diffusion of innovations. So this is addressed first, in *Chapter 2: The variability and variety of diffusion models*, where, following Geroski (2000), we survey several different types of diffusion model: *epidemic*, *probit*, *stock* and *evolutionary* models. Each modelling approach focuses attention on different reasons for diffusion, including: imitation of neighbours, personal responses to a changing environment and responding to the level of adoption in the population itself: a form of feedback loop. In addition, diffusion modelling provides an opportunity to compare and contrast three approaches to computer simulation: system dynamics modelling, which uses difference equations to represent stocks and flows, discrete-event simulation, which represents the occurrence of random events, and agent-based simulation, which represents individual agents, such as people, who interact with each other and an environment. Different simulation models focus one's attention on different aspect of the world one is trying to model, and can lead to different thoughts about how best to act in that world. We choose agent-based simulation over the other approaches, since it makes it easy to represent heterogeneous decision makers, as seen in the probit model, who can interact with each other according to some social network, as seen in the epidemic model, and who, potentially, could base decisions on aggregate properties, as seen in the stock model. Agent-based

simulation also allows the representation of random variation, or stochastic processes, and we provide a demonstration of why it is important to include this in diffusion models. Random variability in real-world cases of diffusion, however, makes it hard to use models and past data on how many people have adopted some innovation to predict reliably how innovation adoption will behave in future.

Focusing on epidemic-type diffusion models, *Chapter 3: Diffusion and path dependence in a social network* looks at how restricting social interactions to a fixed network of relations affects the spread of innovations. Actual networks could be the reporting relations in an organisation, friendship and family ties or relations of geographical proximity. *Social network analysis* has become a familiar social-science tool in recent decades, helped by electronic means to collect data on who interacts with whom, and by the availability of PCs to calculate statistical metrics from those data. Particular interest has been paid to the structural properties of networks, including how many links there are between members of a network, how far information about an innovation has to travel between any two people, and how often a neighbour of a neighbour of oneself is also a neighbour of oneself. The first two properties affect how quickly information about an innovation travels through a networked population. The third property affects information travel whenever would-be adopters want more than one of their neighbours to have adopted before they themselves will adopt. Different networks have different structures and present these properties to varying degrees. Hence the networks have varying effects on diffusion.

Less well known is the relation between network structure and what is known as *path dependence*. This is the property of diffusion whereby adoption decisions at one time affect the chances of later adoption, particularly important when multiple innovations are spreading through the same network and competing for adopters. Our network model shows that the most likely outcome between two competing innovations varies with network structure. Another diffusion model is the *information cascades model* (Bikhchandani, Hirshleifer and Welch, 1992, 1998), in which decision makers use previous adoption decisions by others as evidence for or against their own decisions, as well as having their own, imperfect private source of information about the innovation. Under some circumstances this can lead to cascades of similar adoption decisions, when a group of decision makers start to follow an emerging consensus. In this way, the model offers an explanation of the existence of fads and fashions, and herd behaviour in crowds. But the possibility of fads undermines the utility of learning from others' adoption decisions. Therefore, decision makers should recognise when actions may be the result of herd behaviour, and

discount these in their own decisions. Under such circumstances, only actions that buck the trend are likely to be counted as evidence, for they must be the result of factors other than just following the crowd. These factors could be previous decision makers following private sources of information, though they could also include mistakes and noise. Consequently, mavericks who perform surprising actions wield influence over later adoption decisions, but only while there exist at least some decision makers who are prepared to be influenced by surprising behaviour: not everyone can be a maverick all the time; a balance must be sought. This balance turns out to interact with social network structure. Applying a social network structure to a population of decision makers both reduces the threat of fads and enables them to track the current value of adopting the innovation.

Collective learning is also a feature of the models in *Chapter 4: Explore and exploit*. This begins with the view of humans in organisations proposed by Herbert Simon and colleagues in the so-called ‘Carnegie School’. Humans are bounded rational decision makers, engaging in routine practices most of the time and employing rule-of-thumb search routines, known as heuristics, whenever problems call for a new combination of routine practices. James March (1991) demonstrated with an agent-based model that collectively an organisation could solve problems and learn in situations where an individual could not (Rodan, 2005). Key to this, however, was a balance in the organisation’s learning practices between *exploration* of new candidate solutions, and *exploitation* of those already found. Explore too little, and you may become stuck with consensus around an inferior combination of practices, the problem of *premature convergence* (Levitt and March, 1988). Explore too much, however, and you fail to benefit from the knowledge already acquired. A later agent-based model by Lazer and Friedman (2007) builds upon the idea of this balance between exploration and exploitation, and shows that if a social network constrains who can learn from whom in the organisation, then the social network structure interacts with the search practices to determine this balance. March (1991) also points out that when in competition with other organisations, changes that raise an organisation’s average or *expected performance* are sometimes inferior to changes that raise its *variability in performance* at the expense of the average. This occurs when *relative advantage* in actual performance, i.e. who came first, counts for more than *absolute advantage*, i.e. by how much they came first. This leads to a distinction between biologists and economists over what constitutes ‘rational behaviour’ (Slobodkin and Rapoport, 1974; Thorngate and Tavakoli, 2005), and has become increasingly important as more and more aspects of our society move towards rewarding people and businesses on

a winner-take-all basis (Frank and Cook, 2010). A simple simulation of a betting game illustrates when the economists' ideal of rationality (maximising your expected utility) might not be your best course of action if your very survival is at stake (Thorngate and Tavakoli, 2005). If it involves some level of risk and cost to oneself to engage in exploring new solutions rather than exploiting existing ones, then the rewards for innovation, absolute or relative, collective or individual, become crucial.

In *Chapter 5: Science models*, some data on innovation point to these rewards being cumulative. This is one of several examples in this chapter of a relatively simple mechanism explaining a pattern observed in empirical data. Analysis of data on academic publications shows evidence of opportunities for new publication tending to go to those authors who have already published, that is a case of 'the rich getting richer'. Likewise, citations of past publications, often taken as an indicator of the quality of the cited publication, tend to reference those publications that are already rich in citations. The mark of such cumulative advantage is a *scale-free* or *power-law distribution* of the frequency of occurrence of publications per author or citations per paper. A simple simulation model can generate data that approximate such distributions, given a particular balance between processes of innovation and imitation when choosing authors or citations for new publications. Another pattern discernible in publication data is network clustering, especially clusters of authors linked by having co-authored papers together, clusters of papers linked by common keywords or topics, and clusters of papers linked by citing the same past papers. Such network clustering reflects the academic fields and subfields that give content to a publication, but is also influenced by social processes. In particular, processes that show forms of *homophily*, or the preference for similar others, can lead to clustering in networks of social interactions, especially when combined with processes of social influence, as agent-based models make clear (Axelrod, 1997b; Gilbert, 1997; Hegselmann and Krause, 2002). Science models, that is, simulations of academic production, can combine these processes to produce synthetic publication data with some of the properties observed in real data, including scale-free distributions and clustering. Could these models help inform policy concerning the organisation of real academics? The models of organisational learning in Chapter 4 were examples of simulations of knowledge creation. The processes of problem solving by heuristic search can easily be added to a science model, but calibrating it so that it still generates plausible-looking publication data is more difficult (Watts and Gilbert, 2011). An aim for such a model, however, would be to address the question of whether processes of cumulative advantage and homophily among academic authors enhances or inhibits the pace of knowledge creation.

For the science models in Chapter 5, processes were sought that would generate patterns seen in quantitative data on publications. What if we started with observations of scientists actually at work? What processes, interactions and contextual factors could we identify that might be represented in a simulation? *Chapter 6: Adopting and adapting – innovation diffusion in complex contexts* uses the literature on actor–network theory (ANT) and the social construction of technology (SCOT). Qualitative data from ethnographic studies of scientists, engineers and others at work portray innovation as a process of satisfaction of highly diverse sets of constraints. Innovative combinations of components, practices and contexts involve trade offs between material costs and physical constraints, but also the different social, political and economic interests of the parties involved. As mentioned earlier, the traditional representation of innovation taught in business schools, the so-called linear model of innovation, divides innovation into stages of first invention or the introduction of the innovation, and later diffusion of that innovation among the population. The properties of the innovation are considered fixed at its introduction. In contrast, the view of innovation to be gleaned from the ethnographic studies is that different people see different things in the innovative technology, practice or project, and they evaluate it differently. Each case of adopting an innovation involves adapting it to the new, unique context. The key to creating a successful innovation, that is an innovation that diffuses widely, seems to involve making it easy to adapt it to heterogeneous contexts. With this in mind, simulations of innovation diffusion should resemble processes of complex constraint satisfaction, rather than the simple epidemics modelled in Chapter 2. An example model is described that expresses in visual form the heterogeneity of adoption contexts. The resulting adoption patterns, however, are far from the adoption curves familiar from the literature and discussed in Chapter 2. Instead, there are interdependencies between the different components of the adoption event, and the order in which these components appear makes a difference to the outcome, that is, the simulated system shows path dependence, a concept already mentioned in Chapter 3.

Chapter 7: Technological evolution and innovation networks continues the theme of complex networks of interdependencies. Simulations of technological evolution and knowledge creation are surveyed. Various analogues are given for the creation of new technologies or knowledge, including percolation on a grid network (Silverberg and Verspagen, 2005, 2007), search for good designs of logic functions (Arthur and Polak, 2006), co-evolution, or parallel searches using genetic algorithms, of good game-playing strategies, and a search for *auto-catalytic*, or self-producing, networks of production rules, inspired by artificial chemistry and the origins

of life (Padgett, Lee and Collier, 2003). The patterns generated by these models include scale-free distributions of innovation size, and the structural properties of the networks that emerge among innovation-producing firms. The models offer explanations for various empirical facts about innovation production and its relation to social organisation. Indeed, they may even lead to explanations of the emergence of new forms of social organisation itself, within which novel technologies, roles and practices are components.

The models of Chapter 7 bring together most of the principles illustrated by those described in the previous chapters. Chapter 8 concludes the book with a recap of these principles, namely stochasticity or random variability, epidemics, heterogeneous agents in changing contexts, social network structure, path-dependent outcomes, herd behaviour and the power of surprising actions, heuristic search and collective learning, cumulative advantage, homophily, innovation as constraint satisfaction, adoption as adaptation, networks of interdependent parts, co-evolution, auto-catalysis, and the emergence of innovation networks. On the basis of empirical evidence from quantitative and qualitative studies, past theorising about innovation has striven to include some of these. But rigorously incorporating them all is a task that perhaps only agent-based modelling can achieve. Employing other modelling techniques means omitting important components of the concept of innovation, and thereby weakens our power to think about innovation. It is the key aim of this book to widen awareness of the tools for thinking available to researchers working today. The global financial and economic crisis that has run during our work for this book is unlikely to be the last time innovation will have a major impact on our lives.