

Chapter 2

Theory

I. Nikolic and J. Kasmire

Abstract This chapter introduces and explains the main concepts that provide the theoretical background on how to model the ubiquitous socio-technical systems that are so important to modern life. First the notions of systems, adaptation and complexity are discussed as individual concepts before addressing complex adaptive systems as a whole. This is followed by a discussion on generative science and agent-based modelling, with special attention paid to how these concepts relate to socio-technical systems. Throughout the text examples of how the theories can be applied to real systems are provided. Armed with a solid understanding of concepts such as observer-dependence, evolution, intractability, emergence and self-organisation, the reader will have the right foundation for moving on to the practical aspects of building and using agent-based models for decision support in socio-technical systems.

2.1 Introduction

As discussed in Chap. 1, this book is about creating agent-based models to improve our decision making in and around socio-technical systems.¹ Agent-based perspectives, on both analysis and decision support, imply a complex adaptive systems approach to systems and modelling, and this chapter lays the necessary theoretical foundations for the concepts and tools needed to create such models.

The main theoretical perspectives used in this chapter, and throughout the book, are complexity, complex adaptive systems and the generative science paradigm. These perspectives can be considered as meta-theories, which could apply to just

¹These systems have also been called large-scale socio-technical systems, complex socio-technical systems (Bonen 1981), socio-technical systems (Geels 2004), large technical systems (Bijker et al. 1987), complex innovation systems (Katz 2006), complex engineering systems (Ottens et al. 2006) and even the impressive sounding “system of systems” (DeLaurentis and Crossley 2005). For a detailed discussion on various application fields and uses of complex adaptive systems, please refer to the work of van der Lei et al. (2009).

I. Nikolic (✉)
<http://www.igornicolic.com>
e-mail: i.nikolic@tudelft.nl

about anything, from ecosystems, multinationals, the Internet, the legal system, earthquakes, slang, ant colonies, the brain, immune systems and the pattern of tagging in social networks. However, we use them to explore the domain of socio-technical systems because these systems abound in human society, are fascinatingly complex and unpredictable, are important to the development and welfare of humanity and life on earth, and have so far defied attempts at analysis using other approaches. This is not to say that these are the only valid tools to analyze these systems, or that these approaches are not valid for investigating other systems, but those investigations will not be found in this book.

2.1.1 Focus

Socio-technical systems, as a class of complex adaptive systems (Kay 2002), consist of many technical artifacts (machines, factories, pipelines, wires, etc.) and social entities (individuals, companies, governments, organisations, institutions, etc.). These are interwoven in networks of physical and social components, applying and responding to selection pressures and racing to adapt in a fast vast mega-coupled fitness landscape. Examples include single production plants, regional industrial clusters, interconnected power grids, multi-modal transport networks, telecommunication networks and global enterprises.

Industrial society as a whole is a collection of co-evolving, large-scale, social and technical systems, with multi-dimensional flows of physical matter and social information. Humans, as both physical and social beings, bridge the gaps between the physical networks, governed by laws of nature, and the social networks, with an added layer of social conventions, formal laws, rules and institutions. The social connections are as real as physical ones, varying in length and formality (everything from an interaction with a clerk in store to governmental regulations) (Williamson 1987). In this way, information, research, investment, consumption, institutions, rules, regulations, policies, habits and social values flow between people and physically affect how mass and energy are allocated and used, which in turn shapes the social systems that depend on those physical realities. For example, mixers, reactors, and heat exchangers of a chemical processing plant exchange (and conserve) mass and energy (Coulson and Richardson 1999), but through indirect connections to global supply networks, also influence the price of goods and labour, demand for products, and the stability of governments along the supply chain. Thus, system content, structure and boundaries shift and evolve without any global or central *coordinator*, order and regularity, instead emerging from widely distributed bottom-up interactions of subsystems, some with centralised control and others fully distributed.

2.1.2 Structure of the Chapter

We begin by discussing what systems are. Only after we understand what a system is can we move on to notions of adaptiveness, and how it relates to systems, be-

fore alighting on the heady topic of complexity and linking them all for a complex adaptive systems look at some real world examples. Finally, we wrap the chapter up with an introduction to the basics of generative science and agent-based modelling as a tool for a generative science approach to complex adaptive socio-technical systems.

2.1.3 Example: Westland Greenhouse Cluster

Throughout this chapter, we will use the Westland greenhouse cluster in the Netherlands as a running example of a socio-technical system in order to clarify the theoretical notions. Greenhouses represent a good example to work with as they have clear human and social components as well as technical components, yet both are entwined in such a way that there is little value in trying to examine either in isolation.

Modern greenhouses derive from various ancient techniques to alter the environment of plants so as to protect delicate species or to extend the fruiting or flowering season of popular plants, documented as early as fifth-century BC Greece (Hix 1996). These techniques included heating the soil or air around the plants, placing protective materials around the plants, or moving potted plants in and out of protected spaces to best catch sunlight while avoiding extreme conditions. These rudimentary techniques have since developed into sophisticated heating, lighting, aeration, and irrigation systems, as well as various configurations of protective walls and windows, and technologies as modern as computer controlled robotic systems to move or turn the plants, sow seeds or pick produce.

Before the Industrial Revolution, only aristocracy could afford the time, energy and resources for protected cultivation. As such, greenhouses were status symbols and owners competed to have the most diverse collections, best examples of rare species, or impressive new ways of operating the greenhouses (van den Muijzenberg 1980). This competition led to many advances in design and technique that spread and developed through academic horticultural societies, journals, books, as well as through the hiring of experienced gardeners and architects (Hix 1996). Following the Industrial Revolution, materials and technologies became available in quantities and prices that allowed commercial enterprises to incorporate some protected cultivation methods, which brought a new focus on quantity and production lacking in the non-commercial origins of greenhouses (van den Muijzenberg 1980). As a consequence of commercial competition, modern greenhouse horticulture businesses, like those in the Westland greenhouse cluster, specialise in a single product, or an extremely limited range of compatible products, with enormous investments in the tools, technologies and processes designed to optimise production (Hietbrink et al. 2008). This is a highly successful greenhouse area, one of the largest areas of greenhouses in the world, and has a highly technologically advanced processes that use large amounts of resources and contribute significantly to the GDP of the Netherlands. It has spawned equally specialised and high-tech transport, processing,

and packaging industries, as well as complicated markets, regulations and subsidy schemes.

2.2 Systems

“A System is a set of variables sufficiently isolated to stay discussable while we discuss it.”

W. Ross Ashby, cited in Ryan (2008)

We open the discussion on systems with Ashby’s quote, which introduces several important concepts, some of which are implied rather than stated. Systems are part of the larger whole. They are coherent entities over some period of time. They have boundaries. More subtly, this quote also suggests that a system is something an observer chooses to observe. All of these concepts, and a few more, will be discussed in this section as we explore the history of systems thinking, the contribution of the systems perspective, and elaborate tricky systems notions such as system boundaries, context and nestedness.

2.2.1 *History of Systems Thinking*

Prior to 1950s, the analytic, deterministic and mechanistic world view prevailed in science (Ryan 2008). People seemed comforted by the idea that everything was ticking over like clockwork, operating under immutable rules, and that given enough time, clever people would be able to measure all the gears, time the cycles, link up the actions and understand how it all worked. There was an answer to everything if only we could see enough detail. This was in no small way linked to the closed system approximations of physics set out in Newton’s laws of motion. Physics and chemistry seemed so orderly, so law abiding, so predictable. You could know the future exactly, providing you start with all the right measurements and formulae, and of course have helpful idealisations.² Modelling something as a closed system makes it very easy to calculate what will happen because the exact number and properties of all interacting elements is known.

This mechanistic view led to the idea that anything at all could be made finite and fully knowable by drawing some boundaries, learning everything there was to learn about everything inside those boundaries, and then expanding the boundaries to repeat the process. Physics and chemistry pervade the rest of existence, so it only stood to reason that the same orderly, law abiding behaviour that applies when we draw the boundaries around atoms and molecules must apply to everything made from those atoms and molecules. At higher levels, the levels of plants, animals, people, fashions, stock markets and technological development, there are regularities

²Please be careful, we seem to have misplaced our frictionless surface.

and patterns that bolstered the belief the rules governing these too could be made obvious and written down. And what hope! We could sweep away all unfairness, poverty, waste, and bad music, among other evils, if only we could find out Newton's three laws of humanity! All of science strove to prove that other disciplines were as well behaved and tractable as physics and chemistry, and that, in time, we could understand it all.

Yet differences were observed between the mechanistic predictions of society based on the behaviour of atoms and the reality of society, and these discrepancies were increasingly difficult to justify as errors in calculation or measurement. The idealisations began to look impossibly distant and systems somehow never seemed to be as closed as they needed to be. Exceptions were found to the laws governing physics, although not on the scale of ordinary life. Self-organising chemistry was discovered, which, while obeying all the rules, remained stubbornly unpredictable in the big picture. And all the rules found for disciplines like biology, psychology, and sociology seem to come with a list of exceptions that makes them almost useless.

Slowly, the scientific community came to appreciate that perhaps law abiding behaviour at low levels did not equate with predictability and control at higher levels. The idealisations and falsely closed systems of physics and chemistry just didn't scale up. Something else was going on, and that turned out to be that the kinds of things people wanted to study were often open systems, where matter and energy flow in and out, and where things inside a system are affected by the environment outside the system. This changed the view of systems from a boundary around all things that interact to a boundary around those things that interact together in such a way that makes them interesting enough to draw a boundary around. Among all of this, the development of General Systems Theory (Von Bertalanffy 1972), strove to argue that open systems were valid scientific study and that a system is not a matter of absolute truth but a matter of point of view, of usefulness for a purpose, and of relative value. The system genie had been let out of the bottle.

Ever since people have been struggling to reconcile the leftover desire for an predictable and eternal clockwork world and the unpredictable world that only looks like clockwork when you chop off everything interesting. This had led to quite a lot of interest in ways to talk about, study, control and influence systems, and to capture the mysterious "what else" that is going on between the movement of atoms and the madness of crowds.

Greenhouse Example Let's consider how the greenhouse horticultural sector might look from a mechanistic perspective, with all the idealisations and assumptions that it entails.

Each greenhouse would be an optimised converter of raw resources (nutrients, seeds, energy, etc.) to finished products (vegetables, flowers, etc.). The basic conversion process would follow an immutable formula, so that a fixed amount of resources would always produce the calculated amount of product. As each greenhouse is optimised, there would be no reason to believe that the greenhouses are not equivalent

or interchangeable, or that the production or efficiency of conversion might change over time if not deliberately changed.

The flows of resources and products would be clear and unambiguous, and the flow of information between all components, such as growers and suppliers, would be perfect and synchronised. Fully informed rational decision makers would always behave identically, so predictions are relatively simple.

Regional governments would control the behaviour of the sector by regulating changes that put in or remove barriers to the flow of resources or products or by altering the conversion formula with taxes, subsidies or the like. These controls would achieve maximum competitiveness in the global markets. Any change in the system structure is assumed to be deliberate, rational and coordinated.

While charming, this account of a regional greenhouse horticultural sector looks more like a game of poker where all players have identical hands, all cards are played face up, and with an arbiter controlling the rules to achieve perfectly distributed bets. Clearly, this does not represent what we observe in reality. Greenhouses are not formulaic black boxes that convert x amount of seeds and water into y amount of tomatoes, nor are the inputs fixed or predictable. Participants are not fully informed, are far less than rational, and behave unpredictably. Regional governments find that their efforts to control systems are ineffective or produce the opposite of the desired effect, changes to the system are uncoordinated and the last thing from deliberate, markets are far from free and power is very unequally distributed along the supply chains.

The systems perspective, which has largely superseded the mechanistic world view, is far more useful here. But what is that system perspective? Follow us into the next section to learn more!

2.2.2 Systems³

Systems are many things to many people. To be clear, although the closed systems of Newtonian physics are also systems, we want to talk about the systems with fluid edges, boundary crossing influences, and inexplicably organised internal elements that none the less confound our attempts to describe them.

If we look in a dictionary, we find several related but independent uses of the word “system”:⁴

- A regularly interacting or interdependent group of items forming a unified whole.
- An organised set of doctrines, ideas, or principles usually intended to explain the arrangement or working of a systematic whole.
- Manner of classifying, symbolising, or schematising.
- Harmonious arrangement or pattern or order.

³Recursive structure is a characteristic of systems.

⁴<http://www.merriam-webster.com/dictionary/system>.

For our purpose of understanding and managing the evolution of socio-technical systems, the first and last definitions are useful. Systems consist of interacting and interrelated elements that act as a whole, where some pattern or order is to be discerned. But we can only say they act as a whole because we see a pattern, or perhaps, we only see a pattern because we define the system as acting as a whole. If we saw a different pattern, we might define the system as a slightly bigger or smaller whole, or if we had already defined a different whole system, we might see a different pattern. The circularity of this definition can be a little boggling, so we will adapt a definition from Ryan (2008). Systems:

1. are an idealisation;
2. have multiple components;
3. components are interdependent;
4. are organised;
5. have emergent properties;
6. have a boundary;
7. are enduring;
8. effect and are affected by their environment;
9. exhibit feedback; and
10. have non-trivial behaviour.

We are getting closer, but these points need a little more elaboration to be sure we all agree on what systems are.

Idealisation Systems are not actual entities, as such, but are idealisations or abstractions of a part of the real world. A given system might seem clear enough, an obvious unit or entity that apparently stands alone, like a greenhouse. But closer inspection reveals that no two greenhouses are exactly alike, with, for example, some producing their own energy while others buy from the energy network. Thus, to get a useful abstraction or idealisation of a greenhouse system, we must either exclude the energy production facilities of some greenhouses or ignore those that buy their energy.

Multiple Components Systems always consist of multiple components, usually guided by the structure of a system. Individual greenhouses might be components in the greenhouse regional cluster system, even while the greenhouse system is composed of individual technology components, like combined heat and power (CHP) units or aeration units.

Components Are Interdependent Systems differ from unorganised *heaps*, which also have multiple components, by the fact that the elements are interdependent and interact. For example, lighting systems can add heat to the greenhouse, even though they are not a heating systems, and both lighting and heating systems use power, which might be generated by a CHP unit.

Organised The interaction and interdependence is not random and unstructured, but follows a certain pattern of interaction. While it is not impossible for each component of a system to interact with all other components, in most systems certain components interact more tightly with a subset of components, and the interactions are of a limited type or direction. For example, tomato growers participate in tomato growers associations, and rarely interact with flower growers, even if they are their direct physical neighbours. The interactions are usually limited to discussions, demonstrations and the like, and typically exclude such possible interactions as backrubs, fight clubs or paid musical performances. Within the limits of the associations, some kinds of interaction are bidirectional, such as discussions, while others, like votes or announcements, are unidirectional, further organising the structure of the interdependence.

Emergent Properties As will be discussed in more detail in Sect. 2.5.2, complex systems display properties that cannot be understood by just looking at the properties of the individual components, but are created as a result of the structure and organised interactions between these components. For example, the price of tomatoes can not be directly determined by just looking at the costs of the facilities and resources used to grow them, nor even by looking at the total production and demand. The price of a tomato is determined by many things, not all of which are operating at the level of tomato production.

Boundaries Every system description must contain an explicit definition of what is in the system and what is outside it. This is relatively easy for closed systems or highly idealised systems, but more realistic attempts to describe social and technical systems quickly find that drawing these boundaries is tricky business. The decisions on what to keep in or out of the system description depend on who is looking at the system, what they observe, and with what purpose they are making this observation. For example, when describing the boundaries of the horticultural system in the Westland, a politically based system might prefer to draw a strict geographical boundary at the edges of the municipality of Westland, while an economically minded description would likely include parts of the network of economic interactions around the greenhouses, even if they were located over the municipality border.

Enduring A system can only be considered a system if it lasts long enough to be observed or discussed, but conversely, it can be so enduring that it is hard to observe or discuss. Close observers will measure and study systems that others dismiss as too transient, but might fail to notice very slow acting systems with changes too long-term to be readily noticeable. For example, we easily consider the horticulture sector as a system because it has existed as a clearly identifiable kind of agriculture for a sufficiently long time that people have studied it as a system. On the other hand, an individual horticultural company that goes bust quickly would be too short term to interest most researchers, and although agriculture itself is a system, it is so long standing and wide spread that observations or discussions of agriculture as a whole are almost meaningless.

Environment Defining a system as an observer-dependent abstraction with observer-dependent boundaries means that the system is a particular interpretation of a particular subset of the real world. The “rest” of the real world is the environment in which the system is situated. All systems communicate mass, energy and information with the environment, so are to some degree “open”, but to simplify the system description, we abstract the environment into only those variables or parameters that are most relevant for the system. For example, when understanding the Westland greenhouse system, we might consider Chinese and US tomato growers as an influence from the environment, while we would probably not consider a music festival in Rotterdam, even though it is right next door, as music festival visitors are unlikely to alter the tomato system to a noticeable degree. However, if it became fashionable to buy a kilo of tomatoes after attending such a music festival, then we would need to include the festival in the environment of the tomato growers.

Feedback The interactions between system components are not only organised, but also contain loops, where A influences B and B influences A in turn. These loops create feed-back and feed-forward mechanisms that give rise to non-trivial behaviour. A positive feed-back loop, where the success of something drives more success, would be exemplified by the introduction of the “Tasty Tom” tomato variety, which was popular enough to generate high demand, driving more production, resulting in access into new markets, where it also proves popular and drives yet more demand. A feed-forward loop example would be the case where an anticipated drop in product prices might encourage growers to maximise quantity of production rather than quality, which results in a market flooded with poor quality products that receive a lower price, driving producers to carry on maximising quantity rather than risk switching to quality focused production.

Non-trivial Behaviour Foerster (1972) describes trivial behaviour as invariant mapping between system inputs and outputs, but the different feed-back, feed-forward and other interaction loops, coupled to inputs from the environment create non-trivial behaviour in complex systems. For example, if the price of tomatoes drops sharply, a trivial behaviour would be a proportional drop in supply, but in reality the supply does not behave trivially. Instead, although supply might be affected, tomato growers must balance the drop in price with long-term contractual agreements with suppliers, a sense of pride or family traditions in producing tomatoes, significant investments that need to be recouped, and many other factors, irregardless of the economic conditions.

2.2.3 *World Views*

Whenever we interact with a system, we are never passive or objective observers. In order to see, we need to choose to look, and although reality is infinite, our powers

of observation are limited. To cope with the onslaught of infinite reality,⁵ we have developed several strategies to whittle reality down to comprehensible observations. We observe everything through the lens of our human nature as well as that of our individual world view, which is not entirely under our own control. At the same time we are, to some degree, in control of our powers of attention and focus, or what it is that we choose to observe and how we do so. The question therefore is not whether there are different world views, but how these world views affect the observations and interpretations of the system it views.

Some of these universal strategies to gainfully reduce the tide of the potentially observable are derived from evolutionary history. We preferentially attend to rhythmic motion because it is generally more urgent to identify animals than plants, and we cheerfully ignore microscopic organisms because our human scale is ill suited to that size. But other strategies are part of the world view, or fundamental and relatively consistent cognitive orientation of an individual or society, and include biases that are not necessarily shared, but are derived from experiences and interests. Classically trained musicians will notice off key notes that most would not, while engineers will be fascinated by mechanical devices that musicians might ignore. These world views determine not only what we pay attention to, but also provide a framework for generating, sustaining, and applying knowledge. For example, economic woes will appear to be the direct result of too much or too little regulation depending on how financially conservative the observer is, with each side arguing that any evidence against their point of view must be faulty.

Of course, both sides could be said to be right because they are essentially disagreeing about how they define the system. No two observers will agree on system boundaries, how to parse the system into relevant parts, or how to rank, measure and track the things determined to be of interest, all of which in turn will determine what we attend to, measure and observe in the future. This is especially apparent with socio-technical systems, as there are no unambiguous boundaries, no natural scales, and few shared contexts to guide our attention, efforts and observations.


Greenhouse Example As we are not greenhouse farmers, we find that we have a very different perception of technology than they do. For example, when deciding on the purchase of a new technology, greenhouse farmers seem to view technology as unique and personal, for which only personal experience can be of help. They tend to value their own experience the most, followed by the experience of other greenhouse growers. They express a distrust of the information provided by the technology producers, academic analysis, or government agencies, perhaps expecting it to be painted in a favourable light or applying only abstractly to the actual performance that can be expected. Essentially, the less the person giving the information is like the grower, the less they trust the information given.

But as academics, we are inclined to view technologies as homogeneous and impersonal, and obviously do not see a problem with trusting the results of scientific

⁵It just keeps coming.

Fig. 2.1

Observer-dependence. By Randall Munroe (<http://xkcd.com/638/>), used with permission



analyses, especially if the data is provided, expecting the analysis to be replicable and representative of the technology performance. Yet academics might also express a distrust of corporate figures, government statements and the anecdotal evidence of greenhouse growers, if there is no scientific data as support. Which just goes to show that greenhouse growers might trust themselves first and those most like themselves second because the information is so specific and individual, but academics do the same because they believe the methods of science eliminate the specificity or individuality of the subject. Perhaps growers and scientists are not so different, despite the orientation of the distrust.

2.2.4 Observer-Dependence

While world view affects all aspects of perception and understanding internal to the individual, observer-dependence affects the externalisation of that world view as it applies to a system under study. Every observer must choose the scale at which the system is observed, including time scales, and thereby determines what is considered the smallest/largest elements or periods of interest. For example, one observer might see people as agents, interacting in a legal system over the course of months or years while another observer sees entire countries as agents, engaged in worldwide political games over a matter of decades. Furthermore, each observer chooses a certain perspective when interacting with a system. For example, a chair can be viewed as an artistic, social, mechanical or chemical entity which determines whether the chair is seen as an example of post modernism, a variable leading to improved classroom behaviour, or as an object defined by properties such as mass, strength and volume. The cartoon depicted in Fig. 2.1 illustrates how one's choice of what and where to observe can support an erroneous idea.

Objectivity

“An objective account is one which attempts to capture the nature of the object studied in a way that does not depend on any features of the particular subject who studies it. The object

has certain properties or behaves in a certain way even if the subject (you as a person) does not see it. An objective account is, in this sense, impartial, one which could ideally be accepted by any subject, because it does not draw on any assumptions, prejudices, or values of particular subjects. This feature of objective accounts means that disputes can be contained to the object studied.”

Gaukroger (2001)

Science is at least implicitly understood to be a method for exposing irrefutable facts, and objective, real truth. Those who disagree with the findings of science, like creationists, do not usually argue that science cannot find the true nature of things, such as mechanisms for cell growth or immune system functioning. Instead they suggest that a particular application of science is faulty or biased in some respect, or that there are aspects of life that science cannot examine for truth, like the ineffable nature of the creator, even if such an investigation can find truthiness in fossil records or geological studies (Colletta 2009; Munger 2008). In effect, almost all of us believe that science is capable of objectivity, although not universally so.

But the concept that science can truly be objective is an unattainable goal, an unhelpful ideal, and a needlessly divisive crutch, even for the most uncontroversial topics or objects of study. Objectivity, seems so attainable, but is elusive to the last because the observer is responsible for selecting the object of study, the aspects to measure or record, the tools or instruments to use, and the methodology to follow, among other crucial choices that determine the final outcomes of the scientific endeavour. Inevitably, some features or qualities that could be measured will be ignored or mis-recorded, and the instruments and methodologies chosen cannot help but be less than perfectly precise or balanced. Furthermore, the interpretation of the data collected, the creation of models to best fit that data, and the predictions of the future, which guide measurements and research to come, are also highly dependent on the scientist rather than the object of study.

Total objectivity is arguably not possible in some, or maybe even all, situations. But do not despair! While science cannot promise objectivity, we can at least be aware that:

- each observer has its own world view that determines what and how observations are made;
- observer-dependence interacts with emergence; and
- observer-dependence affects the process of model creation.

We can not reliably and objectively determine what effect these will have, but awareness of these issues makes one better equipped to approach observer-dependence. For a more controversial approach to objectivity in science, please refer to the work on Post Normal Science by Funtowicz and Ravetz (1993).

Greenhouse Example Subjective analyses come from very different sources. For example, two successful greenhouse growers looking to purchase the same property in order to increase their production will offer different bids. They will have taken different aspects of the same property into consideration and arrived at different

subjective evaluations. For example, one will take the relatively high energy consumption to be more important than the benefits of an excellent location near an auction house, while the other will value the location highly and dismiss the importance of the current energy demands.

Another, and far more subtle, example of objectivity vs subjectivity would be the measurement of the temperature in the greenhouse. Decisions must be made as to whether the temperature should be measured in one or more locations, where these locations are, how often the measurements are made, what time of day they are made, and the type and sensitivity of the device employed to measure and record the temperature. While we might think that temperature is objective, there is no really objective way to measure it, and the choices made will influence the data gathered.

Reductionism and Holism

“The utility of the systems perspective is the ability to conduct analysis without reducing the study of a system to the study of the parts in isolation.”

Ryan (2008)

Reductionism is the idea that system behaviour is determined by the behaviour of the system components, and is best exemplified by the idea that a thing is nothing more than the sum of its parts. Reductionism can be understood as an attempt to achieve objectivity, by reducing the systems to smaller, more observable, and therefore more objective, pieces. According to reductionism, understanding all of the parts is equivalent to understanding the thing. Related to this idea, downward causation says that the upper levels of hierarchical structures constrain the actions of the lower levels, or that to understand all of the parts of a thing, you need only understand the thing. Regardless of whether a thing can be understood only in terms of its parts, or the parts can only be understood in terms of the thing, both reductionism and downward causation quickly lose meaning in nested systems, where each of the parts are in turn made up of sub-parts. Taken too far, this leads to large-scale behaviours or events, like consciousness or stock market crashes, being explained in terms of the laws governing sub atomic particles.

The opposite stance would be holism, where a system, be it physical, social or biological cannot be determined or explained by its component parts alone. Put more colloquially, the whole is more than the sum of its parts. This view is linked to upward causation, where the parts that make up a system are not constrained in any way, but the system is constrained by those parts, so that the lower level system components provide all the possible behaviours that a system can have. Like reductionism, holism can be taken to ridiculous lengths by claims that it makes no sense to examine the component parts of a system at all and that only a study of the system in its entirety leads to understanding.

The kind of reductionism that absurdly tries to explain any phenomenon based on the smallest possible components is called greedy reductionism (Dennet 1996), and while the explanations of the smaller components are not necessarily untrue, they are remarkably unhelpful for explaining the higher level system. Likewise, extremely holistic approaches are unlikely to be of any use, especially for addressing urgent

or current issues because an exhaustive study of whole systems is time consuming and difficult, if not impossible, and extremely inaccessible to anyone who has not studied the same system.

More fruitfully, the extremes can be avoided by recognising that the links between the thing and its parts are influential but not completely causal. If the behaviour of the system and its constituent parts influence or constrain, but do not completely determine each other, then there is a clear benefit to looking inside and observing the parts and interactions while also looking at the higher levels as well. Yet avoiding extremes of reductionism and holism demands an understanding of other problems, such as the observer-dependence of delimiting the boundaries or defining the contexts of the systems under study.

Greenhouse Example A reductionist stance would be that the performance of a greenhouse is only a function of the technologies that are in it. An extreme holistic stance would claim that each greenhouse has a unique performance, and it irrelevant that all top performing greenhouses are all using a particular technology for heating. A useful reductionist approach will perceive the connection between the performance of the greenhouses with the technologies implemented within in, but will also recognise that the location, technology interactions, state of maintenance, management style of the owner, as well as factors outside of the greenhouse, are also relevant.

2.2.5 *System Boundaries*

“The real challenge posed by the systems idea: its message is not that in order to be rational we need to be omniscient but, rather, that we must learn to deal critically with the fact that we never are.”

Ulrich (1988)

There is no system with an outer system boundary through which no energy, matter or information penetrates to influence the internal workings of the system, although some systems are simple enough to be usefully modelled as if such boundaries existed. The systems of interest to us, however, are not so easily idealised, so we must decide which parts, relations or influencing factors are not known or suspected to influence the system strongly enough to be worth the effort of including. And we must do so in full awareness that an arbitrary boundary is drawn around what seems like a useful subset of a larger system (or systems) when we observe a cluster of activity that warrants further study.

Although in reality, everything influences everything else in some way, Ulrich is right to stress that we do not need to know how everything is connected. We really only need to be aware that the boundaries we decide to set will reflect our needs and goals rather than the true nature. If we want to examine how rising fuel prices affect the greenhouse horticulture sector in the Netherlands, we might or might not want to include representations of transport, horticulture and agriculture in other

countries, technological development for energy efficiency or international markets. All of these things, and many more, will be connected to the greenhouse sector, so there is no inarguable system boundary to be drawn. But if too many connections are included then clear relationships and influence will be harder to elucidate and the model will be no more enlightening than simple observations of the real world would have been.

2.2.6 System Nestedness

When the highest level of one system is also the lowest level of a larger system, then the systems are vertically nested. Each system can be viewed as an isolated whole, or can be viewed as composed of other systems and residing in an environment made of the next higher level system. Deciding whether a given level will be represented as a unit or as composed of smaller units can be understood as a sort of vertical system boundary. The level of observation must therefore also be made without pretense that it is the only possible level at which the system could be observed. Instead, it is the level at which the expected observations are most likely to lead to improved understanding of a given question about the system.

Although conceptually arranged in hierarchies, Hollings points out that these “hierarchies” are not top-down sequences of authoritative control, but rather, semi-autonomous levels formed from the interactions among a set of variables that share similar speeds (and, we would add, geometric/spatial attributes)” (Holling 2001). These nested or hierarchical arrangements are a sort of conceptual shorthand based on the way evolution tends to develop stability from frequent interactions, both in time and space. These stable interaction patterns appear as structured units which interact with similar units and serve as building blocks in larger structured interactions. For example, the inhabitants of a town interact much more frequently with each other than they do with the inhabitants of another town, so each town could be considered a system, embedded in a larger system for the region or country. But even though the towns are usefully idealised as separate systems, the residents of each town are not constrained or prevented from interacting with the other, they are just less likely to do so. Instead, in nested systems, the subsystems overlap, and it is this overlap and the interaction that it enables that given rise to complex behaviour Alexander (1973).

Because the lower level stable structures are much more likely to interact on the same scale, the interactions at other levels seem remote and simplified, the more so the more distant the level. Therefore, “three levels of a hierarchy, one up and one down, are generally considered sufficient for analysis of the focal level” (Ryan 2008). Although, if there is any disagreement as to what is the focal level, there will of course be disagreement as to what the upper and lower bound of observation should be.

It is important to note that systems can be nested in time, physical space and social relations, among other possible ways, and that every system can belong to

more than one larger system as well as be composed of more than one arrangement of smaller systems.

Greenhouse Example A greenhouse farmer can belong to a physical neighbourhood, inside of a town, inside of a district, inside of a country, as well as belonging to more than one growing association, inside of larger trade unions or industrial sectors. The farmer is also likely to belong to family units, inside of a larger extended family as well, and to a local religious group, club or team inside of larger associations. Inside the greenhouse, systems can also be nested. A temperature regulation feedback loop is a system nested within the greenhouse heating system. The heating system is nested within the climate control system of the greenhouses, together with the aeration, lightning and irrigation system. The greenhouse is a system nested within the district heating system and power grid of the region of Westland, which in turn are nested within the Dutch and ultimately European power grid.

2.3 Adaptive

To be adaptive is to have the property of adaptation, or improvement over time in relation to environment. The environment need not be physical, as social, technical, and cultural environments can also cause adaptations.⁶ Adaptation is not the same as change in response to a stimulus because adaptations are specific kinds of changes in response to specific types of stimuli. The changes must be improvements (how to determine what is an improvement is covered in the next section) as changes that make something worse or merely different while being neutral in respect to the environment are not adaptations.

Further, the changes must be in response to stimuli from the habitat or environment. These stimuli can be purely environmental, like temperature, terrain, or the availability of resources to which adaptive entities can be come better suited. The stimuli can also be contact or interaction with the great diversity of other entities, ranging from direct adversarial interactions, such as predation, competition for resources, and parasitism, to beneficial interactions such as cooperation and symbiosis.

Lastly, the stimuli from the habitat must be constant or reoccurring. Gravity is static, and of course everything that has adapted has adapted to deal with gravity. Likewise, entities can adapt to the dynamic but periodic (within a lifetime) patterns such as climate and diurnal rhythms, growth and decline of population numbers, cycles of resource availability, seasonal migrations, and many others. These relatively predictable and ordered environmental stimuli can force adaptations, but

⁶The specifically improved features, behaviours or traits are also called adaptations, but it is preferable to refer to these as adaptive traits to avoid ambiguity.

catastrophic change events cannot. Some events are so disruptive that a great majority of species go extinct, areas become uninhabitable or unusable, and built environments are completely destroyed. Nothing can adapt to such extremely rare, sudden and utterly devastating global catastrophes.

Adaptations are not just change, or even change in response to stimuli, but neither is adaptation the same as evolution.

2.3.1 Adaptation Versus Evolution

While adaptations are improvements in response to environments, evolution is the algorithmic process that produces these improvements, best summed up by the famous maxim: “Vary, multiply, let the strongest live and the weakest die” (Darwin 1985). An enormous body of knowledge now exists on the evolution and co-evolution of biological systems following the publication of Darwin’s book “On the Origin of Species” (Darwin 1985), which has now developed into one of the best researched and supported scientific theories today. Although Darwin did not understand the specifics of how DNA, mutation, or some selection pressures worked, he quite rightly surmised that evolution will occur whenever certain conditions are met.

The first of those conditions is that there must be differences between things, or variation. Variation occurs, for example, through the addition of totally new material, either through creativity or merely as a result of copying errors, and through the combination of existing designs or genes in new ways.

The second condition is that variations must be replicable or heritable in some way, so that even if offspring or copies are not perfect, they are at least more likely than not to have some of the same variations of the parent or original. This condition of replicability seems to be in conflict with that of variation, with one demanding differences and the other similarities. Of course, the key is that neither is absolute, balancing each other out so that descendants are more similar to the progenitor than to others, without being identical.

The final condition is that there must be some determiner of which replicable variations are better than others. Selection, or selection pressure, is the force in the environment that determines how well suited one variation is to the environment, sifting them mercilessly into the winners and losers, the quick and the dead. Darwin noticed that many animals did not live long enough to reproduce, while some were very successful and had many offspring. Of course we see the same pattern in products, music groups, sports teams, companies, and a myriad of other entities that become successful or not as a matter of the selection pressures acting on them. Only those variations with a leg up on the competition last long enough to produce copies or offspring to populate the future.

Crucially, all three conditions are equally important for the proper functioning of the algorithmic process, and it reveals that evolving things are more than just a collection of matter and chemicals. They are “boring stuff organised in interesting ways” because there is also information, encoded in the structure of the matter

and chemicals, and in the and relationships between these structures and the environment.

2.3.2 *Evolution—More than just Biology*

Traditionally, evolution has been most readily recognised and commonly accepted in the Darwinian process of natural selection that drives changes in the genetic material of living beings to adapt to their physical habitats. But, as an algorithmic process it is domain neutral and can explain non-genetic changes shaping non-biological adaptations as well. Darwin thought languages were also evolving through the same processes as organisms, and scholars such as Mandeville and Harth (1989), Hume (1962) and Smith (1963) considered that industry, law, politics, economies, markets and manufacturing were also shaped by evolutionary forces. Not having physical matter or DNA as biologically evolving entities do, languages, industries, laws and the many other non-biological entities must use some other system, and the leading theory (Dawkins 1990) suggests that *memes*, cultural, self-replicating entities, analogous to genes, are responsible for the spread, construction and evolution of culture. In the words of Dawkins (1990) and Dennet (1996):

“A meme is the basic unit of information which spreads by copying from one site to another and obeys, according to Dawkins, the laws of natural selection quite exactly. Meme evolution is not just analogous to genetic evolution. It is the same phenomenon. Cultural evolution simply uses a different unit of transmission evolving in a different medium at a faster rate. Evolution by natural selection occurs wherever conditions of variation, replication and differential ‘fitness’ exist.”

Proposed examples of memes are traditions, technologies, theories, rules and habits. This book can be seen as a meme that might or might not survive over time, depending on its usefulness to the persons who are aware of its contents and find it useful enough to tell others about it.

Ziman argues (David 2000; Jablonka and Ziman 2000) that both biological and socio-technical entities display variation, replication, and selection through a succession of generations, determined by reproduction cycles in organisms, and cycles of learning or imitation in social systems. Further, Ziman observes (David 2000) many phenomena that arise as a consequence of the evolutionary algorithm in both domains, including diversification, speciation, convergence, stasis, evolutionary drift, satisfying fitness, developmental lock, vestiges, niche competition, punctuated equilibria, emergence, extinction, co-evolutionary stable strategies, arms races, ecological interdependence, increasing complexity, self-organisation, unpredictability, path dependency, irreversibility and progress. While there is a lot of support for the fact that social and cultural things are also evolving by the same algorithmic process as biological things, there are several points or criticisms that merit further discussion, although most of these stem from ill-advised attempts to use a strict biological analogy as the model for non-biological evolution.

First, there is no clear relationship for social and cultural artifacts that compares to the organism and gene relationship. This immediately proves a deal-breaker for

many critics of memetics, but there are good reasons not to dismiss non-biological evolution out of hand. The relationship between organisms and genes was not at all clear when Darwin proposed his theory, but that didn't stop further investigation. Additionally, the link between genes and organisms is not as well understood as it may seem because of the complex, epistatic relations between genes, preventing any simple one to one cause and effect.

Second, the idea that genes have a special role in evolution only holds for advanced organisms. Not all evolving organisms have DNA, and the very first replicators were not organisms with any kind of genetic code. They were most likely simple self replicating molecules, with no distinction between genotype and phenotype. It is not clear if socio-technical evolution has moved beyond the early stages to arrive at special roles for the non-biological equivalent to the gene.

Third, many researchers cheerfully state, without clear evidence, that technical and social artifacts are not generated randomly but are purposefully designed, although randomness is so important that many designs are described as serendipitous flashes of inspiration (Roberts 1989), accidents that prove useful or the lucky result of arbitrary task assignments during the design process. Further investigation reveals that even those revolutionary, clever designs were built on foundations of rigorous testing, trial and error and incremental advances on previously successful designs, which looks quite a lot like the random mutations in each generation of biological organisms.

Fourth, biological evolution operates in generations that are strictly vertical (no one can be their own parent) while social and cultural evolution has been suggested to move horizontally or reverse vertically, as when someone learns something from a parent, and then teaches that parent an improved version of the lesson. However, this confuses the relationship between the people that teach and learn with the things taught or learned. Children can certainly learn from and teach their peers and parents, but a thing can never be taught without having been learned first. Thus, the things that are actually evolving through non-biological evolution are also strictly vertical, in that they must be learned before being taught, just as organisms must be born before they can reproduce.

Fifth, the speed of evolution is much faster in culture, since there is no need for the genetic transfer taking place over generations. It happens in the "meme sphere", the shared human cultural space, and not in the biosphere. Although entirely true that the speed seems to differ, this can hardly be considered a criticism.

And finally, created artifacts are not alive and do not reproduce in the sense of creating offspring. For some people, aliveness seems to be a necessary condition for evolution, despite the fact that viruses clearly evolve, requiring a new flu jab every year, but are not often considered alive. Lee Cronin, an inorganic chemist, has suggested that anything that evolves should be considered alive, but the opposite was once considered commonly accepted knowledge. Thus, Darwin compared evolution by natural selection to the development and change in languages, because "everyone knew" that non-living things could undergo gradual change processes while living things were immutable.

Although memes appear to be one of the best ways of approaching socio-cultural evolution, they may not be as necessary as many might think. It may be preferable to

look at socio-technical evolution on its own rather than through the lens of biological evolution because the systems appear to be two examples of a generic evolutionary algorithm instead of one being real evolution and the other an analogy.

2.3.3 Adaptation in Its Many Forms

Adaptations always start with what is currently available for use, and improve it or apply it in new ways that do not decrease fitness in the immediate term. The best adaptations are those that use the tools at hand the best, not those that can identify what the best of all tools would be. But adaptations are not just the physiological traits of exotic animals that come to mind so readily. There are actually three levels at which different kinds of adaptation can be found, all of which take advantage of different tools and operate on different time scales. These are the individual, cultural and biological scale.

The shortest time scale for adaptation is the individual lifetime, giving us individual adaptiveness. The best example of individual adaptiveness would be learning (Argyris and Schon 1996), although other examples include muscle development in response to constant use, improved immune response after exposure to pathogens, and changes in appearance or behaviour, such as developing a tan in a sunny location and decorator crabs that add bits of seaweed to their shells. In the non-biological realm, individual adaptations would include the unique wear and tear that makes something perform better as it “breaks in”, the addition of new words in the personalised T9 dictionaries on mobile phones, or the gradual increase in predictive power of smart systems as they find patterns in their input sensors. Anything that an individual could do, in its own lifetime, to become better suited to its environment is an individual adaptation, but these are not replicable.

If an individual adaptations comes to be imitated or reproduced as a consequence of learning, then it forms the basis of cultural adaptations. Cultural adaptations take place at a slower time scale than individual adaptation because they necessarily involves at least two occasions of individual adaptation. Humans do many things that fall into this category, so much so that it can be difficult to see that the many ways that adults teach youngsters to stay safe, gain access to food, communicate with others or avoid danger are cultural adaptations. Animals too have cultural adaptations, although less universally agreed upon, these include explicit lessons, as when chimpanzees teach their offspring to break open nuts or angle for termites, and also non-explicit cases of imitation such as songbirds learning to copy the local dialect, foraging habits, food preferences, nesting sites, etc. Social and technical examples of cultural adaptation are also rife, and include schools of thought on everything from economics to the right way to serve tea, designs of houses, models for approaching the design of tools, measuring devices, infrastructures and systems, and just about everything else that people teach or learn not directly related to survival.

The longest time scale for adaptations is also the most well known. Biological adaptations develop over many generations and are not learned, neither individually

or culturally, but relate to inborn traits or instinctual behaviours. Biological adaptations can be quite spectacular, such as the angler fish's glowing and twitching lure which helps him attract food in the dark depths of the ocean, or quite unremarkable, such as the proper functioning of internal organs. Not having biology, socio-technical artifacts may not be evolving at this level, although further development, such as in self-replication robots, may require the replacement of "biological" evolution with something that operates at this level and timescale but includes both biological and non-biological evolution.

An adaptation that first appears at one level, can potentially move to another. If an important lesson is learned by an individual, they may teach another or be imitated without explicit teaching, moving the adaptation from the individual to the cultural level. If the environment allows for a Baldwin effect⁷ (Weber and Depew 2003), then the lesson may become a genetically inherited instinct, resulting in the same outward behaviour but without the effort or risk involved in the learning process by moving the adaptation from the individual or cultural level to the biological.

2.3.4 Direction of Adaptation

Adaptation and selective pressures often appear to "want" to go in a particular direction, toward what is called an attractor, and can display self-reinforcing effects that seem to accelerate the direction of adaptation. This can be represented visually through fitness landscapes (see Fig. 2.2), a description of the conceptual environment of an individual or species, where every point in the landscape corresponds to one possibility in the space of possible, with the fitness of each possibility represented as altitude. The more fit possibilities are the attractors, depicted as the peaks of hills. New variations are made by stepping from one location to another, and selection pressures reward those steps that move uphill, or the adaptations, and ignoring lateral steps as neutral changes. Steps that move downhill are punished as deleterious or maladaptive changes.⁸

While an evolving entity located on a plain between hills could go in many directions without changing fitness, once in the basin of attraction for a peak, the only way to adapt is to continue climbing the hill. Progress up the hill then appears to accelerate due to the self-reinforcing effects of adaptations, where upward motion means that there is less room for lateral movement. If the hill also grows steeper then each step equals faster movement toward the peak. Step size, or the rate at which things can change, becomes an issue too. Bigger steps means faster progress toward the peak, until the peak is closer than the size of the step, at which point

⁷A theory exploring how acquired or learned behaviour can become integrated into an organism's genetic markup.

⁸Alternatively, the attractors are sometimes depicted as the valleys in the fitness landscape, based on the premise that things can only roll downhill, and do so quite naturally.

overshooting the peak means no progress is possible. On the other hand, very small steps are likely to ultimately get closer to the peak, but will take much longer.

Evolution can only locally optimise so evolvees may get driven up a hill that is not the highest on the landscape, but only coincidentally the first encountered. However, movements that descend the hill to search for a better hill are impossible because even a temporary decrease in fitness will be quickly killed off by the short sightedness of selective pressures. As such, hills can only be climbed, never descended, giving the appearance of inevitability or a teleological drive. This misleading interpretation derives from focusing on the well adapted and successful. Viewing only the victors of ruthless competition suggests that becoming better is somehow the meaning of life, the universe and everything. From a wider perspective, it is clear that there is no goal or purpose, because there are far more losers than winners, and that the competition is endless.

Theoretically, socio-technical evolution allows for long jumps, from the top of an attractor hill to what may prove to be at least equal altitude on the slope of another hill by means of imitating what appears to be a more promising design or solution. However, it is not clear that imitation can really be considered to be “jumping” from one peak to another. It depends on the point of view as to what it is that is walking up the hill. If a company is walking up the hill, then abandoning one product design in favour of copying that of a rival would certainly be a case of adaptation, but both companies would be on the same hill and the copier would just be following in the path already taken by the other. If instead the products are seen as the hill climbers, then they are indeed climbing different hills, but copying the competitor would be the extinction of the design at the top of the lower hill while a new competitor appears simultaneously on the surface of the other hill. While fitness landscapes with attractor hills are useful metaphors, they require careful consideration about what exactly the evolvee is and what the landscape might look like so as not to get mixed up.

2.3.5 Coupled Fitness Landscape

In a system, all elements interact, but an adaptive system entails that those elements not only interact, but adapt to each other, reacting to selective pressures to become better suited to the system as a whole and the environment in which the system is situated. But every adaptation changes the environment and selection pressures acting on the rest of the system, so leads immediately to slightly different selection pressures. Since selection pressures are what shape the hills and valleys of fitness landscapes, the changes from every adaptation means the hills are alive, constantly jumping up or falling flat, moving around and changing shape to reflect the new selection pressures that are pushing towards the new attractor hill tops.

Everything is connected and nothing is stable. Referring to evolution as co-evolution emphasises that nothing exists or evolves in isolation. Every action of every element in an evolving system will have some effect on other elements. A comprehensive overview of the co-evolution literature is beyond the scope of this work,




Fig. 2.2 Deformation of a fitness landscape

although it has already been thoroughly described in the literature of biology (Futuyma 1983; Jantzen 1980; Thompson 1994).

As nothing adapts in isolation, the fitness landscape can be recast as a coupled fitness landscape. This works particularly well in identified “arms races” where two species are each adapting to the last adaptation of the other and upping the selection pressure to adapt further, ratcheting up the height of the hill in each other’s landscape. A classic example is that of the cheetah, which has adapted to run faster than any other land animal, and its usual prey, the gazelle, which cannot run as fast over the short distance but can maintain a very high speed over a much longer distance than the cheetah.

Hordijk and Kauffman (Kauffman and Johnsen 1991; Wilds et al. 2008) use the x and y axes of a coupled fitness landscape to represent the possible ranges of properties of two interacting species, such as the cheetah and gazelle. The z axis represents the combined fitness landscape of the two species, with peaks and valleys for combinations of different properties of the two species. The coupled fitness landscape is dynamic, with each adaptive step taken by one species distorting the fitness landscape of the other, and vice versa. This is illustrated in Fig. 2.2, where, going from left to right, the fitness landscape is deformed as species evolve and acquire new traits (e.g. higher top speed, more endurance, faster reflexes). If the gazelle adapts to run faster, the cheetah must adapt in some other way to deal with faster gazelles, reducing the gazelles fitness again. The responses and counter-responses can not be predicted in advance, as it might be higher speed, or better camouflage, or better sensory perception that provides the next temporary advantage. Of course, coupled fitness landscapes can be applied to socio-technical evolution as well, as companies battle to devise the next ingenious way to compete for customers, as new computer viruses develop sophisticated techniques to escape detection while virus detection companies also grow better at detecting, and as corporate loopholes are closed, only to allow some other loophole or tactic to be exploited.

But even coupled fitness landscapes are too simplified to deal with an entire adaptive system. They can only represent a few adaptations at a time, those assumed to be important, in a theoretically idealised population, in a way that can only look at the known past. Individuals, products, designs, companies, solutions or species,

each with unique combinations of features and traits, can be represented on its own fitness landscape. When selection weeds out the losers, one evolvee may lose out, despite being the highest up the hill for a given trait, because of a fatal step down the hill for another trait. So every unique trait or feature could have its own fitness landscape, with all of them coupled to capture the fitness landscape of the evolvee, which in turn would be part of a larger fitness landscape that captures interactions between the evolvees. Just as the levels of nestedness in system boundaries can be understood to go on forever, so do the levels at which fitness landscapes can be coupled. It all gets incomprehensible very fast.

Irreversibility Attractors and fitness landscapes are theoretically related to another concept, called path dependency, also known as high switching costs (or sunk costs) (Economides 1996), group think (Janis 1982), and lock-in (Teisman 2005). Path dependency is captured by the idea that “history matters” (Buchanan 2000) because past decisions influence the future decisions to be made, which leads us directly into the concept of irreversibility.

Steps only go laterally or uphill, never downhill. But even a step laterally can not be reversed in an adaptive system because the landscape changes after every step. Whenever a move is made or an interaction is established, all other conceivable moves or interactions possible at the last step are no longer possible, although there is a whole new set of moves or interactions available. Any living or evolving process involves thermodynamically irreversible processes, so path dependency is “baked in” to reality at every level.⁹ Thus any system that changes involves irreversible processes.

This irreversibility or path dependency applies to the system’s overall behaviour, which can manifest itself in many ways. Physical systems lose mass or energy (Prigogine 1967), while social systems lose information. These losses also cause shifts in the landscape, affecting the future possibilities.

2.3.6 *Intractability*

Evolution is an algorithmic process (Dennet 1996) of variation, replication and selection across the evolutionary design space. Computational theory (Hartmanis et al. 1983) states that evolutionary problems are intractable,¹⁰ that is, that future steps can *not* be calculated any faster than the time required to take that step. Intractability implies that the outcome of the evolutionary ‘program’ can only be found by completing its execution, which is to say, we just have to wait and see because predictions are impossible.

⁹This inherent reversibility is also called the “arrow of time”, but if you start thinking that time arrows are baked into the fabric of reality you might be mixing your metaphors.

¹⁰Although normal English usage of “intractable” means uncooperative or stubborn, we are specifically using the computational complexity theory definition.




Fig. 2.3 An intractability path

Thus, adaptive systems are impossible to predict with any reliability or exactitude, which means we face no small task when trying to understand and steer the evolution of adaptive systems, as discussed in Chap. 1. It can be mathematically proven that we will never truly know the precise effects of our actions, and we must thus act accordingly. An illustration of this process is presented in Fig. 2.3: Let us imagine a system being at some arbitrary point O in the system's history. At time A , something happens to move the system towards point A , forever excluding all the states toward which the system could have evolved but which are no longer possible. At time B , another interaction event happens, again excluding countless possible future states. As the system progresses in time, across points C , D , E , F , etc., more and more of the astronomically large number of possible system states are not able to come into being. Of course, at each time step, the same astronomical number of new possible states continuously becomes possible, as can be seen at point H .

We quickly realise that adaptive systems in nature are incredibly complex. Every component interacts with every other component in a possibly infinite number of ways (if not actually infinite, then certainly astronomically large and probably also Vast¹¹) closing off some options forever and shifting the infinite possibilities of the next interaction in the future in mega-coupled fitness landscape. As if that were not enough, “the system as a whole causally influences the state of its own constituents, which in turn determine the causal powers of the whole system” (Kim 1999), which is philosophically quite problematic, despite seeming to occur quite regularly wher-

¹¹Vast differentiates the super astronomically large from just ordinary large. For example, 10^{50} is a very, very large number. However, 1000^{1000} is Vast (Dennet 1996).

ever things are not run by philosophers who say self-causation is impossible. Some common examples include the mind-body link of psychosomatic disease like ulcers or hypertension, self regulation in social systems, the creation and endurance of norms and institutions, and, in fact, any adaptive system.

But before we jump ahead and look at how complex things behave, let's see more about what we mean by complexity.

2.4 Complexity

“I shall not today attempt further to define the kinds of material I understand to be embraced within that shorthand description ‘hard-core pornography’; and perhaps I could never succeed in intelligibly doing so. But I know it when I see it, and the motion picture involved in this case is not that.”

Stewart (1964)¹²

Complexity is, perhaps surprisingly, like pornography in that you can't really define it, but you know it when you see it. Definitions are observer-dependent and subjective, so rather than starting by trying to define what complexity is, we would first like to discuss what complexity is not, by looking closely at simplicity and complicatedness.

2.4.1 Simple

The most basic definition of complexity is that it is “not simple”. Being the conceptual opposite of complexity, simplicity therefore requires a more careful examination, although it should come as no surprise that simplicity is also an elusive, slippery and uncooperative concept.

Nothing in the real world is ever really and truly “simple”. Instead simplicity is relative, much like “big” or “wealthy”, so a thing can only be judged as more or less simple than another thing. These comparative measures are necessarily observer-dependent, so the same thing can be judged as simple by one observer but as not simple by another, and at the same time they can both be complex. Already in the 13th century Dominican philosopher and theologian Thomas Aquinas argued in his influential work *Summa Theologica* that God is infinitely simple, but will look complex to the finite human mind, because every observer cannot help but see through the lens of unique perspective.

Setting aside the relative nature and the observer-dependence that this entails, simplicity can still be broken down into structural and functional simplicity. We can evaluate which of two structures that perform a given function is the simpler

¹²Concurring opinion in *Jacobellis v. Ohio* 378 U.S. 184 (1964) regarding possible obscenity in *The Lovers*.

structure. Likewise, for a given structure, we can compare which of two functions that it can perform is simpler. However, trying to determine the relative simplicity of two or more structures, each of which may or may not perform two or more functions, starts to lose meaning because simplicity in function and structure appear to be largely, if not totally, mutually exclusive. Thus, something both functionally and structurally simpler than another thing will be rarer than hen's teeth as any increase in simplicity for one measure will be cancelled out by the other. Thus, simplicity does not apply to a thing, but rather only to aspects of that thing in a particular context, which brings us back to observer-dependence and subjectivity. This may be useful to bear in mind when looking at complexity later in this chapter.

Functional Simplicity

“Civilisation advances by extending the number of important operations which we can perform without thinking about them.”

Whitehead (1911)

Let's look at the light switch. They are easy to operate without much thought and creating bright, constant, useful light can be considered an important operation. The function of a light switch is straightforward and easy to explain, so light switches represent functional simplicity. But that functional simplicity requires quite a lot of non-simple structure, like a distant power source, a network of cables to transmit the power, various protections and fail safes to prevent dangerous surges or shorts, and an electrician to setup the switches in the house, among other things, in order to make turning on a light as simple as flicking a switch.

Before all of the structure of power grids and standard light switches was established, turning on a light involved a lot of steps, each of which would have been fairly easy and straightforward, but which required more work and a particular order of operation and which allowed opportunities for failure. Whitehead equates the advance of civilisation with the collapse of several functions, each requiring effort and attention, into one simple function by means of non-simple structures, mechanisms, procedures and tools to sequence, execute and monitor all of the sub-functions.

Simpler function can be very positive, worthy of motivating Whitehead's admiration, as when complex structure provides simpler function in medical devices (Burdulis et al. 2010). Manufacturing and industry have also benefited from the development of complex structure and complex processes which allow for impressive improvements in performance and waste reduction (Aldrich and Whetten 1981; Schonberger 1982). These benefits of simple function arise from bundling multiple necessary structures and functions together so that the relations between the required parts and steps can be more automatic, integrated and efficient, but simplicity of function is not always desirable. For example, Miller (1993) describes how over time, organisations and businesses increase the complexity of organisation and the simplicity of function to their own detriment by maintaining no longer necessary or successful structures. Further, when faced with complex situations and choices, people fall back on simple but suboptimal options (Iyengar and Kamenica 2010). The tendency to maintain complex, but no longer optimal, structures or to choose

suboptimal solutions when faced with complex choices is related to the irreversibility and path dependencies created in adaptive systems. An adaptation that moves up a hill may be beneficial, but reaching the top of a hill always carries a certain risk of becoming trapped without any chance of further movement.

Structural Simplicity The simplest possible structure is a whole, without internal divisions or parts. Unfortunately, nothing known to man is truly whole and non-divisible. Atoms, derived from the Greek word *ατομοζ*, meaning indivisible, are a prime example of the lack of indivisibility, as they are now known to have internal parts which can be split, exchanged and recombined.¹³ The closer we look at what we assume to be whole or structurally simple, the more internal divisions we find. It appears to be turtles all the way down,¹⁴ although at the Planck length, we lose the ability to distinguish where one turtle ends and the next begins.

Cheerfully choosing to ignore the internal components, we can call the structure simple when it is reasonable to do so, which of course is entirely dependent on the observer and the context of observation. Water, for example, does have internal structure, but the arrangement of hydrogen and oxygen atoms is arguably irrelevant if we want to look at whether we need more or less water for growing tomatoes, so we can ignore the internal structure and treat water as simple.

However, the simpler the structure, the more likely it is to have many functions. Water, rocks, planes, wheels, and other simple structures are useful, even necessary, for loads of functions. Rocks, for example, can be projectiles, weights, tokens, building materials or pets, when googly eyes have been glued on, and many other things as well. These simple structures appear to have simple functions, but they are usually only one part of a complex structure to perform a larger function. For example, rocks are only useful as tokens in a complex system of symbolic representation including ideas of value, delayed reciprocity and the exchange of goods and services, while rocks as projectiles are not often just launched for their own sake, but as part of achieving military aims, hunting, games, sporting competitions or scientific experiments.

Thus, even simple structures, when repeated, linked or incorporated into sequences, quickly become complex. Classic examples include the repetition of simple DNA structures, which generate complex gene regulation functions (Tautz et al. 1986), and languages, where structural simplicity permits such complex functions as expressing everything ever said or written (Ferguson 1968).

Occam's Razor The law of parsimony, often known as *Occam's razor*, is often understood to advocate the simplest explanation when choosing from several competing hypotheses. However, this "simple" rule can make finding the simplest explanation very difficult. Whoever is holding the metaphorical razor must decide whether to favour the explanation with the simplest function or the explanation with

¹³And new divisions for those parts, and for the parts into which they are divided, are always lurking at the edge of theory.

¹⁴http://en.wikipedia.org/wiki/Turtles_all_the_way_down.

the simplest structure, which are unlikely to be the same explanation for any phenomenon of interest. In fact, most of the competing hypotheses will be almost impossible to compare for simplicity when both functional and structural simplicity are included, especially if the measure includes all of the associated and supporting structures and all of the alternative or possible functions.

Greenhouse Example A structurally simple greenhouse, essentially just a transparent enclosure, is not specific to any particular plant species, so has multi-functional potential and can be used to grow almost any kind of plant, or even multiple types of plants at once. There is even potential to add chickens or mushrooms for increased, complex function. But such a simple structure requires quite a lot of additional, albeit also simple structures, such as watering cans, harvesting tools, baskets or barrows for carrying the produce, and quite a lot of manual labour to operate. A modern high-tech tomato greenhouse on the other hand, which comes with specifically spaced racks, CO₂ transport tubes, temperature sensors, lighting, heating, aeration, pollination and watering systems, tomato harvesting robots and automated transport trucks with safety systems can only accommodate tomatoes.¹⁵ The increasingly complex structure means that it has the simpler, more specific function of “growing tomatoes”, which it can do extremely efficiently and effectively, but requires far more complex internal structure to achieve it.

As a second example, let’s compare an average smartphone vs Casio Model AS-C. The Casio AS-C only calculates a few limited mathematical operations, while the smartphone has many, very diverse functions including a calculator. From the perspective of a single function, perhaps calculating the amount everyone must pay toward a shared bill at a restaurant, they are equivalent in functional simplicity. They can both divide the total cost of the bill by the number of people at the table, so the smartphone will look far more structurally complex, even annoyingly so as you might have to go through several menus to get to the calculator function while the Casio AS-C need only be plugged in. However, if all of the functions available on a smartphone are considered, then it appears far simpler to carry one small device in place of carrying a series of simple but separate devices such as a telephone, calculator, camera, mp3 player, laptop, etc.

2.4.2 *Complicated*

Before we go on to discuss what complexity is, we want to first introduce a special type of complexity, known as complicated. Complex and complicated are both non-simple, but the important distinction between them is not an inherent quality, but is

¹⁵Modern greenhouses can of course be adapted to grow other crops. Some of the structures and systems would become completely redundant, others would need to be added, and still others would require some adjustments. Like the structurally simple greenhouse, there is potential for more, but the path dependency means that it will be much more costly to switch to a new crop if there is a risk that high cost investments, like a tomato picking robot, will become utterly useless.

a matter of process, change and experience. George Whitesides¹⁶ argues that simple things are:

- reliable, predictable;
- cheap (money, energy, etc.);
- high performance or value/cost; and
- stackable, able to form building blocks.

His view of simple does not distinguish between structural or functional simplicity, instead highlighting elements of each. Reliable and predictable, for example, apply far more to simple structures¹⁷ than to complex structures, but, like the light switch, even complex structures can become predictable, cheap, high performers which can be used as elements in larger structures.

As Whitehead remarked earlier, the hallmark of civilisation is the change in effort required to do something important. This move is not from the complex to the simple but from the complicated, or both structurally *and* functionally complex, to the simple, where either the structure or the function is judged as simple. Many new, cutting edge, experimental or unfamiliar socio-technical innovations are not predictable or reliable, not cheap to make or use, not high performing and not useful as part of larger systems. These systems, made of many parts, involving long sequences of actions, demanding training, and requiring a high level of vigilance to maintain or control, can not be judged as structurally or functionally simple, and so are complicated. Complicated systems are argued to be more difficult to understand than complex systems (Allen et al. 1999), possibly because, as brand new entrants to a fitness landscape, the direction of motion and the effects of any steps taken are uncertain.

What counts as a complicated system is highly observer-dependent.¹⁸ Cars and airplanes are generally considered to be good examples, but others might be the constantly changing rules in Formula One racing, the unwritten and ever shifting norms of fashion, etiquette, and cool music, the fluctuating and labyrinthine financial regulations, or the steady stream of upgrades to high tech software programs. These examples are clearly not structurally simple as they have thousands of rules, elements, subsystems and interacting parts, the removal, misuse or malfunction of any of which could result in a non-operational vehicle, the imposition of a ten second penalty, a social faux pas, astounding financial losses (or gains, if you can spot the loopholes), or surprising software bugs. But in addition to being structurally complex, they are also functionally complex because they require significant and unending effort, attention, training, vigilance and maintenance. Driving a car or taking advantage of the best corporate tax schemes is just not as easy or effortless as flicking a light switch.

¹⁶http://www.ted.com/talks/george_whitesides_toward_a_science_of_simplicity.html.

¹⁷How often do really simple structures fail? A rock, for example, tends to be flung from a catapult with a high degree of reliability.

¹⁸What else did you expect?

The epistatic relations between the many structural elements of complicated things mean that the parts, even if optimised individually, may perform in unpredictable and sub-optimal ways when put together, so any new design will be complicated. But when complicated systems, as a whole, are used over a long period of time, subject to relatively constant selection pressures in steady conditions, they become seamless, effortless and almost invisible as a part of the background of everyday life. Engineers or other experts often idealise either the functions or the structure of complicated systems as relatively simple, downplaying the importance of variations, interactions, or the effect of the environment to see the system as more isolated, mechanistic, and with out any surprising behaviours or flaws. While these idealisations are useful in the design process, to the unfamiliar and non-engineers, the complicated systems remain bafflingly complex in every way. Yet as a design matures, it is tested in many new ways, is subtly refined, grows more prevalent, reliable and inexpensive, performs better in relation to the value, and can serve as a component in larger, newer and more complicated systems. Structural elements may be removed if they can be reliably expected to exist in the environment, simplifying the structure, or the structure may adapt to be more complex, but better integrated, so that the function is simpler. Even if the structure or function does not actually change much, the familiarity of continued use will affect the observer-dependent judgements so that either the many parts come to be viewed as a simple unit or the functions come to be viewed as straightforward and easy.

Greenhouse Example Automation is at the leading edge of an increase in greenhouse complication. Automatic tomato and flower picking robots can not only pick the produce, but also place it on automated conveyor belts which feed into automated packaging machines. These systems have many specific parts, such as conveyors, chutes, clamps, sensors, motors and robotic arms. When working correctly, operation is very efficient, but there has simply not been enough time yet to test every possible situation or combination of factors. Consequently, not all of the “bugs” have been worked out and the failure of something as simple as a ball bearing in one part of one subsystem has the potential to disrupt the entire operation.

However, automated watering systems used to be cutting edge, full of bugs and with enormous potential for risk, but are now seen as commonplace. The continued use has resulted in predictable, inexpensive and highly valuable performance, so they are no longer complicated, but functionally simple parts of a larger complexity.

2.4.3 Complex

Finally, we come to discuss complexity. We already know quite a bit about complexity because real world systems are inherently complex, despite the tendency to idealise them as isolated, mechanistic and fully knowable. Adaptive systems are even more complex because the relationships between elements are so pervasive, subtle, and impermanent. Change is inevitable, and every change rewrites the rules

of the game in some small way. We also know that complex is the opposite of simple, but that new, complicated additions to the system adapt themselves and force the adaptation of their environment until they are so embedded that they look simple when viewed in the right way. And finally, we know that viewing in the right way is key to seeing a whole made of parts or a part in a larger whole, as complex or simple, or as caused by rules or as the cause of those rules. The importance of views makes complexity so infuriating impossible to define because, as Mikulecky (2001) states:

“Complexity is the property of a real world system that is manifest in the inability of any one formalism being adequate to capture all its properties. It requires that we find distinctly different ways of interacting with systems. Distinctly different in the sense that when we make successful models, the formal systems needed to describe each distinct aspect are *not* derivable from each other.”

Formalisms, or formal systems of capturing statements, consequences and rules, that are not derivable from each other, such as mathematics and psychology, capture different truths about a system. To really describe a complex system more than one formalism, incompatible as they are, must be employed because only the multiple viewpoints of different formalisms can come close to seeing the system as a whole and a part, complex and simple, the cause and the effect, *at the same time*. Checkland goes even further to say that “human activity systems can never be described (or ‘modelled’) in a single account which will be either generally acceptable or sufficient” (Checkland and Checkland 1999, p. 191). Knowledge from various domains and disciplines must be integrated to begin to describe the properties and behaviour of a system in a more adequate, acceptable and sufficient way. As most people master a limited number of disciplines and formalisms, the increased attention to complex systems will demand an increase in interdisciplinary cooperation, although every account will still face criticism, probably from those whose formalisms or models were not included, as insufficient, which must be balanced against the increase in complexity from the inclusion of additional formalisms.

Dynamics One important truth of complexity is that it happens in many dimensions at the same time, and one often overlooked dimension is time. Many attempts have been made to understand why we have the complexity we see in the world, especially as complexity involves so much apparent simplicity due to the balances between complex structures and complex functions over time.

Smith and Szathmáry (1997) argued that important transitions in how information is transmitted represent the key points in the development of complexity, typically when simple structures with multiple functions developed more structural complexity by apportioning each new structure with fewer, simpler functions. Allen et al. (1999) on the other hand, suggests that simple structures with simple functions multiply and compound until reaching a critical point after which the parts restructure into a hierarchy. They consider the addition of new structures at any level to be an increase in complicatedness but that the increase in levels, or the deepening of the hierarchy, to be an increase in complexity. Although they use complicated in




Fig. 2.4 Power law observed in the edit frequency per user on wiki.tudelft.nl

a different way than we have so far, their use also supports the idea that new additions to a system are complicated but that over time the system adjusts and this complicatedness disappears into the total system complexity.

Self-similarity or Scale Invariance Fractals are non-Euclidean, irregular, geometric structures where each individual part is, at least approximately, a reduced-size copy of the whole fractal. This recursive self-similarity is true of complex systems because they are nested, with each level being the lower level of a larger system, or the higher level comprised of smaller systems. But not only are complex systems self-similar, or scale invariant, in structure, but also in behaviour, so that the same patterns, shapes and proportions hold true of the output of the system, no matter the scope of the perspective. An example that you might find outside your front door would be the formation and propagation of cracks and tears in concrete slabs which follow power law proportions, also known as the Pareto distribution, the 80/20 rule or long/fat tail. The frequency of the cracks varies as a power of an attribute, such as the size of the cracks. Thus, very large cracks are relatively rare, numerically overwhelmed by the small cracks, yet the big cracks, rare as they are, overwhelm the total size of all the small cracks added together. This relationship is true if you look at the cracks in just one square meter or at all of the cracks in the entire street. Given the ubiquity of complex systems, it should come as no surprise that these scale invariant relationships are observed in a wide range of phenomena, from craters on the moon to the distribution of wealth in an economy and edits in wikis, see Fig. 2.4.

But scale invariance works in the dimension of time as well as space, so that the relationships between the frequency of an event and some attribute of that event hold the same relationships at any time scale. For example, avalanches occur at any size, from the catastrophic collapse of entire hillsides to the tiny movement of small clumps of earth or snow. The likelihood of an avalanche is in power law proportion to the size of the avalanche, so the relationship between the frequency and size of avalanches observed is the same regardless of whether you look at data for a year, ten years or 10,000 years. Other examples of scale invariance in time include the frequency and duration of network outages on the Internet, the frequency and

number of journal articles citations, considered in the network of all citations among all papers, and the distribution of word use in natural languages.

Greenhouse Example Greenhouses display scale invariance in several ways. As greenhouses are nested systems, they reveal self-similarity at various levels. For example, a greenhouse growers association might have committees devoted to certain topics or aspects of greenhouse operation. When moving down to large individual greenhouses that belong to that association, you might find an individual manager devoted each of those same topics, replicating the structure of the next level up. In smaller greenhouse, there might not be managers dedicated to one topic or aspect of operations, but there will be some replication of the structure, even if it is only that paperwork is clustered into file folders in roughly the same divisions as the committees at the association level.

But greenhouses also display scale invariance in behaviour, space and time. Greenhouse owners, like all people, show power law distributions in the number of contacts they maintain and the frequency of use of those contacts, so that a very few greenhouse growers have large networks of contacts, while most greenhouse growers have far fewer contacts, and that each grower uses some of their contacts very frequently, and the remaining contacts only very infrequently. Further, these relationships are held if you look at any scale, from the local and present to the national and historical. Many aspects of greenhouse operation or behaviour, from the size of greenhouses to the energy use of the devices inside the greenhouses, reveal the complexity of the greenhouse systems by displaying scale invariance and self-similarity.

2.5 Complex Adaptive Systems

Putting the three previous sections together gives us complex adaptive systems, which John H. Holland (Waldorp 1992) defines as:

“[...] a dynamic network of many agents (which may represent cells, species, individuals, firms, nations) acting in parallel, constantly acting and reacting to what the other agents are doing. The control of a complex adaptive systems tends to be highly dispersed and decentralised. If there is to be any coherent behaviour in the system, it has to arise from competition and cooperation among the agents themselves. The overall behaviour of the system is the result of a huge number of decisions made every moment by many individual agents.”

This interest in and acceptance of complexity heralds the rise of a new paradigm, and researchers are aware that there is something important going on, although there is not yet a consensus as to what exactly it is. As a new scientific paradigm, complex adaptive systems is a lens for looking at the world and the way it operates that allows, even requires, a multitude of perspectives and formalisms. No single approach or description will be adequate to capture the richness of complex adaptive systems interactions at many dimensions and across several levels, creating dynamic emergent patterns from local interactions between system components (Holland 1996; Kauffman and Johnsen 1991; Newman 2003).

Although we have already been treating greenhouses as examples of complex adaptive systems throughout the chapter, they are also socio-technical systems with both physical and social co-evolution. As complex adaptive systems is displacing older paradigms that equated understanding with simplification, explanation with a single description, and strict isolation between physical and social elements of systems, engineers need to be aware, but not overwhelmed or discouraged by the links between the physical and social. This means that the design of artifacts now faces new dilemmas, as technologies are seen to be influenced by as well as influence the people who interact with or use the technology.

Greenhouse Example Greenhouses are uncomfortably hot places, and the work that needs done nevertheless requires a high level of physical exertion. As a result of incalculable factors, this is currently perceived as undesirable work, reducing the supply of readily available labour. As a consequence, automated or robotic systems are increasingly attractive to greenhouse growers as they do not feel stigmatised by societies disdain for sweaty conditions or heavy lifting. But the use of tomato picking robots, automated packaging machines and self-guiding product carriers reduces, but has not yet eliminated, the need for human employees. The fewer jobs remaining are now more monotonous, lower skilled, and more isolated, further suppressing demand and wages for greenhouse jobs and reinforcing the perception that these are unappreciated, difficult and low paid jobs.


Did the designers of automated greenhouse systems consider the effect their innovations might have on the social elements of the socio-technical system? Would the results be different if they had? Or are attempts to directly influence the non-physical aspects impossible or unethical anyway?

Of course, co-evolution is intractable. We will never know what would have happened if society had instead perceived a hot, sweaty, physically demanding greenhouse job as a noble, enviable and rewarding position. Robots might be far less attractive if young and strong people competed for the chance to spend their days lifting weight and sweating out skin impurities while being paid to produce food for the benefit of all of society. But again, unpredictable as complex adaptive systems are, engineers in this alternate reality would be busy devising some other tools that influenced wages, labour or other social aspects of the system, with perhaps no net difference.

2.5.1 Chaos and Randomness

One of the basic mechanisms at play in all complex adaptive systems, and one of the reasons we will never know if engineers in that alternate reality could really have a net effect on greenhouse horticulture as a sector, is chaos. While not all chaotic systems are complex adaptive systems, all complex adaptive systems contain chaotic elements. Chaos is a large field of study, that we will not attempt to cover exhaustively here. Instead we will highlight and discuss the main points relevant

Fig. 2.5 The Mandelbrot set fractal



to socio-technical systems. For more background please refer to Gleick (1997) or Kellert (1993).

Chaos can be defined as complex behaviour,¹⁹ arising in deterministic, non-linear dynamic systems, when relatively simple processes or rules are repeatedly applied. Chaotic systems display, among others, two special properties:

- sensitive dependencies in initial conditions
- characteristic structures

Repetition Chaos arises from the repetition, iteration or recursion of simple rules, formulae, processes, or mathematical functions, such as logarithmic maps or fractals. For example, repeatedly reevaluating the complex function $Z = Z^2 + Ci$ gives rise to the Mandelbrot (Mandelbrot 1983) fractal, as seen in Fig. 2.5. The iterations of the simple processes of selection, replication and variation, allow for chaos to develop in adaptive systems, driving some of the complexity of complex adaptive systems.

Deterministic The state of any dynamic system changes over time, according to some rule or procedure. If the changes have no trace of randomness and are instead completely controlled by the rules or procedures, then the system is deterministic, meaning that a given cause always has a clear and repeatable effect. Fractals, both deterministic and chaotic, are not in any way random. In fact, randomness is far harder to produce than it seems as randomness is completely without cause and contains absolutely zero information. No known model is capable of producing true randomness, and your computer cannot produce an authentically random number.²⁰ The only suspected source of true randomness in the universe is the decay of radioactive atoms driven by quantum fluctuations (Green 1981).

Initial Conditions But if chaos is deterministic and randomness is not the source of complexity in complex adaptive systems, why are they intractable and unpredictable? While being non-random and fully deterministic, the iteration of rules on

¹⁹Chaos is a complex behaviour, but chaos is not the only mechanism driving the complexity of complex adaptive systems, nor is chaos the same as complexity.

²⁰Get your truly random numbers here: <http://www.fourmilab.ch/hotbits>.

the system magnifies the minute differences between two starting conditions, potentially leading to very different outcomes and the appearance of unpredictability. Often referred to as the *butterfly effect*, this sensitivity to initial conditions means that seemingly insignificant differences, such as the rounding off of numbers in calculations, tiny errors in measurements, or a change in what appears to be a totally unrelated factor, can be sufficient to set the system into a different state, making them appear to act randomly and without reason. With only finite information on the starting conditions, the exact state of a chaotic system cannot be predicted with any certainty, and the uncertainty grows with the distance of the forecast. This is why specific weather predictions beyond a week are no better than guesses.

Attractors Characteristic structures, the second property of chaotic systems above, means that chaotic systems tend to converge towards certain points or regions in the systems state space, over time, called attractors. Thus, while sensitivity to initial conditions means that the exact states cannot be predicted, attractors mean that some very large sets of initial conditions converge on a single chaotic region and this convergence can be predicted with some reliability. Usually, system outputs contain multiple attractors, and some contain repellers as well, which the system seems to be unable to approach. Dynamic systems will have attractors that shift over time, displaying varying intensities and duration of attraction, all influenced by the complexity and adaptations of the system.

These are the same attractors that form the hilltops in fitness landscapes. In adaptive systems, the attractors are a consequence of selection pressures that act on differences in fitness, but non-adaptive systems have attractors too if the rules of the system interact in such a way as to create one or more basins of attraction.

Instability and Robustness When a chaotic system suddenly changes from one attractor to another with only minimal parameter changes, it is called instability. There are many examples of instability, where an apparently consistent and predictable system appears to suddenly change gears and move with a unsettling inescapableness along what previously appeared to be an unlikely path. The human heartbeat, for example, displays a change from one attractor to another when it rapidly increases after hearing a sudden noise. Large crowds are notoriously unstable, suddenly erupting into riot, moving in new directions, and forming crushes or stampedes under certain conditions that do not appear to be very different from the conditions in which crowds of the same size act peaceably. In structural engineering, a structure can become unstable when an applied load crosses a threshold and the structural deflections magnify stresses, which in turn increases the deflections.

Instability can be seen as the opposite of robustness. Note that robustness is not the same as stability, as systems can be simultaneously robust and not stable.²¹ A system is robust when it is close to or at an attractor because very few parameter changes can cause a deviation from the path to the attractor (Callaway et al. 2000).

²¹The terms stability and instability seem like each others opposites, when in fact instability is opposite robustness.

However, robustness is not a general concept, as large changes in some parameters cannot make the system deviate from its path to an attractor, while only very slight changes in another parameter might cause the system to change to another attractor entirely. Robustness is a measure of how the system performs under stress, when confronted by extreme inputs or shocks from the environment, only for particular variables. The Internet, designed to function even if large parts of it are destroyed, is robust against physical attacks, power outages and disruptions, but is weak against sudden rushes of net traffic which can spread the disruption from overloaded sites. Body temperature is robust against changes in temperature of the environment, but can be disrupted by illness, anxiety or even stress. A less positive example is economic lock-in, when a customer is so dependent on a supplier for products and services that the switching costs of moving to another supplier outweigh the benefits.

The combination of chaos, instability and robustness are important concepts for complex adaptive systems and socio-technical systems. Importantly, robustness and instability need to be seen in relation to specific parameters, and to be viewed in context, as the ability to change suddenly or to resist large changes can both be seen as good and bad. Any attempt to engineer, shape or steer a complex adaptive system must be careful to analyze exactly which parameters the system is robust or unstable in relation to, as large intentional changes in the wrong parameter can have little or no effect, while small, accidental and unintended ones can dramatically affect the system.

Greenhouse Example Greenhouses in the Westland often use combined heat and power units to produce heat, electricity and CO₂. At certain times of the day, they produce more electricity than they need, so sell this extra power back to the regional electricity grids, making the power generation capabilities of the region distributed and therefore robust against a catastrophic loss of power. On the other hand, sometimes the greenhouses need electricity, but cannot help producing superfluous heat and CO₂ as well. This wasted production contributes to the total greenhouse gas emissions of the region, which appear to be very resistant to all efforts at reduction, indicating that a very strong attractor makes the system robust to changes that would result in a net decrease in CO₂.

Instability in the greenhouse horticulture sector is readily apparent in the prices paid for products, which can change suddenly and drastically. While the prices paid for the goods can change suddenly, switching from one attractor that drives prices up to another that drives them down, the invested time, effort and money that has already gone into producing the flowers or vegetable is far less unstable, so a mismatch between production costs and selling prices is always a risk. Unfortunately, the limited shelf life of the products means greenhouse farmers are unable to wait for prices to improve.

2.5.2 Emergence, Self-organisation and Patterns

Emergent behaviour or emergent properties are overall system behaviour of a complex adaptive system. Emergent behaviours contain *no magic* because they are only

the motion toward attractors, or away from repellers, although they are rarely obvious or predictable. Instead, the apparently magic new characteristics or phenomena are only the logical consequences that become apparent once the organisational structure and interactions of the system are constituted (Crutchfield 1994; Morin 1999). These phenomena *cannot be deconstructed solely in terms of the behaviour of the individual agents* (Jennings 2000) and *would not arise if isolated from the organising whole* (Morin 1999). Indeed, the emergent properties of systems are lost when the system is broken down into parts and parts removed from the system lose the emergent properties they previously possessed.²² Although an emergent property cannot be found in any of the component parts, nevertheless, emergent properties can appear or disappear with the gain or loss of a single element, depending on where the emergent behaviour is in relation to the various attractors of the system, as chaotic systems are always instable or robust to particular parameters. Examples of familiar emergent behaviours include traffic jams, for which it makes little sense to examine the actions of individual cars, schooling, swarming or flocking behaviours in social animals, which generally have no centralised control and yet behave cohesively as a group, or stock markets, which aggregate the actions of many traders, all of whom have limited knowledge and operate under regulations, yet lead to wildly different results from one day to the next.

Emergent behaviour tends to be easier to recognise or simpler to understand—and potentially more insightful—than the collection of processes that cause it, leading many emergent phenomena to remain as “black boxes” to observers. Human consciousness, for example, is argued to be an emergent behaviour (Dennet 1996) and cannot be understood in terms of the individual parts of the brain. Institutions, governments or corporate boards display unpredictable, emergent system output when establishing policies because “a decision is an outcome or an interpretation of several relatively interdependent streams within an organisation” (Cohen et al. 1972), each with incomplete information and unclear priorities (Lindblom et al. 1980). Economic literature also suggests that externalities are undesired emergent properties, although not all externalities are entirely negative, as when neighbourhood house prices are increased through the many, distributed actions of dedicated home gardeners. Emergent properties are what we look for when studying socio-technical systems and the evolution of these systems. Although perhaps not often framed as such, most of the decisions we take in life, and certainly the decisions made by authorities, are geared toward bringing about or enhancing desired emergent properties, like sustainability, while preventing the undesired ones, such as pollution.

Greenhouse Example The selling price of a given greenhouse horticultural product is emergent, and depends on the interaction between costs of growing substrate, greenhouse technologies, labour, and energy, but also on the behaviour of other tomato growers, the demands and leverage of supermarkets or other retailers, the

²²Consider, for example, how the emergent property of living is lost if an organism is dissected, and how an amputated part ceases to be alive after removal.

shopping and consumption patterns of consumers, the implementation of legislation, the effect of foreign competition, new food fads and many other things. The prices are not determined by any centralised agent, nor can the final price be attributed directly to any one cause or action. Were the current system to be dismantled and all the parts isolated, for example if no communication or delivery of product were allowed between growers, retailers and consumers, not only would the prices not be set as an emergent property, but the entire concept of price would have no meaning. The participants in the system change all the time, as greenhouse companies start up, close, change hands or switch to new business models, as do the retailers, the consumers and the legislators, so the system is robust to the loss of any single participant. Yet it could all collapse entirely, or change beyond recognition, if some element, like the capacity for refrigeration were to be removed, or some as yet unknown new element were to be added that gave fresh products an unlimited shelf life. As it is, the interactions that determine the selling prices are contingent on the expected shelf life of the products and a sudden change in that would rewrite the rules entirely.

Self-organisation A particularly important and interesting form of emergent behaviour, is self-organisation, the process by which a system develops a structure or pattern without the imposition of structure from a central or outside authority, or when a system displays a different output as a result of internal processes (Prigogine and Stengers 1984; Kay 2002). For example, in morphogenesis, an embryo develops toward a fully functional organism by self-assembling from a single fertilised cell (Campbell 2002) while autopoiesis means that societies develop and impose limits to individual choice, which provides a more predictable, self-steering system (Luhmann 1995).

Structure and organisation can be very beneficial, durable and self-reinforcing, so self-organisation can be an adaptive response. But self-organisation also occurs in systems that are not generally considered to be adaptive, such as crystal growth, galaxy formation, micelles and cellular automata. Thus, self-organisation is not adaptive on its own, but is potentially adaptive, depending on the environment and the current state of the system. Influencing socio-technical systems and the evolution of these systems seeks to match up self-organisation and adaptation as much as possible. Self-organising behaviours that are also adaptive to the pressures we want apply to a system are essentially “for free”, allowing relatively small modification of system components and their interactions to achieve a great degree of the desired organisation and regularity in the system.

Patterns Patterns are merely something we observe as standing out in contrast with background “noise”. Living organisms have evolved the ability to detect regularity, even to the point that we are sometimes unable to *not* see a pattern in coincidences (think about conspiracy theorists), to ignore a change from one pattern to another (as when we can’t sleep in a new place because we miss the sounds of traffic that we are used to), or to avoid tapping our foot when a catchy tune comes on. The regularity in a pattern makes it possible to compress it for more efficient understanding, storage or transmission. The pattern also allows for better than chance

predictions about what comes next, which is impossible if there is no regularity in the data.

Computers, on the other hand, do not see patterns unless told to look for them, instead taking all the input as a whole, which is more time consuming, but less likely to produce faulty analysis. However, life being limited by time, evolution favours the efficient storage and transfer of information (DNA and natural languages are both good examples), faster response times (better not to wait until you see a predator if you could recognise the pattern of footsteps behind you instead) and the ability to develop predictions or hypotheses (“food has been found here this time last year, so could be found here again now”), even at the expense of potential loss of accuracy or the risk of false pattern detection. As patterns involve repetitions, they appear quite often in nature and evolving systems, and the detection of these repeated patterns also occurs regularly.

But these patterns do not just appear in nature, they emerge. That is, the regularity of interactions in the system leads to repetition of properties, behaviours, and structures, which are all detected as emergent patterns or as organisation in a dynamic system. The different system levels are apparent to us, not because they are true and real distinguishable levels, but because we see regularities in the interactions, over time or space, that lead us to observe a pattern. Thus, any pattern detected is highly observer-dependent and would not be seen, or would be seen slightly differently, by another observer with different access to the information or a distinct world view. As the system continues to evolve, the system levels may appear to break apart, grow, shrink or add additional levels as novel patterns of interaction emerge. And as the patterns emerge, they can serve as the input for the emergence of new patterns because the system itself capitalises on any regularities present, amplifying the patterns across different system levels.

2.6 Modelling Complex Adaptive Systems

Knowing all of this about systems, adaptations and complexity is great, but how can we use this knowledge? One proposal is that through modelling these systems in light of the principles of complex adaptive systems, we can better understand the specific systems and how to interact with them in order to achieve goals. In this section we will first discuss aspects of modelling complex adaptive systems in general, before moving to the theory behind agent-based modelling. Practical aspects of the creation of agent-based models are described in Chap. 3 in great detail.

2.6.1 *What Does a Model of a Complex Adaptive System Need?*

Models are the formalisation of a modeller’s interpretation of reality, and so are not one, but *two* steps removed from the real world. The challenge therefore is to

carefully balance two important, yet conflicting, needs when modelling a complex adaptive systems. The model must be complex enough to represent the system as well as possible and it must be as simple as possible in order to facilitate a greater understanding or the ability to change the system.

The first requirement is formally expressed by Ashby's Law of Requisite Variety (Ashby 1968), a commonly used formulation of which states:²³

“a model system can only model something to the extent that it has sufficient internal variety to represent it.”

Thus, to be a successful model *of* a complex adaptive systems, the model must also *be* a complex adaptive systems. This need for accuracy and complexity in the model is directly at odds with the other need, which is to be simple enough to give insight into the system that is not attainable through observation of the system itself. Thus, the model must simplify reality in order to be useful, but not so much so that it is useless. Just as a map is a useful simplification and miniaturisation of a place, a model must be gainfully simplified by, among other things, defining system boundaries, level of observance and context.

“All models are wrong, some are useful ...”

Box (1979)

As already discussed above, each model is two-fold simplification of the reality, and, as a consequence, is wrong. However, even when wrong, a model can still be useful if the simplifications are only where appropriate to achieve the task at hand. A model will not be useful if it ignores crucial aspects of the real world system just because they are difficult or ideologically unpalatable. Likewise, replicating too much detail at a very low level, or refusing to include essential details from lower levels are examples of how greedy reductionism or extreme holism are unhelpful oversimplifications and can undermine the system representation. Therefore, we must find that tricky balance between the accuracy needed to reproduce complexity and the simplicity needed to gain any novel insight by following some clever advice. It has been said that the usefulness of a model can be estimated by the speed by which it is replaced. The models that provide the most insight and teach us the most tend to be replaced the fastest as a consequence of their high utility. Thus, we should follow the advice of a clever man when building models:

“Everything should be made as simple as possible, but no simpler.”

Attributed to A. Einstein

In order to do so, we have found that every complex adaptive systems model should contain the following three main properties:

Multi-domain and Multi-disciplinary Knowledge Although any particular model can only be considered a single formalism, they can be used in multi-domain and multi-disciplinary ways to capture multiple formalisms. For example,

²³<http://pespmc1.vub.ac.be/REQVAR.HTML>.

they can be developed with insight from multiple viewpoints, fields of study or experts with varied interests. They can also be used in context with work founded in non-derivable formalisms for a balanced approach to the inherent complexity of the topic, or as one part of a series of models, with each one incorporating new aspects that make it a slightly different formalisation.

Generative and Bottom up Capacity The central principle of generative science is that phenomena can be described in terms of interconnected networks of (relatively) simple units, and that finite, deterministic rules and parameters interact to generate complex behaviour. Most of generative science relies on the idea that “If you did not grow it, you did not explain it!” (Epstein 1999) and thus seeks to “grow” a given macroscopic regularity from an initial population of autonomous agents or to explore the range of behaviours a well understood, well described population of agents is capable of under different conditions. While not every behaviour that can be grown is necessarily also explained, the generative science approach means that if done well and founded on a rich theory of complexity and complex adaptive systems, modellers can attempt to build understanding from the bottom up.

Adaptivity The model must also be adaptive, with a capacity to evolve over time. There must be selective pressures to respond to, and a way to introduce variations for the pressures to act on. Ideally, the selective pressures should also be capable of shifting in response to the changes in fitness, making the entire system adaptive, although this is far more difficult to analyze and interpret and may push the balance of the model toward capturing the complexity of the system at the expense of insight that might be gained by simplifying the adaptivity.

Modelling Options Reviewing the many modelling techniques available would be outside the scope of this work. Some, like statistical thermodynamics, and pattern recognition tools such as neural networks, are unsuitable for modelling complex adaptive systems, as they are clearly not generative. Others, such as computable general equilibrium (Jones 1965; Leontief 1998), dynamic systems (Rosenberg and Karnopp 1983; Strogatz and Henry 2000), and system dynamics (Forrester 1958; Forrester and Wright 1961), are based on mathematical models based on a top-down paradigm and on a assumption of static system structure. Discrete event simulation (Boer et al. 2002; Boyson et al. 2003; Corsi et al. 2006; Gordon 1978) comes closer, being capable of both generative and dynamic system behaviour, but fails to suit our needs as the entities are very passive representations and cannot quite capture some of the necessary decision making. Agent-based modelling (Jennings 2000; Rohilla Shalizi 2006), however, has everything we need with its explicitly bottom-up perspective. The individual agents, whose algorithmic nature allows many different formalisms, act and react according to internal rules to produce the over all emergent system behaviour.

2.6.2 Agent-Based Modelling

Of the presented tools, agent-based modelling is the most suitable for modelling a complex adaptive systems because it is the only one that satisfies Ashby's requirement. In the words of Borshchev and Filippov (2004):

The "agent-based approach is more general and powerful²⁴ because it enables the capture of more complex structures and dynamics. The other important advantage is that it provides for construction of models in the absence of the knowledge about the global interdependencies: you may know nothing or very little about how things affect each other at the aggregate level, or what the global sequence of operations is, etc., but if you have some perception of how the individual participants of the process behave, you can construct the agent-based model and then obtain the global behaviour".

Before we get into the nitty gritty of exactly what agent-based modelling is and does, we should explore a bit of its past. The first inklings of distributed computation that later came to underpin agent-based modelling appeared in the 1940s when John von Neumann conceptualised the Von Neumann machine (Von Neumann and Burks 1966), a theoretical device capable of self-replication using raw materials from the environment.²⁵ The notion was further refined by Ulam with the creation of the computer implementation which he called *cellular automata* (Burks 1970).

Constraints on computer power at the time meant that cellular automata remained as mere mathematical curiosities until Conway published his "game of life" (Conway 1970), a 2D cellular automata. The game of life demonstrated the extremely broad spectrum of behaviour that could arise from very simple rules. Thomas Schelling's segregation model (Schelling 1971) further advanced the possibilities. Although initially, played on a paper grid with coins, the segregation model showed some aspects of complex adaptive systems which were more clearly realised after it was later transformed into an agent-based model.

These early frontrunners paved the way for diverse and numerous explorations of the possibilities as computational power grew throughout the 1980s, including Robert Axelrod's prisoners dilemma model (Axelrod 1980) and Craig Reynolds' Boids (Reynolds 1987), a bird flocking simulation. The 1990s saw the spread of such models coinciding with an increase in the ready availability of vast amounts of computational power, while tools like Swarm, NetLogo and Repast lowering the programming barrier, allowing the agent-based modelling field to grow explosively. Currently, the field can boast a number of dedicated journals and with a series of high level articles appearing in journals such as Nature (Buchanan 2009; Farmer and Foley 2009) and The Economist (Economist 2010).²⁶

²⁴than System Dynamics, Dynamic Systems or Discrete Event Simulation.

²⁵It is interesting to note that Van Neumann's idea is close to becoming a reality some 60 years later with the rise of open source 3D printers, which are currently able to build 90 % of themselves.

²⁶Perhaps the economic crisis of 2008 could have been avoided or minimised if a complex adaptive systems approach was more widely used or appreciated. Although, we will never know, irreversibility being what it is.

Computer power is now ramping up and the theories of systems, complexity and generative science are gaining traction in science. Although previously too difficult to reason about, much less model, agent-based models provides a new tool for the exploration of these topics. The focus is on the interactions of the agents, which Stuart Kauffman says is “a thing which does things to things” (Rohilla Shalizi 2006). Furthermore, Rohilla Shalizi (2006) states that:

“An agent is a persistent thing which has some state we find worth representing, and which interacts with other agents, mutually modifying each other’s states. The components of an agent-based model are a collection of agents and their states, the rules governing the interactions of the agents and the environment within which they live.”

Another perspective is provided by Tesfatsion (2007):

“In the real world, all calculations have real cost consequences because they must be carried out by some entity actually residing in the world. ACE²⁷ modelling forces the modeller to respect this constraint. An ACE model is essentially a collection of algorithms (procedures) that have been encapsulated into the methods of software entities called ‘agents’. Algorithms encapsulated into the methods of a particular agent can only be implemented using the particular information, reasoning tools, time and physical resources available to that agent. This encapsulation into agents is done in an attempt to achieve a more transparent and realistic representation of real world systems involving multiple distributed entities with limited information and computational capabilities.”

2.6.3 *What It Is and Is not*

So agent-based modelling is a method, or approach, which examines the interactions of “things” or “entities” rather than a particular thing or collection of things to be replicated. While modellers may have an instinct for what is or is not an agent-based model, the reality is that it falls, along with several related fields that also focus on interacting things, on a spectrum with fuzzy boundaries and confusing overlaps. Before proceeding to examine our agents, it is important to differentiate where it lies on the spectrum in relation to other “thing-centric” fields, and whether the distinctions between the concepts are useful for a given investigation.

Agent-Based Model *What happens when ...?* Agent-based models are constructed to discover possible emergent properties from a bottom-up perspective. They attempt to replicate, in silico, certain concepts, actions, relations or mechanisms that are proposed to exist in the real-world, in order to see what happens. Generally, agent-based modelling has no desired state or task to be achieved, instead merely describing the entities and observing how they interact in order to

²⁷ Agent-Based Computational Economics; essentially agent-based modelling with agents containing economic decision models.

explore the system's possible states. An agent-based model can examine how farmers might adapt to climate change (Schneider et al. 2000), the co-evolution of autocatalytic economic production and economic firms (Padgett et al. 2003) or the behaviour of an abstract economy (Kauffman 2008). The model acknowledges that reality consists of many components acting, relatively autonomously, in parallel, and that no specific predictions can be made, but that patterns, tendencies, and frequent behaviours shown in the model may be relevant to the real world. While an agent-based model generally has no set state to achieve, replicating some real world phenomenon to a desired degree of accuracy means that some models become less about seeing what happens and more about seeing what it takes to make something specific happen.

Multi-agent System *How can I make a . . . ?* Multi-agent systems are often, but incorrectly, used interchangeably with agent-based modelling because they also use discrete, parallel, autonomous components (or agents) to examine system emergence. The main difference is that agent-based modelling sets up agents believed to have crucial characteristics of real world analogs to see what happens when they do whatever they do, while in a multi-agent system agents are set up with exactly the characteristics, connections and choices that they need to achieve certain desired emergent states. For example, it be used to develop process design using component collaboration and local information (Hadeli et al. 2004), the design of an advanced e-commerce agent (Lee 2003), a predictive control system for transportation networks (Negenborn et al. 2006), and the design of cooperative agents in a medical multi-agent system (Lanzola et al. 1999). A multi-agent system is an attempt to control emergent problems, like traffic control or agenda synchronisation, that are not best solved by top down approaches but must resolve all conflicts (i.e. no traffic jams or conflicting appointments). While usually trying to solve a given problem rather than replicate the behaviour of actors in real world situations, it can sometimes look quite a lot like an agent-based model if they problem to be solved involves exploring the unpredictable behaviour of human-like agents.

Artificial Intelligence *Can he do this . . . ?* Artificial intelligence can be seen as zooming in on the agent. Consciousness, learning, object detection and recognition, decision making and many other facets of intelligence can be considered emergent properties and artificial intelligence researchers are attempting to replicate these, much as agent-based models seek to replicate the emergent properties of industries, economies and cultures. Although often studied in isolation, groups of artificial intelligence agents, usually called distributed artificial intelligence, would be a return to the level of zoom that allows for emergent properties between agents instead of only within agents. This would be almost indistinguishable from a multi-agent system if the distributed artificial intelligence agents were trying to solve a particular problem or achieve a certain state, as when teams of intelligent, problem solving robots try to play a game of football. Alternatively, if distributed artificial intelligence agents are instead left to their own devices while researchers observe their output, like any communicative systems they might develop, then they start to look a lot like an agent-based model (Honkela and Winter 2003).

Object-Oriented Program? Rohilla Shalizi (2006) states that:

“While object-oriented programming techniques can be used to design and build software agent systems, the technologies are fundamentally different. Software objects are encapsulated (and usually named) pieces of software code. Software agents are software objects with, additionally, some degree of control over their own state and their own execution. Thus, software objects are fixed, always execute when invoked, always execute as predicted, and have static relationships with one another. Software agents are dynamic, are requested (not invoked), may not necessarily execute when requested, may not execute as predicted, and may not have fixed relationships with one another.”

In essence, agents may be built with object oriented programming software, and when given very simple rules and limited behavioural options, they behave very like objects. But at heart, agents are designed to be unlike normal objects because they flagrantly ignore the usual programming goal of eliminating repetitive or unnecessary elements. By having multiple and similar agents or components, many of which may not have any actions or whose actions seem ineffective, pointless, counterproductive or irrelevant, a simulation using agents cannot be elegant, streamlined or minimal code. However, as with all bottom-up approaches, the messy, repetitive, unexpected relations are important, and the surprisingly concise results and solutions can only be seen as more than the sum of the parts.

2.7 Anatomy of an Agent-Based Model

Our dissection of an agent-based model begins with a schematic overview, presented in Fig. 2.6, followed by a detailed description of the Agent, its states and its behaviour rules, before a look at the Environment. Finally, we detail the structure and organisation of agent interactions and aspects of time. While very theoretical at this stage, these notions will become concrete in Chap. 3, which discusses in detail the process of creating a model with the anatomy described here.

2.7.1 Agent

Agents are reactive, proactive, autonomous and social software entities, a computer program or “an encapsulated computer system that is situated in some environment, and that is capable of flexible, autonomous action in that environment in order to meet its design objectives” (Jennings 2000). Agents are:

1. Encapsulated, meaning that they are clearly identifiable, with well-defined boundaries and interfaces;
2. Situated in a particular environment, meaning that they receive input through sensors and act through effectors;
3. Capable of flexible action, meaning that respond to changes and act in anticipation;
4. Autonomous, meaning that they have control both over their internal state and over their own behaviour; and




Fig. 2.6 Structure of an agent-based model

5. Designed to meet objectives, meaning that they attempt to fulfil a purpose, solve a problem, or achieve goals.

(adapted from Jennings 2000).

The agent is the smallest element of an agent-based model, the atomic element of a generative theory, and some would even say that the “agent is the theory”. An agent is able to perform actions on itself and other agents, receive inputs from the environment and other agents, and behave flexibly and autonomously because, as shown in Fig. 2.6, an agent consists of both states and rules.

2.7.1.1 State

An agent’s state is the specific collection of parameters that defines an agent (Wooldridge and Jennings 1995), or all of the relevant information about what this agent is at this moment. The internal, local and global states of each agent, any of which can be static or dynamic, contribute to its overall state.

The internal state belongs to the agent and only the agent, and covers the current values of all the possible properties that an agent could possibly have. An agent representing a greenhouse would have an internal state composed of all the values for current growing capacity, energy use, owner, and financial balances, among many other possible properties, while a light switch agent would have one of two possible internal states: on and off. The internal state can be private, public, or a mixture, if

only some properties are observable by other agents or if a property is observable by only some other agents.

However, an agent's actions are also dependent on the actions and inputs from others with which it interacts. Thus, the local state consists of the internal state (private and public) plus all of the publicly observable states of the agents that our agent is interacting with. This puts the internal state into a context and gives some sense to the values, allowing the agent to act based on not only its internal state but the relationship that internal state has to the immediate surroundings.

Finally, the global state is comprised of the internal state, the local state and all of the relevant states in the whole of the observable or influencing environment. With these three states, you can already see how every agent is a complex system, embedded in nested networks of influence that act on various time scales and levels of interaction. The agent uses its internal, local and global states as the basis for applying behavioural rules in order to produce actions.

2.7.1.2 Changing States

Rules Rules, or the “internal models” (Holland 1996) of agents describe how states are translated to actions or new states. Rules should be understood as mechanical decision rules or transformation functions, rather than the more colloquially used social notions of rules as regulations or agreements.

The rules of agents in models of complex systems are usually based on an assumption of rationality or bounded rationality (Simon 1982). For example, a common decision rule might be that agents attempt to maximise some utility, but the agent may or may not have access to information about the other agents with which they interact, may or may not be able to record the outcome of previous actions in order to learn, or may have limits on the computation allowed to process any information in order to mimic the limits that human decision makers face. Decision rules specify what an agent will do with the information that they have access to, as well as how they will perform any actions. Rules can be static or dynamic, and may depend on the internal, local and environmental states. Importantly, agents could choose not to perform an action, either because the rules allow for inaction or because the rules call for probabilities, noise or random elements that alter the normal actions. There are several types of decision rules often used in agent-based modelling. These are:

Rule based Rule based decision rules are the most common, usually in the form of nested if-then-else structures. These are the most common type of decision rules, as they are very easy to implement, and directly couple observed behaviour into decision structures.

Multi-criteria decision making Another common technique used for agent decision making is multi-criteria decision making. It is a technique that allows different choice options to be compared, for example by assigning weights, enabling the agents to have preferences or probabilities. For example, a greenhouse agent might weight the emissions of a CHP unit more heavily than the price

when making purchasing decisions, among other factors, resulting in purchase choices that are not readily obvious.

Inference engines Also known as expert systems, inference engines take facts (states) and decision heuristic to construct a decision tree in order to reach conclusions about which action should be taken. Such systems are often used when an agent needs to base a decision on a lot of real world data, and are often found in engineering, financial and medical applications.

Evolutionary computing When agents need to find a optimal solution in a very complex or large solution space, genetic algorithms can be employed. Agents generate a large number of solutions, evaluate their fitness against a fitness function, select a group of “best” solutions and apply genetic recombination on them in order to make better ones. Such techniques can be very computationally expensive.

Machine learning Neural networks can be used when an agent needs to make decisions based on patterns. Neural networks work as classifier systems, allowing the agent to determine into which category a observed pattern falls, and therefore decide which action is appropriate. Neural networks are also computationally intensive and may require a training period, to allow the agent to learn the categories and which actions are best for each category, before any decisions can be used.

Actions Actions are the actual activities that agents perform based on the application of decision rules on their states. For example, if a light switch agent has a rule that says “If state = off and time = 8:00, set state = on”. Thus, the rule uses the agent’s own internal state (off or on), a global state (the time) and the decision rule to perform an action that changes its own internal state.

Actions can also be directed at other agents. For example, a tax collecting agent might have a rule such as “For each greenhouse grower agent, if yearly profits exceed 10000, deduct 10 %”. This agent then would consider none of his own internal states, but would look at the states of other agents (local), some sort of calendar (global) and act directly on the financial balances of the grower agents (the internal state of another agent).

Of course, agents can also *not* act, which can be understood as an action as well.

Behaviour The agent behaviour is the overall observable sum of the agent’s actions and state changes. It is an emergent property caused by the interaction of the internal, local and environmental states and the decision rules. Overall system (or model) behaviour is an emergent property of the interactions between all of the agents behaviours and the environment.

Greenhouse Example This incomplete example of the agent anatomy of a greenhouse model is drawn from a complete model of greenhouses developed in a step-wise fashion in Chap. 3.

In this model, there is only one type of agent, representing a greenhouse grower. These agents have internal states composed of the values of various properties, such

as what technologies they currently own, what opinions they hold about those technologies, how much money they have, how much profit they earned last season, and what kind of crops they produce. Some of these, such as what technologies they own, are public, while others, such as the exact amount of money owned, are private. The local state also contains the neighbours with whom the agent has communicative links, and the publicly available properties of those agents, such as what technologies they own. The global environment further consists of some environmental properties, such as the price of electricity.

The greenhouse agents have rules that govern how they purchase technologies and how they form opinions about them. For example, agents have a rule like “When a currently owned technology expires, purchase a new technology in the same technology class which has the highest opinion and which is lower in cost than the current money owned.” Thus, when a greenhouse agent’s heater breaks, they purchase a replacement heater, and that heater is the best they can afford, according to their own opinions of which technologies are best.

The actions an agent can take include selling crops, buying technologies, updating account balances, and updating their opinions of technologies. The agent behaviour includes not only the actions they take, but also the changes in technologies, opinions and account balances over the course of the simulation run.

2.7.2 Environment

Agents must be somewhere, and that somewhere gives the agents input through sensors and receives the output or effects of the agent action. This somewhere, in which the agents “live”, is the environment, and it contains all the information external to the agent used in the decision making processes and provides a structure or space for agent interaction. The environment contains everything, including other agents, that affects an agents, but which is not the agent itself, which means of course, that an agent’s environment is context dependent.²⁸ Agents can affect the environment and be affected by it as a consequence of the specific rules they use for actions.

2.7.2.1 Information

Because the environment, in a strict sense, is context dependent, we will use a looser, less context dependent definition here. Rather than consider each agent’s environment uniquely to see how they differ, much the way internal states differ between agents, we focus on the similarities, so that the environment is a shorthand for the information, structure and goings on in the global states. Thus, we can say that the

²⁸Ah, our old friend, context dependency. In this case, it means that no two agents will have exactly the same environment, because every agent will be in the environment for other agents, but not for himself.

environment provides all the things an agent needs to know, and all the ways to do the things that it does, and that are not contained in the agent itself or in its immediate neighbours. The environment will have some elements provided by the model itself, while others could be set by the modeller, or can be emergent.

Environmental information provided by the model itself tends to be quite dull, although crucial. Things like the passage of time, which is so important that we discuss in Sect. 2.7.3 below, is an environmental aspect provided by the model that agents use to make decisions. Many rules use elements of time to make decisions or dictate actions. For example, an agent might have a rule that says “Every turn, make and sell a variety of products. After 10 turns, review the best and worst selling products, adjust the probability of making those products, eliminate the worst selling and introduce a brand new product.” The environment provides the measure of how many turns, or time ticks, have passed, providing information on whether this turn is a turn for making and selling, or for the added action of reviewing sales and production.

When an aspect of the environment is provided by the modeller, it can be considered a model parameter. These parameters can be static or dynamic, and whole sets of these parameters are often considered scenarios that the modeller is interested in. For example, agents might use some global variable, such as temperature, to make decisions. The modeller can set value of the temperature to make a static, global variable. If the modeller runs the model several times, each time with a different temperature, then each model run represents a scenario, or experimental condition under testing. The modeller could also design the model so that the temperature is dynamic, either varying across a set range, following historical data that has been fed into the model, or determined by a probability function. Different ranges, different sets of data, or different probability functions would represent different scenarios for experimentation.

Aspects of the environment could be provided by emergent properties of the model itself rather than directly by the modeller. Emergent properties arise through interaction between agents in the simulation, and may or may not be directly accessible to individual agents, depending on how they must be calculated and whether or not the calculations rely on private or public states. For example, market prices are emergent, and no single agent can determine the prices individually, because agents cannot usually access the supply or demand of particular products, nor the willingness of other agents to pay a given price for a product. Nevertheless, the distributed effects of many private actions and calculations results in a specific price per product, and agents may use that price to determine actions.

2.7.2.2 Structure

The environment also provides the structure in which the agents are situated, which can also be static or dynamic and which can be set by the modeller or derived as an emergent property of the model. If it is not relevant to the decision making rules, agents need not know anything about that structure because the environment will

provide only the information needed by the agents when they need it. For example, a rule might be “Ask all your neighbours if they have any tomatoes to sell. If they do, buy the cheapest”. The agent can simply ask the environment “Who are my neighbours?”, and then use the answer to ask those other agents about their tomato availability. If it does not matter to his decision rules, then he will never know if he has more or fewer neighbours than another agent or if the neighbours are the same as the last time he asked for tomatoes.

Although the agent may not know what structure in which he is situated, that structure can significantly affect the model performance. In agent-based modelling, we can distinguish four main types of structures investigated so far, a unstructured soup, a regular space, a small-world network and a scale-free network. The different structures and network topologies each share some characteristics with human social networks, which may or may not be important, but also influence many things, from the time it takes the simulation to reach some important behaviour to the amount of processor power that is needed to run a simulation for a given number of steps. Furthermore, multidimensional networks are possible so that a set of agents can be connected in more than one way at the same time so that communications flow along some connections, money along others, physical materials along yet others, etc. Choosing one structure over another, or multiple structures, allows the modeller to bring more complexity and realism to the model, an important part of modelling a complex adaptive systems. Nevertheless, the more complex and realistic the structure, the more difficult it is to statistically analyze and the more likely it is to obscure any relations in the model.

Soup A popular and statistically easy to analyze structure is the mean-field organisation, or “soup”, which covers completely random organisations (i.e. agents are equally likely to interact with all other agents) or fully connected organisations (i.e. agents interact with all other agents). A limited number of agents, or a high number of interactions, means that the agents will be fully connected, while a large number of agents, fewer interactions, or a high degree of agent turnover would lead to random interactions, potentially affecting model behaviour. Soups have a short average path length, meaning that there are few steps needed to connect any two agents, which is also true of real-life human networks. But they have a low clustering coefficient, meaning that there are no subsets of agents that interact more frequently with each other than with other agents, which is not true of human social networks. Soups tend to proceed quite rapidly, reaching plateaus, converging on behaviours, or running until stop conditions are reached much faster than other structures, although they typically use more memory resources as well.

Space Another classic structure for organising the agents is the space or regular structure, in which agents are connected to a set of neighbours but not to all the other agents. They are usually arranged in a regular pattern such as a square or hexagonal grid which may be cast onto a toroid to prevent edge effects. Agents can also be situated in a more physically defined space within a GIS map, where neighbourhood is defined in terms of actual distance, a common approach in fields such as spatial

planning and geography. This structure provides a sense of close and far between agents, and because neighbouring agents are connected to the most of the same agents as their neighbours, this structure displays a high clustering coefficient, as do human networks. The high clustering means that local subsets of agents reach convergence or display coherent behaviours quickly. This rapid local convergence can lead to slow global convergence, but reduces the computing power or memory needed to run the simulation. Further, as spaces lend themselves to analysis with the usual methods of statistical mechanics (Baronchelli et al. 2005, 2006) they have been a popular starting point for many researchers.

Small-World Networks Discovered more recently, small-world networks start with a regularly structured space and randomly replace a small number of the local connections with long distance connections. After a very small number of rewirings, the network takes on the short average path length characteristic, and fast global convergence times, of soups without sacrificing the high clustering coefficient, and efficient resource use, of spaces. Despite being discovered quite recently, small-world networks have been subject to much investigation, especially as they relate to regular and random networks. By including both a short average path length and a high clustering coefficient, small-world networks are more realistic than either soups or spaces, but are more difficult to analyze statistically.

Scale-Free Networks Scale-free networks have high clustering coefficients and short average connection length, as do small-world networks, but they also have “hubs” that can dramatically speed up the timescales in the model. These hubs stem from a power law degree distribution,²⁹ which distributes connections, probability to interact, or popularity among the agents such that a very few agents are highly connected while the vast majority have few connections, matching real-life human social networks. The scale-free networks show sharp transitions toward convergence, as do real life populations, and they reach convergence almost as quickly as soups, while using memory resources efficiently. Other structures can be converted to a scale-free network by incorporating growth with some sort of preferential attachment strategy, so scale-free networks can be used in dynamic structure models. While scale-free networks have importantly realistic traits and behaviours, the presence of hubs can introduce problems with robustness and instability if not well considered. For example, if only a few agents are chosen at random to be active each turn, then the overwhelmingly numerous non-hub agents will tend to be selected. If those agents then select one of their neighbours as the recipient of the action (reactive agent), then the hub-agents will tend to be chosen as the recipient quite often, because they are in the set of neighbours of almost every other agent. Thus, depending on the action taken, the hub agents will have experiences and behaviours different than other agents, even if they share all the same decision rules, and may display quite different behaviour. It is not unrealistic, but the implications are not

²⁹Random interaction soups and small-world networks only have a Poisson degree distribution, meaning that extremely popular agents are as common as extremely unpopular agents.

trivial as the divergent behaviour may speed up convergence, promote robustness, or spread instability, depending on how the model works.

Greenhouse Example The environment in the greenhouse model includes some information provided by the model, most notably, time. This is complicated and important enough to have it's own greenhouse example later, so we now move on directly to detail the other informational and environmental factors in this example. Although not directly used as input to the decision making rules in this model, the environment provides the agents with a social structure that they use to interact. The greenhouse growers are placed in a scale free network that mimics the structure of natural human social networks. A few agents are highly connected and acts as hubs for information about opinion on technology, while the majority of agents are in the periphery of the network with a limited number of neighbours. Were it allowed to vary, the structure of the agent interactions would be a model parameter. For this model, however, other structures were not deemed to be realistic scenarios.

Some of the model parameters that can vary, and which are also sources of information from the environment, include how many agents are in the simulation, how big the greenhouse sizes can be, and with what probability the agents are likely to heed the advice of their neighbours when making technology purchase decisions. One possible scenario might be the case of few greenhouse grower agents, ranging in size from 1 to 25 hectares, and with all neighbours opinions valued highly. Another scenario might be many growers, with greenhouse size ranging from 10 to 20 hectares, and with growers paying very little attention to the opinions of their neighbours.

The environmental information also includes information that is not determined by the model itself, nor given as a parameter by the modeller, but which is an emergent property of the model behaviour. In this greenhouse example, agents compare their own profit to the average profit made by all greenhouse in order to determine how satisfied they are with their technology choices. As growers do not have access to the profit of their neighbours, much less to the profit of all the greenhouse in the simulation, they rely on the environment to calculate the average profit from the private information of all growers.

2.7.3 Time

The final aspect of agent-based model that needs to be discussed is the issue of time. Time can be considered a part of the environment, but is so ubiquitous that it demands special attention and unique considerations. Real world complex adaptive systems take place in a continuous real time, and with elements truly acting in parallel. If we are to satisfy Ashby's requirement, we must ensure that these aspects are properly represented in a model, and we must understand them well if we are to represent them.

Discrete Time While reality takes place in real, continuous, time, agent-based models are forced to happen in the discrete time of computers. All conventional computers work with timed instruction clocks, performing rounds of operations within each time step. This reflected by the use of a *tick* as the smallest unit of time. Simulations can play with discrete time by redefining how much time a tick is meant to represent, with no theoretical lower or upper limits. However, if the time needed to compute a single tick is longer than the amount of real time that tick is meant to represent, then this places a particular practical limit on one way that simulations are often used. A simulation that runs slower than reality is not very useful for predictions.

Assumption of Parallelism While the discrete time is significantly different from reality, the main problem is parallelism. Although real world complex adaptive systems are massively parallel, only very recently have multi-core computers enabled the performance of more than one task at a time.³⁰ In order to represent the parallelism of the real world with a serial processing device, all actions are scheduled to occur one after the other, but are assumed to happen at the same time. The disjoint between what actually happens and what is assumed to happen can create significant problems. For example, when simulating bird flocks, each bird constantly observes its neighbours speed and position and adjusts to them. However, one bird has to go first, and he can only observe the *previous* states of all other birds in order to decide how to move itself. Then next bird to move observes the *current* state of the first bird, and the *previous* state of the other birds, and so on through all the birds. One option to deal with this might be that the birds always move in the same order, so that each bird at least has an internally consistent relation between his observations and actions, although his observation-action relation will be unlike that of any other bird. Alternatively, the birds could observe and act in random order, so that no two turns in a row will have the same observation-action relationship, but these will be roughly consistent between all birds. Or the modeller can choose that a tick takes place in two parts, first observing the position of all other birds, and then moving to where they think they need to be. These decisions are not trivial, as they can provide the agents a “glimpse into the future” or purposefully restrict information by only allowing access to the inherently outdated. The explicit management of the order of agent interaction *over* and *within* time is performed by the scheduler.

Scheduler Schedulers are the central controller of a system that *simulates* a system without any central control.³¹ The scheduler progresses the ticks and ensures that all “parallel” actions are executed. Most commonly, this involves randomising the iteration order of agents at each step. If not properly randomised or otherwise controlled, first-mover advantage modelling artifacts may appear. For example, a greenhouse agent has to observe a market, find the best price, and sell all of his

³⁰Multitasking on single core computers is achieved by very quickly switching between tasks.

³¹Finally! Now we know the dirty little secret of agent-based modelling

produce. If the same agent always goes first, then he will always be able to sell his entire stock at the best price that could possibly be offered, giving him a distinct advantage over the poor agent who always goes last and has to settle for whatever nobody else wanted. Depending on the exact implementation of the modelling software, the scheduler might allow for very fine grained control of specifying the order of particular actions of particular agents, making sure that for example market clearing agent always goes last, or that the order of actions of a particular agent are randomised.

Greenhouse Example In the greenhouse model, each tick represents one year. At each tick, agents perform some actions, such as selling produce and balancing their books, and the environment “ages” the technologies. The agents perform other actions, such as buying replacement technologies, only when triggered by one of their technologies reaching its expiration age. Thus, the agents use information directly from the model in the form of ticks, information indirectly from the model, in the form of their technology ages, and their own decision rules to act in time.

Furthermore, each time tick is composed of several phases that all agents perform in a random order before moving on to the next phase. The first phase sees all agents sell their produce and calculate their own profit, at which point the environment calculates the average profit for the round. In the next phase, agents compare their own profit to the average and use the difference to form opinions on the technologies that they own (more profit than the average means they form a positive opinion of their technologies, lower than average means they form poor opinions). After that, agents share their current opinions with neighbours, before incorporating the opinions of their neighbours back into their own opinion of the technologies. Finally, agents check to see if any technologies have expired, and if so they use their current account balance and their newly updated opinions to purchase a replacement.

Importantly, the agents all form their own opinion before sharing any opinions with neighbours. If some agents shared their opinions based solely on their profit compared to the average before other agents got the chance to compare their profits and form an opinion, then the opinions of the first agents to speak would be more influential than later agents. The opinions of the first agents would influence others, who would then repeat their compromised opinions to others, influencing them. If this were totally random, perhaps it would all balance out to prevent first-mover advantage, but by separating the actions of opinion formation, sharing and reconsideration, no opinion has any undue influence from first-mover advantages.

In this model, the representation of time as multiple phases of action per “one year” tick cannot change. Other simulations might give the option to vary the representation of time, the order of interactions, or the representation of parallelism, and thus would be model parameters as part of a scenario. The possibility of varying such fundamental model details as the conceptualisation of time might be considered unique formalisations. However, to truly be a distinct formalisation, any possible variance would need strong justification and thorough theoretical support, so slapping a variable on for time is no substitute for proper multi-formalism.

References

- Aldrich, H., & Whetten, D. (1981). Organization-sets, action-sets, and networks: making the most of simplicity. In *Handbook of organizational design* (Vol. 1, pp. 385–408).
- Alexander, C. (1973). A city is not a tree. *Surviving the city: a sourcebook of papers on urban livability*, p. 106.
- Allen, T., Tainter, J., & Hoekstra, T. (1999). Supply-side sustainability. *Systems Research and Behavioral Science*, 16, 403–427.
- Argyris, C., & Schon, D. A. (1996). *Organisational learning II: theory, method and practice*. Amsterdam: Addison-Wesley.
- Ashby, W. R. (1968). Variety, constraint, and the law of requisite variety. In *Modern systems research for the behavioral scientist*. Chicago: Aldine.
- Axelrod, R. (1980). More effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(3), 379–403.
- Baronchelli, A., Dall'Asta, L., Barratt, A., & Loreto, V. (2005). Topology induced coarsening in language games. *Physical Review E* 73, 015102.
- Baronchelli, A., Loreto, V., Dall'Asta, L., & Barratt, A. (2006). Bootstrapping communication in language games: strategy, topology and all that. In *Proceedings of the 6th international conference on the evolution of language* (pp. 11–18).
- Bijker, W., Hughes, T., & Pinch, T. (1987). *The social construction of technological systems: new directions in the sociology and history of technology*. Cambridge: MIT Press.
- Boer, C., Verbraeck, A., & Veeke, H. (2002). Distributed simulation of complex systems: application in container handling. In *Proceedings of SISO European simulation interoperability workshop* (pp. 24–27).
- Bonen, Z. (1981). Evolutionary behavior of complex sociotechnical systems. *Research Policy*, 10(1), 26–44.
- Borshchev, A., & Filippov, A. (2004). From system dynamics and discrete event to practical agent based modeling: reasons, techniques, tools. In *Proceedings of the 22nd international conference of the system dynamics society* (pp. 25–29).
- Box, G. (1979). Some problems of statistics and everyday life. *Journal of the American Statistical Association*, 74(365), 1–4.
- Boyson, S., Corsi, T., & Verbraeck, A. (2003). The e-supply chain portal: a core business model. *Transportation Research, Part E*, 39(2), 175–192.
- Buchanan, M. (2000). In *Ubiquity: the science of history or why the world is simpler than we think*. London: Weidenfeld.
- Buchanan, M. (2009). Economics: Meltdown modelling. *Nature*, 460(7256), 680.
- Burdulis, A., Fitz, W., Vargas-Voracek, R., Lang, P., Steines, D., & Tsougarakis, K. (2010). *Surgical tools facilitating increased accuracy, speed and simplicity in performing joint arthroplasty*. US Patent App. 12/776,701.
- Burks, A. (1970). *Essays on cellular automata*. Champaign: University of Illinois Press.
- Callaway, D., Newman, M., Strogatz, S., & Watts, D. (2000). Network robustness and fragility: percolation on random graphs. *Physical Review Letters*, 85, 5468.
- Campbell, N. (2002). *Biology*. Redwood City: Benjamin Cummings.
- Checkland, P., & Checkland, P. (1999). *Systems thinking, systems practice: includes a 30-year retrospective*. New York: Wiley.
- Cohen, M., March, J., & Olsen, J. (1972). Garbage can model of organizational choice. *Administrative Science Quarterly*, 17(1), 1–25.
- Colletta, L. (2009). Political satire and postmodern irony in the age of Stephen Colbert and Jon Stewart. *The Journal of Popular Culture*, 42(5), 856–874.
- Conway, J. (1970). The game of life. *Scientific American*, 223(4), 4.
- Corsi, T., Boyson, S., Verbraeck, A., Van Houten, S., Han, C., & Macdonald, J. (2006). The real-time global supply chain game: new educational tool for developing supply chain management professionals. *Transportation Journal*, 45(3), 61.

- Coulson, J., & Richardson, J. (1999). *Coulson & Richardson's chemical engineering*, Stoneham: Butterworth/Heinemann.
- Crutchfield, J. (1994). The calculi of emergence: computation, dynamics and induction. *Physica D*, 75(1–3), 11–54.
- Darwin, C. (1985). *The origin of the species*. Baltimore: Penguin.
- David, P. (2000). Path dependence and varieties of learning in the evolution of technological practice. In *Technological innovation as an evolutionary process* (p. 119). London: Cambridge University Press.
- Dawkins, R. (1990). *The selfish gene*. London: Oxford University Press.
- DeLaurentis, D., & Crossley, W. (2005). A taxonomy-based perspective for systems of systems design methods. In *2005 IEEE international conference on systems, man and cybernetics* (Vol. 1).
- Dennet, D. (1996). *Darwin's dangerous idea: evolution and the meanings of life*. New York: Simon & Schuster.
- Economides, N. (1996). The economics of networks. *International Journal of Industrial Organization*, 14(6), 673–699.
- Economist, T. (2010). Agents of change. *The Economist*. <http://www.economist.com/node/16636121>.
- Epstein, J. (1999). Agent-based computational models and generative social science. *Complexity*, 4(5), 41–60.
- Farmer, J., & Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256), 685–686.
- Ferguson, C. (1968). Absence of copula and the notion of simplicity: a study of normal speech, baby talk, foreigner talk and pidgins.
- Foerster, H. (1972). Perception of the future and the future of perception. *Instructional Science*, 1(1), 31–43.
- Forrester, J. W. (1958). Industrial dynamics—a major breakthrough for decision makers. *Harvard Business Review*, 36(4), 37–66.
- Forrester, J., & Wright, J. (1961). *Industrial dynamics*. Cambridge: MIT Press.
- Funtowicz, S., & Ravetz, J. (1993). Science for the post-normal age. *Futures*, 25(7), 739–755.
- Futuyma, D. (1983). Evolutionary interactions among herbivorous insects and plants. In *Coevolution* (pp. 207–231). Sunderland: Sinauer.
- Gaukroger, S. (2001). *Francis Bacon and the transformation of early-modern philosophy*. Cambridge: Cambridge University Press.
- Geels, F. W. (2004). From sectoral systems of innovation to socio-technical systems—Insights about dynamics and change from sociology and institutional theory. *Research Policy*, 33(6–7), 897–920.
- Gleick, J. (1997). *Chaos: making a new science*. New York: Vintage/Ebury.
- Gordon, G. (1978). The development of the general purpose simulation system (GPSS). In *History of programming languages I table of contents* (pp. 403–426).
- Green, P. (1981). A new look at statistics in fission-track dating. *Nuclear Tracks*, 5(1), 77–86.
- Hadeli, W., Valckenaers, P., Kollingbaum, M., & Brussel, H. V. (2004). Multi-agent coordination and control using stigmergy. *Computers in Industry*, 53(1), 75–96.
- Hartmanis, J., Sewelson, V., & Immerman, N. (1983). Sparse sets in np-p: exptime versus nexptime. In *STOC '83: proceedings of the fifteenth annual ACM symposium on theory of computing* (pp. 382–391). New York: ACM.
- Hietbrink, O., Ruijs, M., & Breukers, A. (2008). *The power of dutch greenhouse vegetable horticulture: an analysis of the private sector and its institutional framework*. Technical report 2008-049, LEI Wageningen UR, The Hague.
- Hix, J. (1996). *The glasshouse*. London: Phaidon Press.
- Holland, J. (1996). *Hidden order; how adaptation builds complexity*. Reading: Addison-Wesley.
- Holling, C. S. (2001). Understanding the complexity of economic, ecological, and social systems. *Ecosystems*, 4(5), 390–405.
- Honkela, T., & Winter, J. (2003). Simulating language learning in community of agents using self-organizing maps.

- Hume, D. (1962). *A treatise of human nature*, vol. 1. Glasgow: Collins.
- Iyengar, S., & Kamenica, E. (2010). Choice proliferation, simplicity seeking, and asset allocation. *Journal of Public Economics*, 94(7–8), 530–539.
- Jablonka, E., & Ziman, J. (2000). Biological evolution: processes and phenomena. In *Technological innovation as an evolutionary process* (pp. 13–26). Cambridge: Cambridge University Press.
- Janis, I. (1982). *Groupthink: psychological studies of policy decisions and fiascoes*. Boston: Houghton.
- Jantzen, D. (1980). When is it coevolution. *Evolution*, 34, 611–612.
- Jennings, N. (2000). On agent-based software engineering. *Artificial Intelligence*, 117(2), 277–296.
- Jones, R. (1965). The structure of simple general equilibrium models. *The Journal of Political Economy*, 73(6), 557.
- Katz, J. S. (2006). Indicators for complex innovation systems. *Research Policy*, 35(7), 893–909.
- Kauffman, S. (2008). *Reinventing the sacred: a new view of science, reason and religion*.
- Kauffman, S., & Johnsen, S. (1991). Coevolution to the edge of chaos—coupled fitness landscapes, poised states, and coevolutionary avalanches. *Journal of Theoretical Biology*, 149(4), 467–505.
- Kay, J. (2002). On complexity theory, exergy and industrial ecology: some implications for construction ecology. In C. Kibert, J. Sendzimir, & B. Guy (Eds.), *Construction ecology: nature as the basis for green buildings* (pp. 72–107). London: Spon.
- Kellert, S. (1993). In *the wake of chaos: unpredictable order in dynamical systems*. Chicago: University of Chicago Press.
- Kim, J. (1999). Making sense of emergence. *Philosophical Studies*, 95(1), 3–36.
- Lanzola, G., Gatti, L., Falasconi, S., & Stefanelli, M. (1999). A framework for building cooperative software agents in medical applications. *Artificial Intelligence in Medicine*, 16(3), 223–249.
- Lee, R. (2003). Ijade surveillant—an intelligent multi-resolution composite neuro-oscillatory agent-based surveillance system. *Pattern Recognition*, 36(6), 1425–1444.
- Leontief, W. (1998). Environmental repercussions and the economic structure: an input-output approach. *International Library of Critical Writings in Economics*, 92, 24–33.
- Lindblom, C., Cohen, D., & Warfield, J. (1980). Usable knowledge, social science and social problem solving. *IEEE Transactions on Systems, Man and Cybernetics*, 10(5), 281.
- Luhmann, N. (1995). *Social systems*. Stanford: Stanford University Press.
- Mandelbrot, B. (1983). *The fractal geometry of nature*. New York: Freeman.
- Mandeville, B., & Harth, P. (1989). *The fable of the bees*. Baltimore: Penguin.
- Mikulecky, D. (2001). The emergence of complexity: science coming of age or science growing old? *Computers and Chemistry*, 25(4), 341–348.
- Miller, D. (1993). The architecture of simplicity. *The Academy of Management Review*, 18, 116–138.
- Morin, E. (1999). Organization and complexity. *Tempos in Science Nature: Structures, Relations, Complexity*, 879, 115–121.
- Munger, M. (2008). Blogging and political information: truth or truthiness? *Public Choice*, 134, 125–138.
- Negenborn, R., De Schutter, B., & Hellendoorn, H. (2006). Multi-agent model predictive control of transportation networks. In *Proceedings of the 2006 IEEE international conference on networking, sensing and control (ICNSC 2006)* (pp. 296–301).
- Newman, M. (2003). The structure and function of complex networks. *SIAM Review*, 45, 167–256.
- Ottens, M., Franssen, M., Kroes, P., & van de Poel, I. (2006). Modelling infrastructures as socio-technical systems. *International Journal of Critical Infrastructures*, 2(2–3), 133–145.
- Padgett, J., Lee, D., & Collier, N. (2003). Economic production as chemistry. *Industrial and Corporate Change*, 12(4), 843–877.
- Prigogine, I. (1967). *Introduction to thermodynamics of irreversible processes* (3rd ed.). New York: Interscience.
- Prigogine, I., & Stengers, I. (1984). *Order out of chaos: man's new dialogue with nature*. Boulder: New Science Library.

- Reynolds, C. (1987). Flocks, herds and schools: a distributed behavioral model. *Computer Graphics*, 21, 25–34.
- Roberts, R. (1989). Serendipity: accidental discoveries in science. In *Serendipity: accidental discoveries in science* (p. 288).
- Rohilla Shalizi, C. (2006). Methods and techniques of complex systems science: an overview. In *Complex systems science in biomedicine* (pp. 33–114). Berlin: Springer.
- Rosenberg, R., & Karnopp, D. (1983). *Introduction to physical system dynamics*. New York: McGraw-Hill.
- Ryan, A. (2008). What is a systems approach? [arXiv:0809.1698](https://arxiv.org/abs/0809.1698).
- Schelling, T. (1971). Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2), 143–186.
- Schneider, S., Easterling, W., & Mearns, L. (2000). Adaptation: Sensitivity to natural variability, agent assumptions and dynamic climate changes. *Climate Change*, 45(1), 203–221.
- Schonberger, R. (1982). *Japanese manufacturing techniques: nine hidden lessons in simplicity*. New York: Free Press.
- Simon, H. (1982). *Models of bounded rationality*. Cambridge: MIT Press.
- Smith, A. (1963). *An inquiry into the nature and causes of the wealth of nations*. New York: Wiley.
- Smith, J., & Szathmáry, E. (1997). *The major transitions in evolution*. London: Oxford University Press.
- Stewart, J. P. (1964). *Jacobellis vs Ohio*, 378 u.s. 184, *Jacobellis v. Ohio*. Appeal from the supreme court of Ohio. No. 11.
- Strogatz, S., & Henry, S. (2000). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. New York: Westview Press.
- Tautz, D., Trick, M., & Dover, G. (1986). Cryptic simplicity in DNA is a major source of genetic variation.
- Teisman, G. *Publiek Management op de Grens van Chaos en Orde: Over Leidinggeven en Organiseren in Complexiteit*. Den Haag, SDU Uitgevers bv.
- Tesfatsion, L. (2007). Agents come to bits: Towards a constructive comprehensive taxonomy of economic entities. *Journal of Economic Behavior & Organization*, 63(2), 333–346.
- Thompson, J. (1994). *The coevolutionary process*. Chicago: University of Chicago Press.
- Ulrich, W. (1988). C. west churchman-75 years. *Systemic Practice and Action Research*, 1(4), 341–350.
- van den Muijzenberg, E. (1980). *A history of greenhouses*. Institute for Agricultural Engineering.
- van der Lei, T. E., Bekebrede, G., & Nikolic, I. (2009). Critical infrastructures: A review from a complex systems perspective. *International Journal of Critical Infrastructures*, 5(4).
- Von Bertalanffy, L. (1972). The history and status of general systems theory. *The Academy of Management Journal*, 15(4), 407–426.
- Von Neumann, J., & Burks, A. (1966). *Theory of self-reproducing automata*. Urbana: University of Illinois Press.
- Waldorp, M. (1992). *Complexity: the emerging science at the edge of order and chaos*. New York: Simon and Schuster.
- Weber, B., & Depew, D. (2003). *Evolution and learning: the Baldwin effect reconsidered*. Cambridge: MIT Press.
- Whitehead, A. N. (1911). *An introduction to mathematics*. Williams and Norgate.
- Wilds, R., Kauffman, S., & Glass, L. (2008). Evolution of complex dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 18, 033109.
- Williamson, W. O. (1987). *The economic institutions of capitalism*. New York: Free Press.
- Wooldridge, M., & Jennings, N. (1995). Intelligent agents—theory and practice. *Knowledge Engineering Review*, 10(2), 115–152.