Chapter 6
Analysis: Frameworks and Theories for Social Simulation

6.1 Overview

The previous chapter focused upon the current methodologies and theoretical implications of agent-based modelling in the social sciences. While many within this growing field accept that agent-based models provide a potentially powerful new method for examining social behaviours and structures, a great debate still continues over the best methods for utilising the strengths of this methodology.

As seen in Part I, such debates are not unique to social simulation. Indeed, artificial life and more conventional forms of biological modelling have faced similar challenges over the past few decades. With this in mind, this chapter begins by placing artificial life within the various theoretical frameworks discussed thus far in Chaps. 3, 4 and 5. In this way the limitations of each framework can be illuminated.

Social science simulation using agent-based models shares a number of constraints and methodological difficulties with biological modelling using the same methodology. Thus, having placed artificial life and biological models within a theoretical framework, social simulation will be subjected to a similar analysis. Finding the most appropriate framework for social simulation will lay the groundwork for Chap. 7, in which one of the more prominent exemplars of social simulation will be subjected to theoretical and methodological analysis. Chapter 7 lays the foundation for the conclusions of Part II, in which our analysis of Schelling's residential segregation model will provide a means to demonstrate the most important elements of a useful modelling framework for the social sciences.

E. Silverman, *Methodological Investigations in Agent-Based Modelling*, Methodos Series 13, https://doi.org/10.1007/978-3-319-72408-9_6

6.2 Frameworks and ALife: Strong ALife

6.2.1 Strong ALife and the Lack of 'Real' Data

As noted by critics of artificial life, social simulation and related methodologies, computational simulations suffer from a perceived lack of 'real' data, or data derived from experimental observation. Part of this inherent difficulty stems from the need for abstraction in many such models; for example, connectionist models of cognitive processes embrace the idea of 'distributed representation' and their potential role in cognition, while generally avoiding integrating those models into larger, more complex neural structures as seen in the brain (Rumelhart and McClelland 1986).

Strong ALife simulations suffer even more strongly from this shortcoming. Ray's Tierra provides an enticing look at a 'digital ecology' composed of evolving computer programmes competing for memory space (Ray 1996), but those creatures are purely artificial constructions. While the parasites and hyper-parasites which eventually evolve in Tierra's world may provide an analogue to real-life parasitic behaviour, the specialised nature of their virtual world is such that analyses of Tierra would be very difficult to apply to the natural world. Ray might argue that his open-ended evolutionary system, which lacks the standard fitness function of many genetic algorithms and instead provides selection only through life and death, evokes real-world interactions in evolving systems. Does the difficulty of verifying such claims confine analyses of Tierra to the realm of mathematical curiosity?

6.2.2 Artificial^1 vs Artificial^2: Avoiding the Distinction

As noted in Silverman and Bullock's description of differing varieties of artificiality in science (Silverman and Bullock 2004), providing a distinction between man-made instances of natural systems and man-made facsimiles of natural systems is important to understanding the goals of a simulation. In the case of strong ALife, researchers aim to produce models that embody Artificial^1, or a man-made instance of something natural; these models claim to create digital life, rather than simply resemble biological life. In this case, the strong ALife researcher falls into the role of a sort of digital ecologist, studying the behaviour and function of real, albeit digital, organisms. Of course, such claims seem remarkable, but in the absence of a complete and verifiable definition of life such claims are difficult to refute.

6.2.3 Information Ecologies: The Importance of Back-stories

The inclination of the strong ALife researcher to study 'real' digital organisms points to the importance of formulating a theoretical backstory for any given simulation model. As per Silverman and Bullock's PSS Hypothesis for Life, presuming that:

1) An information ecology provides the necessary and sufficient conditions for life.
2) A suitably-programmed computer is an example of an information ecology.

...then the strong ALife researcher may claim, under such a framework, that their simulation represents an information ecology, and thus is a digital instantiation of a biological system. Whether or not the low-level functions of that system match those of real-life, carbon-based biological systems is immaterial; the only criterion is the presence of an ecology of information in which genetic material competes for representation, and under this criterion the position stated here is justifiable.

To use our central example, if we construct a bird migration model in Alife fashion using individual interacting agents, we may wish to demonstrate that we are in fact performing some form of empirical data collection, rather than simply investigating an interesting mathematical system. So, we claim in this case that are simulated birds do in fact demonstrate an information ecology; perhaps our agents evolve and reproduce, this producing a dynamic of information amongst the agent population. If we follow Langton, and are willing to class ourselves in the strong Alife camp, then signing up to the PSS Hypothesis for Life may be a good course of action for us to take. In that case, our bird model becomes a true information ecology, and thus presents an opportunity for empirical data-collection in this virtual population of birds.

6.3 Frameworks and ALife: Weak ALife

6.3.1 Artificial^1 vs. Artificial^2: Embracing the Distinction

In contrast to strong ALife, weak ALife faces some initially more daunting theoretical prospects. This seems somewhat paradoxical; after all, the strong ALife researcher seeks to equate digital organisms with natural organisms, whereas the weak ALife researcher seeks only the replication of certain properties of natural life. However, while the strong ALife researcher may justify their investigations into a digital ecology by signing up to an appropriate theoretical backstory, however far-fetched, proving the relation between a natural system under investigation and a computational model based on that system is a more difficult problem and requires more in-depth justifications.

Returning to our central example, recall our researcher who wishes to model the flocking behaviour of migrating birds. While a great deal of experimental data exists upon which one can base such a computational study, the researcher must choose which elements of that data provide a useful background for the simulation and which do not. Data about bird migration varies greatly across different species and climates, and the researcher must identify the most salient forms of collected data to use as a basis for the model. These choices will in turn inform the construction of the model, and the related points of theory regarding migration that must be incorporated into that model.

As Chalmers suggests (Chalmers 1999), these abstractions reveal the inherent theory-dependence of artificial life as an enterprise; the researcher's choice of abstractions may conform to their specific theoretical biases. In order to make these choices effectively, and to draw a strong correlation between the digital system and the natural system, one must find a framework which embraces the inherently artificial nature of such simulations and uses the advantages of the digital medium effectively to draw useful experimental conclusions.

6.3.2 Integration of Real Data: Case Studies

6.3.3 Backstory: Allowing the Artificial

The integration of a suitable theoretical backstory into weak ALife research seems a more difficult task than for strong ALife. As Bryan Keeley describes, weak ALife can only hope to be functionally related to natural life, producing behaviours that are analogous to those displayed in biology (Keeley 1997). However, establishing a clear relationship to a natural system is not always straightforward, particularly when the artificial system under consideration bears less resemblance to the fundamental make-up of the natural world.

For some researchers and theorists, these artificial worlds present a tantalising opportunity to examine characteristic evolutionary behaviours in a simplified environment, one amenable to study; given that natural evolution is far more difficult to observe than an abstracted evolutionary algorithm, this naturally seems an attractive prospect for those who wish to observe evolution in action. Ray (1994, 1996) goes so far as to assert that digital forms of evolution can produce the same diversity of forms that we observe as a result of natural evolution (though perhaps given his statements regarding Tierra as a whole, this position is, for him, milder than most).

Unfortunately for Ray, while Tierra does display an impressive variety of self-reproducing digital entities that display unexpected behaviour, the simulation tends to get caught in an eventual evolutionary cycle in which certain forms repeat. This contrasts strongly with real-world evolutionary behaviour, in which the general level of complexity in evolving populations tends to continue to increase over time. Similarly, Avida and other simulation environments developed by the ALife community suffer the same problem, preventing the community from replicating the sort of staggering diversity seen in nature (Adami and Brown 1994). Bearing this in mind, can the researcher be certain that these artificial evolutionary systems are truly functionally related to natural evolutionary systems? Given that the overall processes of evolution are still under constant debate and revision, and the innate difficulty of providing appropriate selection pressures in an artificial environment, how much do these systems coincide on a 'given level of abstraction' (Keeley 1997).

In cases such as this, one must be careful in defining a theoretical backstory linking the natural system to the artificial. A clear statement of the mechanisms at work in the simulation and how they relate to similar theorised mechanisms in natural evolution seems most helpful; given that there is no fundamental metric to determine just how abstracted a given model is when compared to reality, a clear statement of assumptions made and potential confounds in the simulation (i.e., difficulties in fitness functions and similar issues) could be helpful in attempting to link the simulation results to empirical results.

In the case of our bird example, such a backstory would need to include information about the assumptions made when constructing our simulation. We would have to describe the real-world instances of bird behaviour that we are trying to replicate, and how these real-world instances have influenced our implementation of the model. Where simplifications have been made, i.e. by simplifying the structure of the agents to facilitate computability and simplicity of analysis, we would need to note this fact and mention how these simplifications may change the character of the behaviour observed in the simulation. In the ideal situation, someone reading a paper describing our bird model should be made aware of the shortcomings of the simulation, where it attempts to reproduce real-world bird behaviour and physiology, and where it makes abstractions for the purposes of making the simulation tractable.

6.4 The Legacy of Levins

6.4.1 The 3 Types: A Useful Hierarchy?

As discussed at length in Chap. 4, the Levinsian framework for modelling in population biology appears generally useful for ALife modelling endeavours. After all, Levins seemed intent upon creating a pragmatic framework for constructing biological models, and since ALife often falls within that remit his ideas remain relevant. However, the extended framework developed from Levins in Chap. 4 seems perhaps more useful within the context of artificial life.

With Braitenberg's Law in mind, the concept of a tractability ceiling placed on ALife seems appropriate, if vexing. With ALife systems spanning an enormous variety of biological phenomena, often incorporating versions of vast and complex biological mechanisms such as evolution, the question of analysis becomes of paramount importance. While we may comfortably classify ALife models within Levins' three types with relatively little difficulty, we remain uncertain how analysable those models will prove to be with only that classification in mind.

6.4.2 Constraints of the Fourth Factor

Indeed, the fourth Levinsian factor appears to place some serious limitations upon ALife systems. As noted in Chap. 3, such systems frequently fall into the Type 3 category, oriented as they are toward producing broad-stroke investigations of generalised populations. Applying those results toward real populations becomes increasingly problematic as tractability concerns become important; without reasonable analysis of the data produced by these simulations, the researcher will have great difficulty applying that data to any understanding of a natural biological system. In essence, even with carefully-designed simulation models, this tractability ceiling prevents highly complex simulations from being productive of great insight.

Thus, our Alife-type bird migration model may run into difficulties if we incorporate too many real-world complexities. If we use agents with neural network controllers, for example, then such networks are very difficult to analyse (recall Beer 2003a,b). As modellers we must judge whether the use of such components in the simulation is justified given the increase of complexity and analytical difficulty. The more we incorporate added elements in an attempt to capture real-world complexity, the more we approach the tractability ceiling.

Even assuming that our model manages to balance Levins' three types appropriately while maintaining a reasonable level of tractability, significant problems still remain. With Braitenberg's Law in mind, agent-based models in general suffer a greater disconnect between invention and analysis than other modelling methodologies, leading to a greater chance of producing impenetrably opaque simulations.

If we then apply agent-based methodologies to the field of social science, in which there are already significant difficulties in constructing models based upon strongly empirical social data, then these problems become increasingly acute. Bearing in mind these discussions of theoretical frameworks in relation to ALife and agent-based models in general, we shall examine how these same frameworks impact upon the agent-based social simulations detailed in the previous chapter.

6.5 Frameworks and Social Science

6.5.1 Artificial^1 vs. Artificial^2: A Useful Distinction?

ALife research focuses on living systems and their governing processes, and as such relies upon definitions and theories of life as a basis for enquiry. Life being difficult to define empirically, strong ALife can, as illustrated earlier, attempt to demonstrate a digital instantiation of life itself when life is defined appropriately within the context of that research. However, when expanding beyond processes relating to simple organisms and populations and attempting to model social structures and processes, additional layers of complexity come into focus. Epstein (1999) provides a comparison to the famous 'Boids' model of flocking behaviour:

> Generating collective behaviour that to the naked eye "looks like flocking" can be extremely valuable, but it is a radically different enterprise from generating, say, a specific distribution of wealth with parameters close to those observed in society. Crude qualitative caricature is a perfectly reasonable goal. But if that is one's goal, the fact must be stated explicitly. . . . This will avert needless resistance from other fields where "normal science" proceeds under established standards patently not met by cartoon "boid" flocks, however stimulating and pedagogically valuable these may be. (p. 52–53)

In essence, Epstein presents a position reminiscent of our earlier discussion of Braitenberg: imitation of a system with a model, as in mimicking the behaviour of flocking birds, is far simpler than creating a model system which can generate that behaviour. Further, such models can confuse the research landscape, perhaps claiming to produce more insight than they are fundamentally capable of providing. In such cases a model which seeks this sort of qualitative similarity is not constructed in such a way as to allow any insight into the root causes of that complex behaviour. Epstein implies later in his discussion that in the case of social science, which involves interacting layers of actors in a society, this problem becomes more acute for the computational modeller.

In such a context, any argument for 'strong social-science simulation' seems difficult to justify; in order to accept that a given social model is a digital instantiation of a social system, one would have to accept that the agents have sufficient complexity to interact with one another, generate a social structure, and react and respond to that structure in an appropriately non-trivial manner. There is a significant danger, as noted by Epstein, of a model of a complex social system falling within the realm of a 'crude qualitative caricature.' Social science then may be said to lie within the domain of Artificial^2: something made to resemble something else.

Fortunately, our examination of Luhmann has revealed the problematic nature of building an Artificial^1 social simulation. With our search for a fundamental social theory relying upon developing a new means for hypothesis-testing and social explanation, creating instantiations of 'digital society' is rather less than useful. Thus, while remaining within the domain of Artificial^2 may seem initially limiting, in fact the social simulator is likely to find it of far greater utility than the alternative.

6.5.2 Levins: Still Useful for Social Scientists?

As demonstrated earlier in this chapter, Levins' framework for modelling in population biology is remarkably applicable to today's more modern computational methodologies. Our updated Levinsian framework developed in Chap. 4 and its consideration of tractability presents an account of a concern common to most varieties of computational models: efficient utilisation of computing resources relies on a relatively tractable and analysable problem. However, when considering the application of such a framework to social simulation, some differing concerns come to light.

As Gilbert and Tierna note (Gilbert and Terna 2000), emergent phenomena in social science reflects a certain additional complexity in comparison to the natural sciences:

> In the physical world, macro-level phenomena, built up from the behaviour of micro-level components, generally themselves affect the components.... The same is true in the social world, where an institution such as government self-evidently affects the lives of individuals. The complication in the social world is that individuals can recognise, reason about and react to the institutions that their actions have created. Understanding this feature of human society, variously known as second-order emergence Gilbert (1995), reflexivity Woolgar (1988) and the double hermeneutic, is an area where computational modelling shows promise. (p. 5)

Thus, Gilbert and Tierna contend that agent-based models can capture this emergent phenomena more vividly than other methodologies which, while providing potentially strong and useful predictions of a macro-level system's behaviour, do not provide an explanation of that behaviour in terms of its component parts. Epstein (1999) takes this statement further, arguing that in certain contexts, successful mathematical models may be 'devoid of explanatory power despite [their] descriptive accuracy' (p. 51). Epstein goes on, however, to acknowledge the difficulties inherent in creating an artificial society with sufficient 'generative power, proposing the use of evolutionary methods to find the rule-sets most amenable to complex emergent behaviour in a given simulation.

Levins' framework, which depends upon tractability as a key concern for the modeller, may at first blush seem insufficient to deal with the methodological complexities inherent in the simulation of social structures. After all, the researcher is dealing with multiple interacting layers of complexity, with some of these emergent behaviours relying upon not just reactions to a changing environment, as in an artificial life simulation, but a cognitive reaction to that environment, which can then influence both that agent and others within that artificial society. With such high-level abstractions taking place in these simulations, how might one quantify the realism, generality and precision of a social model?

To clarify, imagine that we have designed and implemented an agent-based model of our migrating bird population. Each agent is capable of moving through a simulated spatial environment, reproducing, and thus evolving through simulated generations. Now, in an effort to capture the social effects at play within a bird colony upon arrival at its destination, we allow each agent to communicate and exchange information with its compatriots. In order to capture this we suddenly need to add all sorts of new elements to our implementation: a means of encoding commuication between individuals, means for choosing the method and frequency of those communications, and deciding how these interactions will affect the agent, its communication partners, and the surrounding environment. How can we capture such effects? How might we model the birds' communications, and how they affect the simulation as a whole? Already we have introduced a number of new factors into the model for which there is little hard empirical information to guide our implementation of these factors. If we cannot identify the level of realism of these new components of the simulation, how is it possible to clarify the position of our new model amongst Levins' four dimensions of model-building?

In essence, without a solid set of criteria identifying the realism of simulated social behaviours and cognition, one becomes reliant on qualitative judgements to decide the validity of a social simulation; in other words, the researcher must determine that the behaviour of that artificial society sufficiently resembles the behaviour of a real society to call it a successful result. Perhaps then social simulations begin skewed away from realism, and further toward generality and precision than other simulation varieties. With realism so difficult to define in a social context, and with the necessary abstractions for social simulation so wide-ranging, the researcher seems best served by striving to provide examples of possible mechanisms for an emergent social behaviour based upon a very general picture of agent-based interactions and communications. Definitive statements regarding which of these mechanisms are responsible for a given social phenomenon may be impossible, but the model could provide a means for illuminating the possible role of some mechanisms in the formation and behaviour of social structures (this being one of Epstein's suggested roles for simulation in 'generative social science').

6.5.3 Cederman's 3 Types: Restating the Problem

If Levins' framework is insufficient to capture some of the methodological complexities particular to modelling for the social sciences, perhaps this framework could be usefully informed by Cederman's own framework of C1, C2 and C3 political models (see Table 5.1). With some examination of this framework, we may be able to draw useful parallels between these three types and those described by Levins.

Cederman's C1 models focus on modelling the behavioural aspects of social systems, or the emergence of general characteristics of certain social systems in a computational context. He cites Axelrod's Complexity of Cooperation as a major foray into this type of modelling, as well as Schelling's residential segregation model (Cederman 2001). This type of model seems related to Levins' L3 models, which sacrifice precision for realism and generality; Cederman's C1 models do not attempt to link with empirical data, but instead seek to demonstrate the emergence and development of behavioural phenomena in a generalised or idealised context.

Cederman's C2 models aim to explain the emergence of configurations in a model due to properties of the agents within the simulation (Cederman 2001). Examples of this methodology include Lustick's Agent-Based Argument Repertoire Model (Lustick 2002, 2006), which provides agents with a complex set of opinions which they may communicate with other agents. The opinions of agents within the simulation can alter through these communications, leading to the apparent generation of social structures within the agent population.

In this case, the comparison with Levins is more difficult; while the agents themselves are designed to be more complex within a C2 model, can this necessarily be pinned down as an increase in either precision or realism? The closest analogue appears to be Levins' L2 models, which eschew realism in favour of generality and precision. While a C2 Cederman model or an L2 Levins model is not concerned with comparison to empirical data, these models do seek to provide a useful framework for describing complex social behaviours in an idealised context, similar to those population biology models alluded to by Levins. In either discipline, the modellers hope to illuminate some of the contributing factors that produce the modelled behaviour, and perhaps stray closer to a useful description of the real-life instantiation of that behaviour than may be expected initially from such an abstracted model.

Cederman's C3 models are the most ambitious in that they attempt to model the emergence of both agent behaviours and their interaction networks (Cederman 2001). He specifically cites ALife as a field which may provide illuminating insights in that regard, given that ALife is concerned with such complex emergent behaviours and structures. As discussed earlier, however, Cederman's categories have no hard borders between them, particularly in the case of C3 models; Cederman himself admits that C1 and C3 can easily overlap. As a consequence, identifying where the C3 models lie amongst a given crop of social science simulations is not always such a simple task, though the C3 categorisation does provide useful context in which to examine how the goals of different varieties of models can affect their construction and implementation.

Once again, though, the comparison with Levins is difficult; allowing for such emergent behaviours as described by Cederman in the C3 categorisation does not correlate directly with Levins' three categorisations. Perhaps the closest analogue here is once again Levins' L3 models, given the focus on emergence of both agent-level and societal-level behaviours; in neither case is the modeller overly concerned with a direct relation to empirical data. Instead the modeller hopes to provide a cogent explanation for the emergence of social behaviours by allowing agents to interact and change due to pressures within the model, rather than due to complex constraints placed upon that model.

6.5.4 Building the Framework: Unifying Principles for Biology and Social Science Models

We have clearly run into a difficulty in comparing Levins and Cederman's respective modelling frameworks. Levins reserves the L1 category for those models which seek a direct relationship to empirical data, leaving generality behind in favour of realism and precision. However, Cederman's 3 types leave out this distinction, instead providing what appear to be two variations on Levins' more general L3 models. Both Cederman's C1 and C3 models appear to leave precision behind, seeking instead to describe social phenomena in an idealised context; C1 models focus purely on the emergence of social structures, while C3 models focus on the emergence of both societal and individual structures and behaviours.

Cederman's view can be used to further qualify Levins' original framework however. Levins himself cites L3 models as the most promising, and his preferred

Table 6.1 One possible Levins/Cederman framework

Modified Levinsian framework	
L1	Precision and realism at the expense of generality
L2	Generality and precision at the expense of realism
L3A	Generality and realism at the expense of precision at one level of simulation
L3B	Generality and realism at the expense of precision at multiple levels of simulation

methodology within population biology; similarly, Cederman cites his C3 models as the most promising within political and social science (though interestingly, Cederman's recent work has strayed towards a version of Levins' L1; see Cederman 2006 and the section below). The question then becomes: how can we characterise Cederman's C1 vs. C3 distinction in terms of the Levinsian factors of generality, precision and realism?

One can envision a further subdivision of Levins' L3 into a L3A and L3B: L3A being characterised by a sacrifice of precision in one level of the simulation (i.e., Cederman's C1 which seeks emergent social behaviours), and L3B being characterised by a sacrifice of precision at multiple levels (as in Cederman's C3, which seeks both emergent social behaviours and interaction networks); Table 6.1 provides a summary of this potential framework.

Alternatively, with the incorporation of tractability into the Levinsian framework as a fourth factor as in Chap. 4, this subdivision may be much simpler. Both Levins and Cederman acknowledge these L3 models to be both the most useful and the most challenging; perhaps then Cederman's C1 and C3 may be characterised by differing levels of tractability. In this way we can maintain this modified Levinsian framework as a more fluid continuum, based upon the determining factor of overall tractability, rather than introducing additional sub-categories for special cases of specific methodologies.

6.5.5 Integration of Real Data

Some researchers involved in computational modelling have attempted to sidestep the difficulties of the inherent lack of 'real data' by attempting to integrate experimental data into their simulations. Cederman's 2006 paper on geo-referencing in datasets for studies of civil wars provides a useful example. In this case, Cederman uses data and maps from the Russian Atlas Narodov Mira, an atlas produced by a 1960s Soviet project aiming to chart the distribution of ethnic groups worldwide. The project then seeks to formulate agent-based models using this data to examine the potential causes of ethnic conflict. As Cederman notes, 'there is no substitute for real-world evidence' when attempting to understand the causes of such conflict; however, the ethnographic data in this case is both limited and quite old.

Of course the worldwide ethnographic distribution has likely changed significantly since the publication of this atlas in 1964, and updating the information contained therein is no simple task. The integration of such data into an agent-based computational model seems like a potentially fruitful method for tying the results of that model more closely to political and social reality. However, with the limitations of this dataset and the difficulties inherent in collecting future data of a similar type, is this data integration still useful as a framing mechanism to place this model on more solid empirical ground, or is it an interesting but ultimately misguided enterprise?

A further difficulty with the Atlas Narodov Mira dataset is that the atlas provides a static ethnographic picture. While it is a remarkably detailed look at a particular time in world history and the related distributions of peoples throughout the globe, the lack of similar data in the following decades leaves us with an inability to directly associate the intervening political and social changes with ethnic conflicts that have erupted in the years since the atlas' publication. Some argue that computational modelling in such a circumstance provides a remarkable capacity for hypothesis testing; by basing a relatively realistic model on such a solid footing of experimental and observational data, the researcher can experiment with varying parameters and initial conditions in an attempt to replicate the ethnic conflicts seen since the atlas was produced. However in such a case the same problems return to haunt the researcher: deciding which abstractions to make can be critical, and deciding how to formalise the influence of complex social and political interactions is far from trivial.

6.6 Views from Within Social Simulation

6.6.1 Finding a Direction for Social Simulation

While the last chapter provided an in-depth examination of the current state of social simulation, and an analysis of its ability to explain and interpret real-world social phenomena, we still have relatively little understanding of the perception of the utility of social simulation within the social sciences. A number of prominent social simulators display great enthusiasm for the pursuit, as would be expected, but how do those viewing this growing trend from other parts of the field react?

With this in mind we will examine some views of the general purpose of social simulation, and descriptions of the perceived problems of the methodology, from within the social sciences. We can then incorporate these analyses with our own discussion of the importance of simulation in a fundamental social theory and further expand our growing methodological and theoretical framework for social simulation.

6.6.2 *Doran's Perspective on the Methodology of Artificial Societies*

Doran in his article for Tools and Techniques for Social Science Simulation provides a coherent examination of the major difficulties facing the artificial society methodology in social simulation (Doran 2000). Doran describes this method as follows:

> The central method of artificial societies is to select an abstract social research issue, and then to create an artificial society within a computer system in which that issue may systematically be explored. Building artificial societies is a matter of designing and implementing agents and inter-agent communication, in a shared environment, and existing agent technology provides a wealth of alternatives in support of this process. (p. 18)

Thus, Doran views social simulation as a means of examining the workings of large-scale, abstract social issues. For Doran simulation provides a way to generate new 'world histories,' allowing the researcher to watch a society grow from a provided set of initial conditions and assumptions.

Of course he notes the significant methodological difficulties inherent in this artificial society approach. He identifies three main problems: searching the space of initial conditions and parameters; describing the results of the simulation in a tractable way; and overcoming problems of instability due to changes in initial conditions. The first and third problems here are highly reminiscent of more general problems common to agent-based modelling as a whole, while the problem of analysis and tractability harkens back to Levins and our initial theoretical explorations of artificial life.

Interestingly, however, Doran posits that the greatest problem facing the social simulator is the largely undefined nature of the computational agent itself. He notes that different uses of agent-based models often incorporate different base properties for the agents within the model, and that despite the simplistic views of what constitutes an agent within social science, there is little overall consensus regarding a definition of agent architectures.

He argues that agents should be defined in computational terms, and thus should be able to emerge from a model in the same way as other more complex phenomena. Of course, this is not the normal course in most agent-based models; in practice such models are designed to include predefined agent structures based upon certain theoretical assumptions. As a consequence, this would not be an easy task for the modeller; constructing a simulation in which agent structures are expected to emerge seems extremely difficult. There would almost certainly need to be some sort of precursors to defined agents to allow such structures to develop, given that the earliest beginnings of our own life and society are so murky to begin with and can provide little clear inspiration. The purpose, however, is sound, as producing simulations that at least approach such possibilities would allow the modeller to step away further from the difficulties produced by theory dependence and pre-defined agent structures.

Doran's view is also interesting in that it meshes with our earlier discussion of Niklas Luhmann and the search for a fundamental social theory. Since in Doran's view, artificial societies aim to examine the general development of human society in an abstract fashion, these societies must be based upon valid assumptions. However, those assumptions are necessarily based upon our own cultural experiences, and thus will be imprinted upon that model in some fashion. A model which eliminates this difficulty, or at least minimises it, would allow for agents to develop in simulation which bear far fewer markings of our own societal preconceptions. For the social scientist looking to develop new social theory, such an approach would be far more fruitful than the heavily theory-dependent alternative.

As discussed in the previous chapter, Doran agrees that artificial societies should strive to develop from the earliest beginnings of societal interaction ('the lowest possible level, even below the level of agent' in Doran's phrasing [p. 24]), and that the mechanisms of society should emerge from that basis. This would avoid producing a simulation constructed around pre-existing assumptions from our own society.

6.6.3 Axelrod and Tesfatsion's Perspective: The Beginner's Guide to Social Simulation

Axelrod and Tesfatsion (2005) lay out a brief guide to social scientists hoping to incorporate social simulation into their current work. Given the authors' prominent position in the current social simulation field, this guide provides an illuminating look at those aspects of agent-based modelling perceived to be the most valuable by those within this area of research.

After beginning with a brief introduction to the basic characteristics of agent-based models (including a brief discussion on creating 'histories,' as described by Doran), Axelrod and Tesfatsion lays out a four-part description of the potential goals of simulation models in the social sciences. They describe each of these potential goals in turn:

1. Empirical understanding: 'ABM researchers seek causal explanations grounded in the repeated interactions of agents operating in specified environments.'
2. Normative understanding: 'ABM researchers pursuing this objective are interested in evaluating whether designs proposed for social policies, institutions, or processes will result in socially desirable system performance over time.'
3. Heuristic: 'How can greater insight be attained about the fundamental causal mechanisms in social systems?'
4. Methodological advancement: 'How best to provide ABM researchers with the methods and tools they need to undertake the rigorous study of social systems through controlled computational experiments?' [p. 4–5]

Here we discover yet another framework underwriting agent-based models in social science. However, unlike the work of Levins or Cederman, Axelrod and

Tesfatsion prefer not to discuss specific criteria which may place a given agent-based model within these categories, preferring instead to provide only examples of each research goal.

Interestingly, despite their initial mention of Doran's described 'artificial societies' methodology, and the general goal of generating artificial 'world histories,' the remainder of Axelrod and Tesfatsion's introduction presents a sharp contrast to Doran's ideas. They stress the potential for agent-based models to produce substantive empirical insights related to real-world societies, rather than produce more general insights about the initial formation of societies and social interaction as a whole.

However, given the caveats presented thus far with agent-based social simulation, Axelrod and Tesfatsion's research goals seem at odds with the conclusions we have reached. If we agree, as discussed in the last chapter, that social simulation may lack some critical elements required to provide social explanation, then hoping to use such simulations to design real-world social institutions may produce disastrous results. Similarly, developing empirical understanding and 'causal explanations' would be equally problematic, particularly as any non-reductive aspects of the society under investigation would be difficult to identify.

Even if we disagree with the alleged difficulties in social explanation for social simulation, Doran's views have illuminated a further difficulty for Axelrod and Tesfatsion's suggested approaches. Without a clear definition of what constitutes an 'agent' in an empirical or normative simulation, the level of correspondence between these agents and human social actors is too ill-defined. What cognitive capacities would agents require to provide that degree of understanding? Surely Schelling-esque simplicity is not going to be the best course for all possible investigations using simulated societies?

6.7 Summary and Conclusions

Thus far we have seen a great variety of proposed theoretical frameworks for the use of agent-based models. Artificial life provided a useful starting point, given its relationship to the long-standing use of mathematical models in various biological disciplines. Comparing the methodological difficulties of artificial life with mathematical models in biology allows us to develop a greater understanding of the problems most particular to agent-based models.

However, applying the resulting methodological discussions to agent-based models in social science is not so straightforward. As discussed in Chap. 5, social scientists must cope with some unique difficulties: social structures are not clearly hierarchical in nature; empirical studies of social phenomena are frequently problematic due to the interacting complexities of individual and collective behaviour; and social explanation may suffer from difficulties similar to those faced by researchers in mental phenomena, as some social phenomena may be similarly irreducible to straightforward low-level behaviour.

Obviously all of these points impact the social simulator and make the job of model-building significantly more difficult. However, as seen during this analysis, these difficulties shift in emphasis depending on the purpose of the model (as we might expect from previous discussion regarding the Levins and Cederman frameworks). Taking an external, Luhmannian perspective, and incorporating Doran's related views, social simulation may be able to provide a unique window into certain aspects of social theory that are otherwise inaccessible through standard empirical means.

This methodology brings us back to the artificial world problem of Chap. 4. If we accept that 'growing' artificial societies from a level even below that of agents is a promising means for investigating potential fundamental social theories, we must likewise accept that these artificial societies may be quite difficult to relate to real-world societies. Given the gaps in our understanding in social science, ranging from unanswered questions in individual behaviour through to questions of high-level social organisation, can such artificial societies bring us any closer to the desired fundamental social theory?

The next chapter will bring us closer to answering this critical question. By examining one of the most prominent exemplars of social simulation, Schelling's residential segregation model, we will be able to illuminate this concern in greater detail. Schelling's model is the very definition of abstraction – and yet it is credited with producing some important insights into the social problem it is based upon. By determining how Schelling's model surpassed its abstract limitations, we can develop a framework which may underwrite future models in social simulation of a similar character.

This analysis of Schelling marks a conclusion of sorts to the arguments and frameworks discussed thus far in Parts I and II. Having discussed modelling frameworks in Alife and biology, and having brought those together with new elements from simulation in the social sciences, Schelling's model will provide a means to demonstrate the importance of these theoretical, methodological, and pragmatic frameworks for the modeller who wishes to push social science forward through simulation. As we shall see, Schelling's simple model was not nearly so simple in its widespread influence and overall importance. The elements which contributed to this success can be seen via its relationships to the frameworks discussed thus far.

References

Adami, C., & Brown, C. T. (1994). Evolutionary learning in the 2D artificial life system Avida. Cambridge: MIT Press.

Axelrod, R., & Tesfatsion, L. (2005). A beginner's guide to agent-based modelling in the social sciences. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics, Vol. 2: Agent-based computational economics*, p. Appendix A.

Beer, R. (2003a). Arches and stones in cognitive architecture. *Adaptive Behavior, 11*(4), 299–305.

Beer, R. (2003b). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior, 11*(4), 209–243.
Cederman, L. E. (2001). Agent-based modeling in political science. *The Political Methodologist, 10*, 16–22.
Chalmers, A. F. (1999). *What is this thing called science?* Buckingham: Open University Press.
Doran, J. E. (2000). Trajectories to complexity in artificial societies. In A. Kohler & G. Gumerman (Eds.), *Dynamics in human and primate societies*. New York: Oxford University Press.
Epstein, J. (1999). Agent-based computational models and generative social science. *Complexity, 4*(5), 41–60.
Gilbert, G. N. (1995) Emergence in social simulation. In G. N. Gilbert & R. Conte (Eds.), *Artificial societies: The computer simulation of social life*, pp. 144–156. London: UCL Press.
Gilbert, N., & Terna, P. (2000). How to build and use agent-based models in social science. *Mind and Society, 1*(1), 57–72.
Keeley, B. L. (1997). Evaluating artificial life and artificial organisms. In C. G. Langton & K. Shimohara (Eds.), *Artificial life V* (pp. 264–271). Cambridge: MIT Press.
Lustick, I. (2002). PS-I: A user-friendly agent-based modeling platform for testing theories of political identity and political stability. *Journal of Artificial Societies and Social Simulation, 5*(3), 7.
Lustick, I. (2006). Defining violence: A plausibility probe using agent-based modeling. *Paper Prepared for LiCEP*, Princeton University, 12–14 May 2006.
Ray, T. S. (1994). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life, 1*, 179–209.
Ray, T. S. (1996). An approach to the synthesis of life. In M. A. Boden (Ed.), *The philosophy of artificial life* (pp. 111–145). Oxford: Oxford University Press.
Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing*. Cambridge: MIT Press.
Silverman, E., & Bullock, S. (2004). Empiricism and artificial life. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, & R. A. Watson (Eds.), *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*. MIT Press/Bradford Books.
Woolgar, S. (1988). *Science: The very idea*. Chichester: Elias Horwood.