Chapter 6
Agent-Based Models – Because They're Worth It?

David O'Sullivan, James Millington, George Perry, and John Wainwright

Abstract We address the question of when the relative complicatedness of spatial agent-based models (ABMs) compared to alternative modelling approaches can be justified. The spectrum of ABM types from simple, abstract models to complicated models aspiring to realism makes a single answer impossible. Therefore we focus on identifying circumstances where the advantages of ABMs outweigh the additional effort involved. We first recall the reasons for building *any* model: to simplify the phenomena at hand to improve understanding. Thus, the representational detail of ABMs may not always be desirable. We suggest that critical aspects of the phenomena of interest that help us to assess the likely usefulness of ABMs are the nature of the decisions which actors make, and how their decisions relate to the spatio-temporal grain and extent of the system. More specifically, the heterogeneity of the decision-making context of actors, the importance of interaction effects, and the overall size and organization of the system must be considered. We conclude by suggesting that there are good grounds based on our discussion for ABMs to become a widely used approach in understanding many spatial systems.

D. O'Sullivan (✉) • G. Perry
School of Environment, University of Auckland, Auckland, New Zealand
e-mail: d.osullivan@auckland.ac.nz; george.perry@auckland.ac.nz

J. Millington
Department of Geography, King's College London, Strand, London, UK
e-mail: james.millington@kcl.ac.uk

J. Wainwright
Department of Geography, University of Sheffield, Sheffield, UK
e-mail: j.wainwright@sheffield.ac.uk

A.J. Heppenstall et al. (eds.), *Agent-Based Models of Geographical Systems*,
DOI 10.1007/978-90-481-8927-4_6,

6.1 Introduction

In this chapter we critically examine the usefulness of agent-based models (ABMs) in geography. Such an examination is important because although ABMs offer some advantages when considered purely as faithful representations of their subject matter, agent-based approaches place much greater demands on computational resources, and on the model-builder in their requirements for explicit and well-grounded theories of the drivers of social, economic and cultural activity. Rather than assume that these features ensure that ABMs are self-evidently a good thing – an obviously superior representation in all cases – we take the contrary view, and attempt to identify the circumstances in which the additional effort that taking an agent-based approach requires can be justified. This justification is important as such models are also typically demanding of detailed data both for input parameters and evaluation and so raise other questions about their position within a broader research agenda.

One inspiration for our approach is found in a brief but challenging commentary by Helen Couclelis (2002). Noting that ABMs add to "the well-known problems of modeling a highly complex, dynamic spatial environment" (pp. 4–5), the additional difficulties of "modeling highly complex, dynamic decision-making units interacting with that environment and among themselves in highly complex, dynamic ways". She continues:

> The question is whether the benefits of that approach to spatial modeling exceed the considerable costs of the added dimensions of complexity introduced into the modeling effort. (Couclelis 2002, pp. 4–5)

Couclelis offers her own answer, when she goes on to say: "The answer is far from clear and, in my mind, it is in the negative" (p. 5). However, Couclelis does leave the door open to further discussion. Others such as Gross and Strand (2000) have argued that capturing micro-scale complexity requires models with the complex micro-structures that the agent-based approach incorporates; in short, a complex world requires structurally complex models. These contrasting perspectives make it clear that an open question remains: under what circumstances is the extra effort of these data- and theory-intensive models rewarded, and why? The aim of this chapter is therefore to establish under which circumstances ABMs really are worth it!

6.2 Horses for Courses: Different Agent Models for Different Purposes

There are many possible ways of classifying ABMs (see Crooks and Heppenstall 2012 for a brief overview). In geographical applications, at the most abstract level, an ABM consists of agents interacting with and in an environment. Various typologies can be constructed on the basis of the nature of the agents and of the environmental representation. Couclelis (2002, p. 4) offers one such classification based on whether the agents and the environment are 'designed' or 'analyzed'. This terminology is somewhat confusing (it derives from an engineering perspective), but may be clearer if we replace 'designed' with *theoretically derived* and 'analyzed' with *empirically derived*. Couclelis goes on to consider the purpose of these different possible combinations of agent and environment type.

An alternative approach to classifying ABMs is to consider three broad styles of model (see O'Sullivan 2008). Arguably, the bulk of academically orientated work to date using ABMs continues to be in the realm of simple abstract models where the focus is on exploring the collective implications of individual-level decision making. Schelling's book title *Micromotives and Macrobehaviour* (Schelling 1978) captures the intention of this approach well (and is discussed by Birkin and Wu 2012). The 'Schelling model' of residential segregation is the most familiar example of this style (Schelling 1969), and has spawned a cottage industry of variants and explorations of how various minor changes to the assumptions underlying the model affect the outcomes (see Fossett 2006, for a detailed exploration of some aspects of the model). In the same vein are Epstein and Axtell's (1996) *Sugarscape* models, Axelrod's work on iterated game theoretic models (Axelrod 1997) and many ABMs of economic behaviour (see Tesfatsion and Judd 2006). Examples of this abstract approach in an urban context include Batty's work on how simple movement and resource exploitation actions on heterogeneous landscapes produce characteristic settlement size distributions (Batty 2005, Chap. 8), and a preliminary model of sprawl presented by Brown and Robinson (2006). The abstract approach is also common in other fields such as biology (see, for example, Ehrlich and Levin (2005)). It is this style of work which is largely responsible for excitement in some quarters around the potential of 'complexity science' to answer general questions about the nature of systems in a wide range of specialist fields (e.g. Bar-Yam 1997).

A second type of ABM is more detailed and locates virtual model agents in a representation of the real world setting of interest. Typically, such models operate at a regional or landscape scale, although this is dependent on the issue(s) that a particular model is addressing. A common application for this flavour of ABM is land-use and cover change (LUCC), often in the context of climate-change scenarios. A recent special issue of *Landscape Ecology* (Milne et al. 2009) gives a sense of the diversity of models in this context, and also of the importance of integrating ABMs with those other approaches. Examples of the type we have in mind are the work of Millington et al. (2008) and Matthews (2006). Here, the goal of developing a model is to understand how expected or possible changes in the behaviour of individual entities arising from the changing policy environment affect landscape-level variables that feedback to both agent behaviour and resulting system-level outcomes (such as, for example, climate change). A different context for models of this kind is the attempt to understand how an urban streetscape or a complicated building design affects the behaviour and paths followed by pedestrians interacting in that environment (Haklay et al. 2001; Helbing et al. 2001; Kerridge et al. 2001). The common thread linking these settings is that the interactions among agents may have more or less dramatic effects on the overall outcomes of the model. In both cases, agent actions change the decision-making environment of other agents, albeit at different spatio-temporal scales, and in different ways. In a LUCC model, more or less permanent changes in the environment are made by agent actors, and these collectively affect future decision-making for all agents at the scale of the whole model. In a pedestrian model, the urban or built environment is fixed, and the agents themselves are a salient and rapidly changing feature of the environment, which affects agents, often only locally.

Thirdly, some of the most ambitious models aim at detailed (i.e. "realistic", although see Dietrich et al. 2003, pp. 7–8, for a more extended consideration of realism in models) representations of both the geographical setting and the processes unfolding in that setting. Such models tend to be driven by the concerns of policy- and decision-makers and revolve around urban, economic, and demographic management applications. The most obvious example of this style of model is the TRANSIMS ABM of urban traffic where every individual vehicle in a large urban system is represented second-by-second (Cetin et al. 2002). Closely related to TRANSIMS is EpiSims, which takes the same approach to epidemic processes in detailed representations of social networks (Toroczkai and Guclu 2007). When models become this large, it becomes difficult to get to grips with their overall structure, or even to consider them as truly single models. The 'model' becomes a framework in which subsystem models are integrated. An example of this approach which has evolved over many years is the SIMPOP family of urban growth models (Sanders et al. 1997; Bretagnolle et al. 2009). The modular and extensible structure of such models is an attempt to cope with the difficulties inherent in extending the scope of individual-based models as they grow to encompass large scale continental or global systems, a problem which is also encountered in using and interpreting general circulation models of global climate.

This last category makes it clear that any typology of ABMs is necessarily highly schematic. The three types of ABM we have discussed are more like points along a continuum of increasing size and complexity than discrete categories. The value of developing such a typology at all is to realize that ABMs are built for a wide variety of reasons across a wide range of disciplines. ABMs, like all models, may be used to explore theories and their possible implications, to understand how particular theories may play out in particular contexts, and to assist in risk-assessment, or policy- and decision-making. This complicates answering the question of whether or not ABMs are useful in any particular application, although it suggests that the answer is "it depends!" (on context, on purpose, on application, and so on). Even so, it is possible to be more specific about the situations where agent approaches are likely to justify the additional effort and cost that their development, analysis and use entail.

6.3 Are Modellers Agent-Based Because They Should Be or Because They Can Be?

While there has been a lot of excitement in recent years about the potential of agent-based methods, it is important to remember that none of the cases cited above is one where agent models are the only possible approach. In most cases, ABMs are a relatively late arrival in a field where there is considerable previous experience with styles of model that adopt a more aggregated approach, and these aggregated models continue to be widely used. Thus, for example, land-use transport models, which are calibrated and run based on transport analysis zones, are much more widely deployed by city governments worldwide than ABMs at the individual vehicle level simulating morning and evening rush hours; see Wegener (2004). What, if anything do ABMs add, and by extension, when should we prefer ABMs over more traditional methods?

At times, it appears that the main motivation for adopting an agent-based approach is simply because it can (now) be done. While the tools available for ABMs (Railsback et al. 2006) are not yet as accessible or as well developed as those for more established approaches such as systems dynamics (Muetzelfeldt and Massheder 2003; Deaton and Winebrake 2000; Eberlein and Peterson 1994), ABMs have surprisingly quickly become a viable approach for the spatial model builder. The increasing ease with which ABMs can be developed, coupled with their intuitively satisfying representational approach, in which each software agent represents an 'actor' whether an individual person (or animal or plant) or an institution (often the barely more aggregated household) has led to widespread enthusiasm for the approach. The appeal is undeniable: it appears obvious that individual-level decision making is the fundamental driver of social systems, or more broadly that the individual-level behaviours of plants and animals drive environmental change. Setting to one side the thorny question of whether or not social phenomena are distinctive in kind from the merely aggregate actions of individuals (see O'Sullivan and Haklay 2000), and hence also the question of whether it is the case that social and environmental systems really are driven entirely by individual-level decision-making, if we *can* represent systems at the 'atomic' level on which they operate, then surely we *should*?

In our view this stance ignores the motives for developing models in the first place. Put simply, the need for a model arises when understanding the world itself is too hard! The danger of wholesale adoption of ABMs is that we simply replace one difficult to understand phenomenon – the world itself – with an equally hard to understand model. This is the difficulty that Couclelis identifies in her commentary. A model that advances our understanding is one that represents what are considered in a particular context the key features of a system and thus enables us to improve our understanding of how that system works. Any gain in understanding of the system resulting from the modelling process derives from our ability to analyze the model and experiment with it. If the model is too complicated to analyze, all we have done is to replace one poorly understood object of study with another, which we know to be incomplete! There are good reasons to believe that using disposable 'fast-and-frugal' models will result in more rapid learning than highly detailed ones (Carpenter 2003), and in most, if not all cases, ABMs are not a 'fast-and-frugal' option.

Considering such issues is at the heart of all model building. However, ABMs are one aspect of a recent trend towards more complicated and detailed models. This trend flies in the face of longstanding conventions in modelling and simulation, which hold that simpler, more parsimonious models are preferable to complicated ones, all other things being equal. The search for parsimony in models is often presented as a logical consequence of Ockham's razor (see Perry 2009). That is not a position we wish to defend. First, it is clear that Ockham's admonishment to avoid the 'unnecessary multiplication of entities' was never intended to guide the development of simulation models! Second, there is no *a priori* reason for assuming that the world is a simple place, when it is patently not!

Careless application of the principle of Ockham's razor might lead us to conclude that a less complicated model is more convincing, just because it is less complicated, although this is not a logically defensible point of view. Ockham's razor is an argument about the capacity of different descriptions of reality to explain observed phenomena, not grounds for always preferring simpler explanations to more complicated ones. Even so, there are good pragmatic reasons for preferring parsimonious models. Such models are much easier to learn from than models with many parameters and sub-models. They are easier and more cost-effective to parameterize, and they are also much less vulnerable to the propagation of errors due to the uncertainties in estimating multiple interrelated parameters (again, see Carpenter 2003).

Based on this observation, the important question is to determine what features agents bring to a model *which make a difference that matters*. This concern is similar to the argument made by Andrew Sayer in his consideration of "the difference that space makes" in explaining social systems (Sayer 1985). Although he is discussing the role of space in social theory, Sayer's arguments seem to us to apply with equal force to the evaluation of models. The basis of the argument is the distinction to be made between *necessary* and *contingent* features of a theory. Some aspects of any phenomena we wish to explain are absolutely central – that is, necessary – to the nature of that phenomena, while others are peculiar to occurrences of those phenomena in particular contexts – that is, contingent on those particular occurrences. A less philosophical way to express the same idea is simply to ask, which features of the phenomena we are interested in are essential? Asking this question is really what building a model is all about. Answering this question in the context of ABMs should focus our thinking on the issue of what the agents in a model are, what they do, and following from this, when they are necessary to any representation of the phenomena of interest. In the remainder of this chapter, we sketch out the circumstances in which agents are more likely to be necessary to an adequate model. In our conclusions, we briefly revisit the idea of contingency and its relevance to this issue.

6.4 What Are Agents? And What Do They Do?

These considerations bring us to the basic question of what adopting an agent-based representation in a model achieves in terms of a simulation. There is general agreement (amidst much debate about finer points!) on the basic characteristics of agents in spatial models. More detailed consideration of the meaning of the defining characteristics of such agents can be found in Crooks and Heppenstall (2012). We consider the most fundamental characteristics of agents in spatial models to be goal-direction and autonomy (Jennings et al. 1998). However, more specific definitions of the concept may add any of flexibility, 'intelligence', communication, learning, adaptation or a host of other features to these two. In practice, whatever way we describe their characteristics, agent actions in models revolve around exercising *choice* among available options in order to achieve defined *goals*.

The outcome of an agent making a particular choice is some difference in either the location of the agent (i.e. the agent moves) or in the environment. In the latter case, the agent alters the attributes of its current location in some way. Depending on the model context, this may involve the agent exploiting resources at its current location (and hence depleting the supply of those resources at that location); altering the state of the location (e.g. changing the land use); acquiring the land at its present location; or, perhaps simply updating its current 'map' of the environment. In each case, there may be an accompanying change in the state of the agent itself, such as when resource exploitation increases the agent's wealth or energy resources.

This account of spatial ABMs (and it is important to note that there are many examples in the literature of aspatial ABMs) has several implications:

- Agents may be mobile, but this is not a necessary feature (models of trees in forests are among the most common types of ABM). However, it is important that each agent has a different relationship with the spatial environment, most simply in terms of a location in the environment. If all agents have the same spatial relationship with the environment (if, for example, every agent has an equal capability to alter every location in the model regardless of the agent's specific location or every agent sees and responds to an aggregate 'average' of the environment), then it makes little sense to formulate the model as an agent model;
- Agents may change their spatial relationship with the environment over time, which may be by moving, or it may be by alteration, acquisition or disposal of locations; and
- Agents are able to evaluate spatial configurations. This ability may be as simple as determining that the availability of some resource at the current location is sufficient for some purpose, or is greater than at neighbouring locations. Alternatively, it may involve a complicated evaluation of the spatial distribution of resources (including other agents) with respect to the current location, relative to a number of alternative locations.

This framework for thinking about agents in a spatial ABM may be illuminated by considering some examples (see also Fig. 6.1):

1. *Pedestrian or other mobile agents* in a model of an urban streetscape or complex building. The primary choice made by such agents is to determine, with respect to their intended destinations, which among the possible next locations they should move to. In most models of this kind, the location of other agents is an important element in the choice, but the decision will also be affected by the agents' local physical environments (e.g. building geometries).

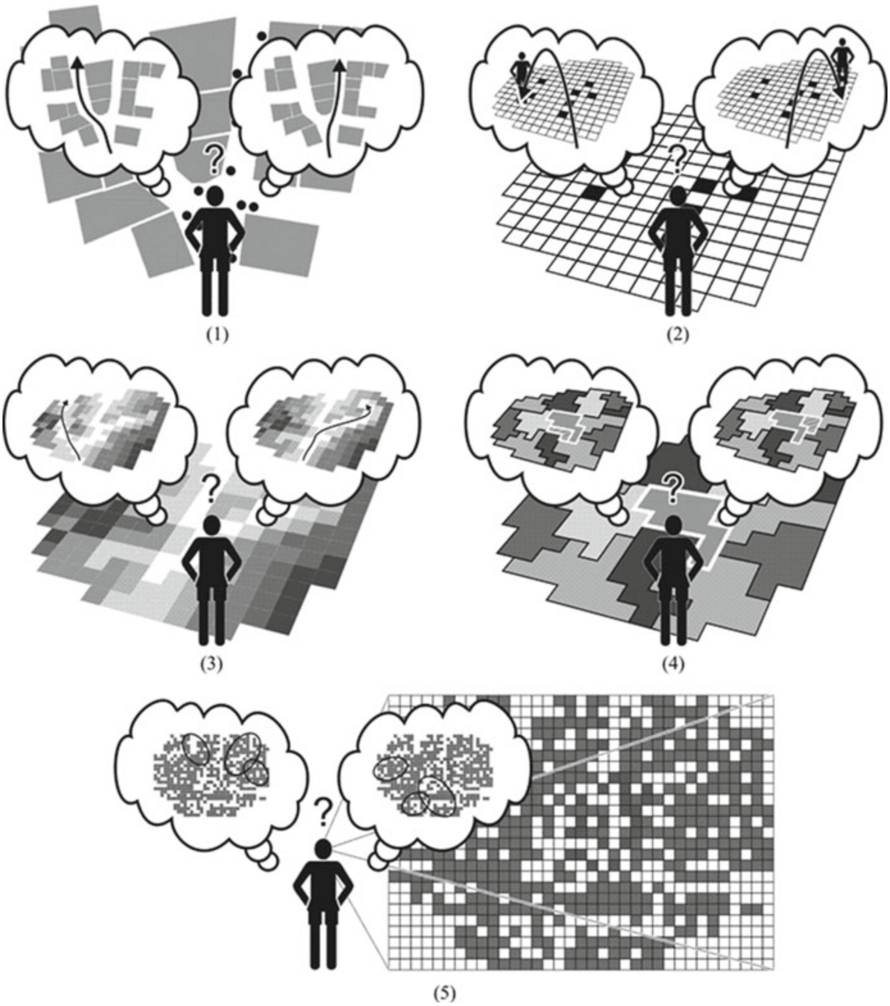

Fig. 6.1 Schematic illustration of the choices facing agents in five different types of model. See text for details

2. *Residential agents* in a 'Schelling-style' model are also primarily concerned with movement, although it is movement of a rather different kind. They evaluate their current and potential new locations, and if one of the new locations considered is preferable in some way to their current location, then they may move there. Again, the locations of other agents in this model influence the choices made by each agent, but the nature of the environment itself does not.
3. *Hunter-gatherer agents* in a model of resource exploitation in which establishment of permanent settlements is an outcome will probably combine aspects of the two previous types of agent, in that they evaluate competing locations, and

will choose to stay or go depending on the resources at those locations. Rather differently to the previous two cases, however, the actions of these agents will alter the environment directly, not just in terms of the location of the agents themselves.

4. *Farmer/land-use change agents*, like the previous type, alter the environment itself, but unlike them are unlikely to move in the process. They may alter their relationship with the spatial environment by acquiring or disposing of land as one aspect of the management of their resources.
5. *Property developer agents* in an urban growth or development model are unlikely to be explicitly spatially located in the way that agents in the previous examples are. Like farmer agents, they will have some attachment to a 'territory', which they are able to grow, change, or reduce by acquisitions, development actions, or sales into a property market. Such agents are likely to have a relatively sophisticated ability to evaluate spatial configurations of currently owned locations relative to the various land uses and land values in the model.

Aspects not explicitly considered in these examples, but highly relevant in practice, are the spatial and temporal *grain* of the model representation, and the relationship between the two. By grain we mean the extent of the smallest units of space and time which are explicitly represented in a model. A fine-grained model might represent second-by-second developments at spatial resolutions of a metre or less; traditionally, such models have been seen as unable to consider extended spatio-temporal domains. A coarse-grained model might operate on large units of space (say several square kilometres) over time periods of a year or more. Grains much coarser than this seem unlikely in practice because ABMs are about the choice-making behaviour of individual living actors. While contemporary societies occasionally aspire to decision-making that takes into account time horizons longer than a year or so (and simulation models are seen as central to this decision-making; see Clark et al. 2001), it is rare for choices to be 'locked in' over much longer time frames than this. Similarly, it is difficult to imagine an ABM model that would be recognizable as such where spatial agents act on 'local' spatial knowledge more wide-ranging than a few kilometres.

Note that we adopt the concept of grain here in preference to spatio-temporal scale because the latter often also implies the overall extent or scope of a model. While the grain of the representation in a model and its overall scope are not independent, it is increasingly common to see unexpected combinations particularly of fine grains with wide extents (for example, Epstein 2009, refers to an epidemic model that explicitly represents the whole population of the Earth as individual agents).

Although as geographers we might wish to grant representation of the spatial aspects priority over temporal aspects, temporal considerations are of at least equal importance, not least because the two are interlinked (both conceptually and computationally). Decisions are usually made by agents over some timeframe of interest, which may in turn imply a relevant spatial grain.

In a pedestrian model this timeframe might be second-by-second, as pedestrians adjust their course to avoid obstacles (including other agents). More generally, mobile agents (whether human or some other animal) will be making decisions at time scales dictated by their mobility on the one hand and their perception of the nature of the spatial distribution of resources on the other. The decision-making timeframe combined with the speed of movement of the agents then effectively dictates a sensible spatial grain for a model of this type – with plausible models capturing spatial detail down to or below the metre range for human agents. In resource exploitation models the timeframe of interest is dictated by context. In a model of hunter-gatherer behaviour with only limited storage of resources, daily or weekly activity patterns and decisions will predominate, and this, combined with rates of movement, will govern how we represent spatial aspects of both the agents *and* their environment. In cases where the mobility of agents is less dominant, as in the farmer or property developer examples above, the linkage between the temporal and spatial grain is less direct, but nevertheless remains important. The key issue in these cases is how rapidly agents change the environment, and how quickly those changes affect the later decisions of other agents. A monthly, seasonal or annual timeframe is likely to be the most appropriate in these cases, since the outcomes of planting or development decisions that take appreciable times to unfold will affect further decision making. In these cases the spatial grain is a product of the amount of change which can be effected by individual agents over the chosen time frames. This in turn will be dependent on organizational features of the agents themselves in particular if they are institutional actors. For example, where property developers are small businesses, we may be interested in development at the level of individual land parcels. Where we are interested in larger corporate actors, the spatial extent of agent actions may be much larger.

In the one highly abstract case we consider above, that of 'Schelling-style' residential relocation models, these considerations are a lot less clear-cut. In such cases, questioning the spatial and temporal grain can contribute to conclusions that may be considered very unflattering to the model under examination; see, for example, Goering (2006). The essentially theoretical, abstract nature of the model comes to the fore and the spatio-temporal grain of the representation is of less relevance than its structure and the overall system tendencies it points to.

6.5 So When Do Agents Make a Difference?

The emphasis we have placed on decision making by agents and the related choice of the spatial and temporal grain in a model helps to address our original question about when it is appropriate to adopt an agent-based representation in a model. If the decisions at the heart of a model are made in local contexts, which depend in turn on the spatio-temporal grain of the model in such a way that every agent decision reduces to the same decision, then an aggregated statistical or mathematical representation may be sufficient. The classic examples from game theory, such as Prisoner's Dilemma and the Tragedy of the Commons fit this template well, and continue to shed light on the overall structure of many social systems and coordination problems.

Where agents' preferences and (spatial) situations differ widely, and where agents' decisions substantially alter the decision-making contexts for other agents, there is likely to be a good case for exploring the usefulness of an agent-based approach. This argument focuses attention on three model features: *heterogeneity* of the decision-making context of agents, the importance of *interaction effects*, and the overall *size* and *organization* of the system.

If agents are the same throughout the system, then, other things being equal, an aggregate approach is likely to capture the same significant features of the system as an agent-based approach. However, it is important to extend our concern with heterogeneity to encompass not just agents but to agents in their (spatial) decision-making contexts. A population of identical agents in diverse contexts can produce somewhat unexpected outcomes as a result of different choices being made in those different contexts, which then alter the options available to all agents at subsequent times. 'Schelling-style' models exemplify this. The opposite case, where every agent makes its choices in the same context but heterogeneity in the agents may produce dramatically different results depending on the degree of heterogeneity, is less familiar. An example is provided by Rand et al. (2002), whose abstract model of urban growth shows that the existence of even small numbers of households with a preference for aesthetic over urban amenity can dramatically accelerate exurban sprawl.

In both of these cases, agent actions result in changes to the decision making context for other agents, an indirect and weak form of agent-to-agent interaction. Some form of agent interaction is necessary at a minimum if an agent-based approach is to be justified. If each agent's decisions make no difference to the subsequent decision-making contexts of other agents, then the generalized pay-off matrices of classical game theory are again likely to provide a sufficient representation of social systems. The stronger any interaction effects are, then the more important it will be to consider agent-based or other disaggregated approaches. In a pedestrian model, interaction is direct. Each pedestrian agent is a significant element in the local environment of many other agents, and decisions made by one agent immediately alter the local decision-making environment of nearby agents. Where the contexts for decision-making are more general, based on aggregate system measures, so that each individual's decisions make only minor differences to the choices of others, then the case for an agent-based approach is less clear.

By the system size, we mean the total number of agents in the system. This aspect relates to the previous point. In large systems, other things being equal, unless interaction effects are strong and direct, it may not be necessary to adopt ABM approaches. In such cases, mean-field approaches provide appropriate representations of system dynamics (Berec 2002). This consideration is closely related to one of the earliest characterizations of the idea of system complexity by Warren Weaver (1948), who distinguishes middle-sized systems of "organized complexity" from small systems of only a few elements on the one hand, and large systems of disorganized complexity explicable in statistical terms (gases are the obvious example) on the other. Systems of organized complexity are those where interaction among elements – more than that, iterative or hierarchical *organization* of the elements – renders statistical explanation inadequate. He wryly notes that the size range of such systems is very broad: "large compared to two, but small compared to the number of atoms in a pinch of salt" (Weaver 1948, p. 539). Taking only system size into account, this aspect may appear redundant in determining the viability of agent-based approaches since all social systems (that we know of!) fall into this broad 'middle' range.

To resolve this issue, we must delve more deeply into the idea of system organization. Where systems are sufficiently 'organized', it may be that intermediate levels of organization are durable enough to form the atomic units on which we should focus in a model, rather than individuals. This fact is already implicit in cases where households rather than individuals are the agents in a model. Similarly, in economic models, firms are often recognized as the appropriate units for representation. In models of large collections of individual actors, perhaps the most important question for the would-be agent-based modeller to ask is not "is an ABM appropriate?" (where the presumption often is that agents should represent individual actors). A more important question may be, "what should the agents in an ABM of this system represent?" If the interactions among individual actors in the real world are substantially channelled via institutions or other social or spatial structures, perhaps it is those institutions or social or spatial structures that should be represented as agents in an ABM rather than the individuals of which they are formed. One way to think about this is to see that in choosing to represent not individual actors as agents but instead some other intermediate level aggregate entity, we are effectively reducing the system size to a point where actions of individual agents make a difference, thus justifying the approach.

All three of our system criteria favouring the adoption of ABM – heterogeneity, interaction, and the combined effects of system size and organization ('middle-numbered-ness') – call for considerable prior knowledge and insight about system characteristics on the part of those developing models. Thus, it would be wrong to draw any universal conclusions from our account to a statement about the usefulness of agent-based approaches in general. Instead, we strongly recommend careful consideration of the system features we have discussed before simply assuming that an agent-based representation is inherently superior. Where consideration of these aspects suggests that an agent-based representation is indeed necessary, then it is worth noting that the resulting model is often one where a full explanation of the model behaviour calls for a historical account of the events in the model. If agents are necessary in the model because they are differentiated from one another, because they interact meaningfully with one another, and because they are able to make a difference to system level outcomes, then in describing and understanding the model, it is likely that Sayer's (1985) contingent effects will be significant. Thus particular agent-agent interactions will matter, and a detailed account of the model 'history' may be necessary for a complete understanding of any particular model run. The difference from the real world target system we seek to understand, is that a model allows repeated runs and enables a probabilistic or general account of the system behaviours and tendencies to be developed.

Our discussion relies on *a priori* understanding or analysis of the system structure, or *post hoc* assessment of whether the resulting model demonstrates the historical-contingent features that would suggest it was the right choice of approach. Neither is a particularly satisfactory or systematic way to decide whether or not to embark on the demanding and potentially costly development of an ABM approach in a particular case. Given the complex nature of the systems and problems involved, it is difficult to see how at least piloting ABM and alternative approaches can be completely avoided (another reason for preferring simple models to complicated ones?), but recent approaches do suggest ways in which the usefulness of ABMs can be assessed, such as pattern-oriented modeling (Grimm et al. 2005; Grimm and Railsback 2012) and the comparison of mean-field and individual-based models (Iwasa 2000).

While we cannot make sweeping general claims from our discussion, it seems clear that human settlement systems are often strong candidates for agent-based representations. This claim is based on the criteria for the usefulness of ABMs that we have identified: heterogeneity, interaction, and system size and organization. Similar arguments can be made about human-environment systems more generally, even in prehistoric settings where the degree of organization of the social systems may be rather more limited. While other approaches remain useful, arguments against building ABMs based on the extra effort involved can be countered because the potential for insight and understanding from building and using such models makes those efforts worth it.

References

Axelrod, R. M. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton: Princeton University Press.

Bar-Yam, Y. (1997). *Dynamics of complex systems*. Reading: Perseus Books.

Batty, M. (2005). *Cities and complexity: Understanding cities with cellular automata, agent-based models and fractals*. Cambridge: MIT Press.

Berec, L. (2002). Techniques of spatially explicit individual-based models: Construction, simulation, and mean-field analysis. *Ecological Modelling, 150*, 55–81.

Birkin, M., & Wu, B. (2012). A review of microsimulation and hybrid agent-based approaches. In A. J. Heppenstall, A. T. Crooks, L. M. See & M. Batty (Eds.), *Agent-based models of geographical systems* (pp. 51–68). Dordrecht: Springer.

Bretagnolle, A., Pumain, D., & Vacchiani-Marcuzzo, C. (2009). The organization of urban systems. In D. Lane, D. Pumain, S. E. van der Leeuw, & G. West (Eds.), *Complexity perspectives in innovation and social change* (pp. 197–220). New York: Springer.

Brown, D. G., & Robinson, D. T. (2006). Effects of heterogeneity in residential preferences on an agent-based model of urban sprawl. *Ecology and Society, 11*, 46–65.

Carpenter, S. R. (2003). The need for fast-and-frugal models. In C. D. Canham, J. J. Cole, & W. K. Lauenroth (Eds.), *Models in ecosystem science* (pp. 455–460). Princeton: Princeton University Press.

Cetin, N., Nagel, K., Raney, B., & Voellmy, A. (2002). Large-scale multi-agent transportation simulations. *Computer Physics Communications, 147*, 559–564.

Clark, J. S., Carpenter, S. R., Barber, M., Collins, S., Dobson, A., Foley, J. A., et al. (2001). Ecological forecasts: An emerging imperative. *Science, 293*, 657–660.

Couclelis, H. (2002). Why I no longer work with agents: A challenge for ABMs of human-environment interactions. In D. C. Parker, T. Berger, & S. M. Manson (Eds.), *Agent-based* *models of land-use and land-cover change: Report and review of an international workshop,* Irvine, Oct 4–7, 2001 (pp. 3–5). Louvain-la-neuve, Belgium: LUCC International Project Office.

Crooks, A. T. & Heppenstall, A. J. (2012). Introduction to agent-based modelling. In A. J. Heppenstall, A. T. Crooks, L. M. See & M. Batty (Eds.), *Agent-based models of geographical systems* (pp. 85–105). Dordrecht: Springer.

Deaton, M. L., & Winebrake, J. J. (2000). *Dynamic modeling of environmental systems*. New York: Springer.

Dietrich, W. E., Belugi, D. G., Sklar, L. S., Stock, J. D., Heimsath, A. M., & Roering, J. J. (2003). Geomorphic transport laws for predicting landscape form and dynamics. In P. R. Wilcock & R. M. Iverson (Eds.), *Prediction in geomorphology* (Geophysical monograph, Vol. 135, pp. 1–30). Washington DC: American Geophysical Union.

Eberlein, R. L., & Peterson, D. W. (1994). Understanding models with VenSim. In J. D. W. Morecroft & J. D. Sterman (Eds.), *Modeling for learning organizations* (pp. 339–348). New York, NY: Productivity Press.

Ehrlich, P. R., & Levin, S. A. (2005). The evolution of norms. *PLoS Biology, 3*, e194. doi:10.1371/journal.pbio.0030194.

Epstein, J. M. (2009). Modelling to contain pandemics. *Nature, 460*, 687.

Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: social science from the bottom up*. Washington, DC: The Brookings Institute.

Fossett, M. A. (2006). Ethnic preferences, social distance dynamics, and residential segregation: Theoretical explorations using simulation analysis. *Journal of Mathematical Sociology, 30*, 185–274.

Goering, J. (2006). Shelling [sic] redux: How sociology fails to make progress in building and empirically testing complex causal models regarding race and residence. *Journal of Mathematical Sociology, 30*, 299–317.

Grimm, V. & Railsback, S. F. (2012). Designing, formulating and communicating agent-based models. In A. J. Heppenstall, A. T. Crooks, L. M. See & M. Batty (Eds.), *Agent-based models of geographical systems* (pp. 361–377). Dordrecht: Springer.

Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., et al. (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science, 310*, 987–991.

Gross, D., & Strand, R. (2000). Can agent-based models assist decisions on large-scale practical problems? A philosophical analysis. *Complexity, 5*, 26–33.

Haklay, M., Schelhorn, T., O'Sullivan, D., & Thurstain-Goodwin, M. (2001). "So go down town": Simulating pedestrian movement in town centres. *Environment and Planning B: Planning & Design, 28*, 343–359.

Helbing, D., Molnár, P., Farkas, I. J., & Bolay, K. (2001). Self-organizing pedestrian movement. *Environment and Planning B: Planning and Design, 28*, 361–383.

Iwasa, Y. (2000). Lattice models and pair approximation in ecology. In U. Dieckmann, R. Law, & J. A. J. Metz (Eds.), *The geometry of ecological interactions: Simplifying spatial complexity* (pp. 227–252). Cambridge: Cambridge University Press.

Jennings, N. R., Sycara, K., & Wooldridge, M. (1998). A roadmap of agent research and development. *Autonomous Agents and Multi-Agent Systems, 1*, 7–38.

Kerridge, J., Hine, J., & Wigan, M. (2001). Agent-based modelling of pedestrian movements: The questions that need to be asked and answered. *Environment and Planning B: Planning and Design, 28*, 327–341.

Matthews, R. (2006). The People and Landscape Model (PALM): Towards full integration of human decision-making and biophysical simulation models. *Ecological Modelling, 194*, 329–343.

Millington, J., Romero-Calcerrada, R., Wainwright, J., Perry, G. W. L. (2008). An agent-based model of Mediterranean agricultural land-use/cover change for examining wildfire risk. *Journal of Artificial Societies and Social Simulation*, *11*(4), 4. Available at: http://jasss.soc.surrey.ac.uk/11/4/4.html

Milne, E., Aspinall, R. J., & Veldkamp, T. A. (2009). Integrated modeling of natural and social systems in land change science. *Landscape Ecology, 24*, 1145–1147.
Muetzelfeldt, R. I., & Massheder, J. (2003). The Simile visual modelling environment. *European Journal of Agronomy, 18*, 345–358.
O'Sullivan, D. (2008). Geographical information science: Agent-based models. *Progress in Human Geography, 32*, 541–550.
O'Sullivan, D., & Haklay, M. (2000). Agent-based models and individualism: Is the world agent-based? *Environment and Planning A, 32*, 1409–1425.
Perry, G. L. W. (2009). Modelling and simulation. In N. Castree, D. Demeritt, D. Liverman, & B. Rhoads (Eds.), *A companion to environmental geography* (pp. 336–357). Hoboken: Wiley.
Railsback, S. F., Lytinen, S. L., & Jackson, S. K. (2006). Agent-based simulation platforms: Review and development recommendations. *Simulation, 82*, 609–623.
Rand, W., Zellner, M., Page, S. E., Riolo, R., Brown, D. G., & Fernandez, L. E. (2002). The complex interaction of agents and environments: An example in urban sprawl. In C. Macal & D. Sallach, (Eds.), *Proceedings of Agent 2002: Social Agents: Ecology, Exchange, and Evolution* (pp. 149–161). Chicago: University of Chicago and Argonne National Lab. Available at http://www.dis.anl.gov/pubs/2002%20Proceedings.pdf
Sanders, L., Pumain, D., Mathian, H., Guerin-Pace, F., & Bura, S. (1997). SIMPOP: A multi-agents system for the study of urbanism. *Environment and Planning B: Planning and Design, 24*, 287–305.
Sayer, A. (1985). The difference that space makes. In D. Gregory & J. Urry (Eds.), *Social relations and spatial structures* (pp. 49–66). London: Macmillan.
Schelling, T. C. (1969). Models of segregation. *American Economic Association Papers and Proceedings, 59*, 488–493.
Schelling, T. C. (1978). *Micromotives and macrobehavior*. New York: Norton.
Tesfatsion, L., & Judd, K. L. (2006). *Handbook of computational economics, volume II: Agent-based computational economics*. Amsterdam: Elsevier/North-Holland.
Toroczkai, Z., & Guclu, H. (2007). Proximity networks and epidemics. *Physica A: Statistical Mechanics and Its Applications, 378*, 68–75.
Weaver, W. (1948). Science and complexity. *American Scientist, 36*, 536–544.
Wegener, M. (2004). Overview of land-use transport models. In D. A. Hensher & K. Button (Eds.), *Transport geography and spatial system* (pp. 127–146). Kidlington: Pergamon/Elsevier Science.