Introduction

Strategic management, in line with other social science fields, is difficult: it is a subject that does not lend itself easily to experimentation, as the environment co-evolves at the same time that specific strategies are being evaluated. Agent-based modelling offers an exciting opportunity to experiment where interactions between other firms and the complex environment can be modelled and robust strategies developed accordingly.

Agent-based modelling provides an opportunity to experiment by way of simulation, where experiments can be undertaken, strategies evaluated, and interdependencies examined without making irreversible commitments of one's own firm. Agent-based modelling builds on the long history of work within complexity science, where entities – in our case firms – interact in a non-linear manner, much as they do in real life (despite many theorists' simplifying assumptions that assume they do not). Strategic management, when defined as the search to gain and maintain competitive advantage, is starting to be examined by researchers using agent-based modelling techniques borrowed from evolutionary biology, physics and other natural sciences. This offers a fascinating opportunity to extend existing research and to develop new models for competition that are not constrained by the traditional limitations of equilibrium or comparative statics.

We investigate in the following way. Section 1 is a review of agent-based models and of complexity science more generally. Section 2 offers a review of antecedents to the world of agent-based modelling from economic and systems models, organizational ecology, and biological models with their conceptualizations of fitness landscapes. Section 3 introduces the *NK* model, perhaps the most successful agent-based model within strategic management and organizational science, and shows alternative models that include competition between firms, the lack of competition being a limitation of the *NK* series of papers. Finally, Section 4 suggests directions for future research building on existing agent-based models that have yet to be fully developed within the field of strategic management.

1 Agent-Based Models

Agent-based models, a class of simulation models (for a general review of simulation, see Robinson 2004), can be traced back to game theorists' research such as von Neumann's (1966) 'work on computers' and later cellular automata such as Conway's *Game of Life* (Gardner 1970). These cellular automata were grid-based simulations whose cells could be in a state of 'on' or 'off'. Simple transition rules changed these states, dependent upon the states of their neighbours, which often resulted in elaborate evolutions of the system. Genetic algorithms have since been used within the strategic management community to 'breed' competitive strategies (Midgley *et al.* 1997).

The birth of agent-based models can be traced back to Schelling (1971a, 1971b) in which a 'general theory of tipping' (Schelling 1978) was developed. Schelling's model showed that if *individuals* (modelled as being on a cellular grid) had even a slight preference of being near neighbours with characteristics similar to themselves, the *system* of individuals would evolve to be segregated. This connection between 'micromotives and macrobehavior' led to a vast range of subsequent models, some of which are discussed later in this Element.

Axelrod (1997:3) describes agent-based modelling as the 'third way' of doing science:

> Agent-based modelling is a third way of doing science. Like deduction, it starts with a set of explicit assumptions. But unlike deduction, it does not prove theorems. Instead a simulation generates data that can be analyzed inductively. Unlike typical induction, however, the simulated data comes from a rigorously specified set of rules rather than direct measurement of the real world. While induction can be used to find patterns in data, and deduction can be used to find consequences of assumptions, simulation modeling can be used to aid intuition.

Bonabeau's (2002:7280) review of agent-based models as a specific method of simulation sums up the methodology quite nicely:

> In agent-based modeling (ABM), a system is modeled as a collection of autonomous decision-making entities called agents. Each agent individually assesses its situation and makes decisions on the basis of a set of rules. Agents may execute various behaviors appropriate for the system they represent – for example, producing, consuming, or selling. Repetitive competitive interactions between agents are a feature of agent-based modeling, which relies on the power of computers to explore dynamics out of the reach of pure mathematical methods.

Bonabeau goes on to describe agent-based modelling as a 'mindset more than a technology' – the fact that a system can be *conceptualized* as a group of interacting actors or agents, be these individuals within an organization or firms interacting with other firms in an industry. For Epstein (Epstein and Axtell 1996; Epstein 1999), agent-based models are 'generative' models, 'growing' societies from the bottom up.

Agent-based modelling has also been used to introduce models that are actually cellular automata models (Goldenberg *et al.* 2010) with the ensuing limitations on topology. For a discussion on the similarities and distinctions between such modelling frameworks, see Robertson (2019), and for a more general introduction to agent-based modelling, see Gilbert (2008), North and Macal (2007), Gilbert and Troitzch (2005), Axelrod and Tesfatsion (2006), Woodridge and Jennings (1995) and Robertson and Caldart (2009). A primer on agent-based modelling is given in the following subsection.

1.1 Agent-Based Modelling: a Primer

Agent-based models, also called individual-based models, are a class of computer simulations where individual components of a system are modelled so that the collective behavior of those agents comprises the behavior of the system as a whole. As an exemplar of agent-based modelling, we review a classic agent-based model, Schelling's (1969, 1971a) 'Model of Segregation'.

Schelling observed that individuals are segregated according to various attributes such as gender, race, income, language, and other classes. While some of this can be explained by the work of a central planner imposing choices on individuals, an alternative explanation exists – that the micro-level behaviour of individuals within the system generates the macro-level system behaviour, in this case the segregation of individuals.

Consider a system of N individuals, of two types. In this example, we can visualize these by supporters of political parties, the light greys and the dark greys. For reasons of modelling simplicity, these individuals are assumed to be arranged in a grid, although this assumption is not required for the segregation behaviour to manifest itself. One can visualize this system as agents of two types representing the agents' political affiliation: dark grey or light grey, as shown in Figure 1. These N individuals are distributed randomly in the grid.

These individuals or 'agents' have a limited knowledge of the world: they can observe the political affiliation of their neighbours, but not the political affiliation of all individuals in the system. In this way, they are 'intendedly rational, but boundedly so' (Simon 1997:88) – that is to say, the individual's behavior is 'subjectively rational if it maximizes attainment relative to the actual knowledge of the subject' (Simon 1947:76).

The behavior of individuals is as follows:

- Each individual observes their neighbours and calculates whether the individual is in a local minority. In the example in Figure 1(b), the dark grey individual at the centre of the large square is in a local minority, as only one out of six neighbours are of the same political affiliation (in numerical terms, $1/6 \approx 17\% < 50\%$). If the proportion of neighbours with the same political affiliation is less than 50 per cent, the individual is 'unhappy', and if this proportion is 50 per cent or above, the individual is 'happy'.

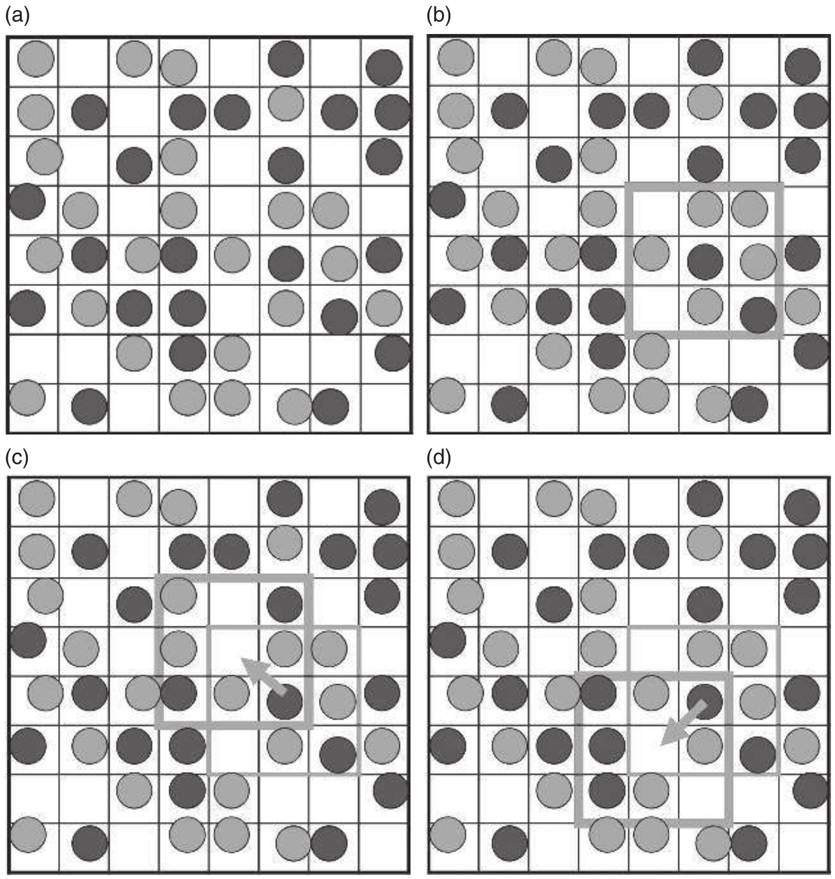

Figure 1 Movement of unhappy dark grey agent showing candidate locations

- Unhappy individuals look for a vacant location where they will be happy. In Figure 1(c), the dark grey individual tests to see whether they would be happy in the location to their north-west. In this case, they would have two out of six neighbours ($2/6 \approx 33\% < 50\%$) with the same political affiliation as their own, which means that, in this location, they would remain unhappy, and this location is not viable.
- This search continues until an individual finds a location where they are happy. In Figure 1(d), the dark grey individual investigates the location to their south-west. In this location, they will have three out of six neighbours with the same political affiliation ($3/6 = 50\% \geq 50\%$), and this individual is happy. They remain at this location.

In each round, some individuals change their state from happy to unhappy, others change their state from unhappy to happy and others remain either happy or unhappy. The model continues until all individuals are happy and equilibrium is established.

The remarkable thing about this model, over and above the micro-level actions causing the macro-level behavior of the system, is that there arises a counter-intuitive result: individuals do not have to desire to be in a local majority in order for segregation to take place.

Consider the initial conditions shown in Figure 2(a), where dark grey and light grey individuals are randomly allocated to locations within a grid. These individuals possess the same threshold of happiness, σ, defined as the ratio of similar neighbours to total neighbours, in this case, 50 per cent. Some of these individuals will be initially happy (denoted by happy faces) or unhappy (denoted by sad faces). The model is allowed to run, with unhappy individuals moving locations, until an equilibrium is reached where all individuals are happy, as shown in Figure 2(b)(i). As can be seen, the equilibrium state is

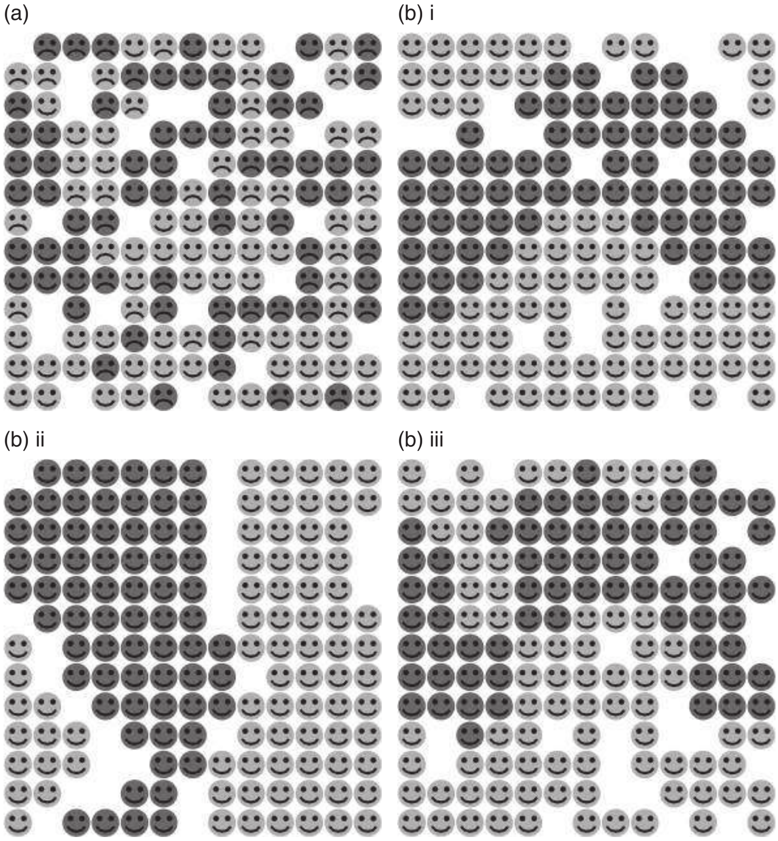

Figure 2 (a) Initial conditions with σ = 50%, (b) equilibrium conditions with (i) σ = 50%, (ii) σ = 66%, (iii) σ = 33%

segregation between dark grey and light grey individuals, *caused by* the micro-level actions of individuals.

But what if we vary σ to a higher number, say 66 per cent? If we do, the initial locations of individuals will be the same, but specific individuals now may be unhappy in their initial location (as they are subject to a higher threshold of similarity). If we allow this model to run, the equilibrium, as shown in Figure 2(b)(ii), still exhibits segregation with a more pronounced boundary between light grey and dark grey individuals. This is an expected result: with a higher threshold above 50 per cent, we would expect segregation to occur, as each individual wants to be in a local majority, and there is a clear driver of macro-level segregation.

The remarkable result from this model, however, appears when we lower the threshold σ to below 50 per cent, say 33 per cent. This means that each individual is happy to be in a local minority. And not just any minority: they are happy to be outnumbered two-to-one by individuals of a different political affiliation. Intuitively, we would expect that in such a tolerant population, segregation would not occur. However, the results shown in Figure 2(b)(iii) demonstrate otherwise: there is *still* segregation between light grey and dark grey individuals, even in a seemingly tolerant population.

This agent-based model and other similar models exhibit *emergent* behavior, which is one of the characterizing features of complex systems, discussed below.

1.2 Complexity Science

The antecedents of agent-based models in strategic management can be traced back to the methods used in complexity science. Complexity science is the study of the interactions between individual actors in a system. While complex adaptive systems have been studied generally without specific applications in mind (Miller and Page 2007), Anderson, in the special issue of *Organization Science* devoted to complexity (Anderson 1999; Anderson *et al.* 1999), produces a review of the use of complexity science in organization science and breaks down the key elements of complex adaptive systems models into the following concepts: agents with schemata, self-organizing networks sustained by importing energy, co-evolution to the edge of chaos and recombination and system evolution. These concepts, however, are not novel in other academic fields. Indeed, in the preface to the special issue Lewin (1999:215) notes:

> [M]any of these ideas are not new. However, they do not simply represent old wine in new bottles … this reframing gives me reason to believe that complexity science has much more to recommend to organization science.

Simon's definition – now more than half a century old (Simon 1962:468) – passes the test of time:

> [B]y a complex system, I mean one made up of a large number of parts that interact in a nonsimple way. In such a system, the whole is more than the sum of its parts.

We are relatively early in the journey of exploration of complexity science, particularly within strategic management. While the *NK* literature introduced in Section 3.1 has dominated agent-based modelling within organization and managerial science, we should not assume that this is the only model worthy of investigation: it is not, and there is scope for many more models to be brought into the domain of strategic management.

Bonabeau (2002) sets out the benefits of agent-based modelling as capturing emergent phenomena, providing a natural description of a system and being flexible. And it is the first concept, that of emergence, where agent-based modelling comes into its own. Emergence is a macro-level phenomenon brought about by the micro-level interaction of its constituent parts. In the preface to his book *Micromotives and Macrobehavior*, Schelling (1978) relates an anecdote in which, when he was giving a large public lecture, the first thirteen rows of an otherwise filled auditorium were left empty. When he asked his host why the seating had been arranged in this way, he was told it hadn't been: the individuals had apparently conspired to leave a gap at the front, their individual actions at the micro level combining into a noticeable effect at the macro/system level – in short, emergent behaviour. Schelling (1978:13) also makes the pertinent observation that to understand specific behaviours in this case is not to develop a specific understanding of, say, auditorium management, but rather the exact opposite: these emergent phenomena go far beyond individual realizations and are something much more fundamental. They involve:

> a kind of analysis that is characteristic of a large part of the social sciences, especially the more theoretical part. That kind of analysis explores the relation between the behavior characteristics of the *individuals* who comprise some social aggregate, and the characteristics of the *aggregate*.

Emergent behaviour can be found in a wide range of social and natural systems, for example, the flocking of birds, which can be modelled by the micro-level interactions of individual birds; the growth of crystals; convection; and the growth of cities.

The study of complex systems is inherently cross-disciplinary, exemplified by transdisciplinary research institutions such as the Santa Fe Institute (Dillon 2001), as the beauty of the study of complex systems is that the classes of problem that can be studied are not – or should not be – confined to the individual discipline being studied. In the case of strategic management, the class of problem is general in nature. The notion of sustainable competitive advantage, for example, can be thought of as a general concept, with parallels seen in the competition between species and within groups in the animal kingdom – where individuals aim to maximize not their performance but their 'fitness' – as well as in the self-organization of chemical reactions demonstrating cyclical behaviour. We should be ready to embrace these different application areas where complexity science has shown its promise and not be afraid to import ideas into management merely because a case does not match exactly the area in which we find our specific strategic and business problem.

Garcia (2005:383) summarizes agent-based models as being useful in the following circumstances:

- Where both micro and macro levels are of interest;
- Where social systems can be described by 'what if' scenarios but not by differential equations;
- When emergent phenomena may be observed;
- When co-evolving systems interact in the same environment;
- When learning and adaptation occurs;
- When physical space and temporal space are of interest;
- When the population is heterogeneous or the topology of interactions is complex and heterogeneous (e.g. social networks).

Anderson (1999) sets out a direction of travel for complexity research applied to management, citing D'Aveni's (1994) notion that the environment in which firms reside is 'hypercompetitive', explained by Anderson (1999:228) as 'nonlinearity lead[ing] both to unpredictable behavior and a rapid rate of change'. The new strategy for organizations under these conditions should therefore be to evolve temporary advantages faster than competitors (Brown and Eisenhardt 1998). Anderson also suggests that organizations can (a) attempt to alter 'the fitness landscape on which individual agents are trying to adapt, [while] strategists can change both the trajectory of emergent behavior and the diversity of behaviors in an organization's repertoire'; (b) 'reconfigure the organizational architecture within which agents adapt'; and (c) 'give executives guidelines to follow in evolving networks of agents'.

Agent-based models have now been widely used in other areas of management including innovation diffusion (Garcia and Jager 2011; Stummer *et al.* 2015), complexities in markets (Gilbert *et al.* 2007), entrepreneurship (McMullen and Dimov 2013), supply chains (Chang *et al.* 2008; Julka 2002; Kaihara 2003; Swaminathan 2007), and artificial financial markets (LeBaron 2000, 2006; Farmer and Foley 2009). They also are employed in marketing, diffusion of information and product adoption, retail location choice, inter-firm relationships and in the choice of marketing mix (Rand 2014), to name but a few more areas.

Traditional economists view the world as moving towards equilibrium or as at equilibrium – or at least as being able to *be* in equilibrium. There is a growing corpus of alternative economics, however, that shows the importance of diverging from the traditional views, and there have been calls for the use of natural science models within management by means of agent-based models (Robertson and Caldart 2008).

Holland and Miller (1991:365) describe complex adaptive systems (CAS) by defining complexity:

> [S]uch a system is complex in a special sense: (i) It consists of a network of interacting agents (processes, elements); (ii) it exhibits a dynamic, aggregate behavior that emerges from the individual activities of the agents; and (iii) its aggregate behavior can be described without a detailed knowledge of the behavior of the individual agents. An agent in such a system is *adaptive* if it satisfies an additional pair of criteria: the actions of the agent in its environment can be assigned a value (performance, utility, payoff, fitness, or the like); and the agent behaves so as to increase this value over time.

They go on:

> [A]ny given level can usually be described in terms of local niches that can be exploited by particular adaptations. The niches are various, so it is rare that any given agent can exploit all of them, as rare as finding a universal competitor in a tropical forest. Moreover, niches are continually created by new adaptations.

Allen *et al.* (2007) contrast economic general equilibrium theory with complexity science. They observe that economic theory can assume that equilibrium is the outcome of knowledge of a system by firms, as opposed to complexity science, where little needs to be known about the system as a whole.

Note that the definition of complexity has changed since the organizational ecology definitions of the 1980s (Hannan and Freeman 1984:162):

> [A]lthough the term complexity is used frequently in the literature to refer to the numbers of subunits or to the relative sizes of subunits, we use the term to refer to patterns of links among subunits. Following Simon (1962), we identify a simple structure with a hierarchical set of links . . . 'nature loves a hierarchy'.

whereas Simon (1962) highlights the *decomposability* of social systems, a topic to which we shall return later in the context of *NK* models.

1.3 Tipping Effects, Cascades, Contagion, and Emergence

Schelling (1973) builds on the game-theoretical prisoners' dilemma and differentiates between *externalities*, the fact that the payoff of an individual depends on the choice of the other player, and *internalities*, the effect of the individual's payoff on their own choice.

Schelling (1973:385) notes that in the prisoners' dilemma game, the internality and the externality are opposed, and the externality outweighs the internality. Schelling then extends the prisoners' dilemma to multiple players, with a parameter k being the number of players out of n players in a multiple-player game that play dominated choices (those that make the dominated choices are better off than those that do not). Schelling then demonstrates the effect of k by using four payoff matrices (these payoff matrices defaulting to the standard two-player payoff matrices when $k = n / 2$). Although these are not explicitly modelled as such, they are in fact multi-agent prisoners' dilemma models:

> [T]he literature on externalities has mostly to do with how much of a good or a bad should be produced, consumed, or allowed. Here I consider only the interdependence of choices to do or not to do, to join or not to join, to stay or to leave, to vote yes or no, to conform or not to conform to some agreement or rule or restriction. Joining a disciplined, self-restraining coalition, or staying out and doing what's natural is a binary choice. (Schelling 1973:382)

Schelling's concept of tipping (Schelling 1971a, 1971b, 1973, 1978) has been extended within sociology by Granovetter (1978) who developed Schelling's ideas of thresholds, while Macy (1991) built on that work, which later was popularized by Gladwell (2000).

Delre *et al.* (2007) studies the timing of new promotions from the perspective of the ability of a networked group of agents to spread an innovation (whether a technology or a product) across the network until all (or a significant percentage) of incumbents adopt the innovation as introduced, all building on the Bass diffusion model (Bass 1969). Delre *et al.* (2007) model the spread of a product across a fixed network of individuals. This is really a transition model from not-adopted to adopted on a fixed network and therefore is on the boundary of agent-based models and network models. A range of promotion strategies are modelled, comparing targeted versus mass promotion. This model is then extended by van Eck *et al.* (2011).

In order to capture the strategic management agent-based literature, the author has performed a systematic PRISMA analysis of the literature (Moher *et al.*, 2009; Vrabel, 2015) following the methodology of Utomo *et al.* (2017). Details are shown in the Appendix.