The Knowledge Engineering Review, Vol. 27:2, 123–136.
doi:10.1017/S0269888912000094

Analysis and synthesis: multi-agent systems in the social sciences

ROBERT E. MARKS

Melbourne Business School, University of Melbourne, Carlton, Vic 3053, Australia;
e-mail: robert.marks@mbs.edu

Abstract

Although they flow from a common source, the uses of multi-agent systems (or 'agent-based computational systems'—ACE) vary between the social sciences and computer science. The distinction can be broadly summarized as analysis versus synthesis, or explanation versus design. I compare and contrast these uses, and discuss sufficiency and necessity in simulations in general and in multi-agent systems in particular, with a computer science audience in mind.

1 Introduction

Although it might appear strange to computer scientists, social scientists in general—and economists in particular—hanker after the formal, mathematical rigour of closed-form, algebraic models, and their lemmas, theorems, and proofs[1]. One reason, I believe, is the possibility of obtaining not only sufficient conditions, but also necessary conditions for the existence, uniqueness, and stability of the equilibria traditionally characterized[2]. When they use computer simulation, social scientists are concerned about validating their models from historical data, and traditionally have separated 'positive' (descriptive or analytical) from 'normative' (prescriptive or synthetic) activities. Only recently has the notion of the 'design economist' gained currency, and, as we discuss, the complexity of designing markets, for instance, means that synthesis ('design') requires computer simulation even more than does analysis.

2 Sufficiency and necessity

Simulation can, in general, only demonstrate sufficiency, not necessity. Since necessity is, in general, unattainable for simulators, such proofs are also unattainable: simulation can disprove a proposition (by finding conditions under which it does not hold) but cannot prove it, absent necessity. But if there are few degrees of freedom, so that the space of feasible (initial) conditions is small, then it might be possible to explore that space exhaustively, and hence derive necessity, as we discuss below.

With some formality, it is possible to show how difficult it is to derive necessity using simulation. A mathematical 'model A' is the conjunction of a large number of separate assumptions embodied in a specific implementation, with several equations that constitute a conglomeration of hypotheses and generalizations, as well as parameters and initial conditions that must be specified. So model A comprises the conjunction $(a_1 \wedge a_2 \wedge a_3 \dots a_n)$, where $\wedge$ means 'AND', and the a_i denote the elements (equations, parameters, initial conditions, etc.) that constitute the model.

[1] This paper draws extensively on Marks (2006, 2007). I wish to thank a referee for his comments.
[2] Another is the influence of physics on early classical economics (Mirowski, 2007).

Sufficiency: If model A exhibits the desired target behaviour B, then model A is sufficient to obtain exhibited behaviour B. That is, $A \Rightarrow B$, or model A is sufficient for the existence of behaviour B. Thus, any model that exhibits the desired behaviour is sufficient, and demonstrates one conjunction of conditions (or model) under which the behaviour can be simulated. Indeed, model A might be thought of as a solution to the problem of producing behaviour B; the fact that other solutions might exist is only of secondary concern if the problem is hard. But if there are several such models, how can we choose among them? A designer would choose based on some criterion. And just what is the set of all such conjunctions (models)?

Necessity: Only those models A belonging to the set of necessary models $\boldsymbol{N}$ exhibit target behaviour B. That is, $(A \in \boldsymbol{N}) \Rightarrow B$, and $(D \notin \boldsymbol{N}) \nRightarrow B$. A difficult challenge for the social science simulator is not to find specific models A that exhibit the desired behaviour B, but to determine the set of necessary models, $\boldsymbol{N}$. Since each model $A = (a_1 \wedge a_2 \wedge a_3 \ldots a_n)$, searching for the set $\boldsymbol{N}$ of necessary models means searching in a high-dimensional space, with no guarantee of continuity, and a possible large number of nonlinear interactions among elements.

For instance, if $(D \nRightarrow B)$, it does not mean that all elements a_i of model D are invalid or wrong, only their conjunction, that is, model D. It might be only a single element that precludes model D exhibiting behaviour B. But determining whether this is so and which is the offending element is a costly exercise, in general, for the simulator. With closed-form solutions, this might not be trivial, but it might seem easier than determining the necessary set using simulation. Yaneer Bar-Yam (2006, personal communication) argues, however, that necessity is only well defined in a class of models, and that 'if anything, the necessary conditions are better defined in a discrete model (such as a simulation model) and are more difficult to establish in differential-equation models, hence the emphasis on their proofs'.

Why is knowledge of the set $\boldsymbol{N}$ of necessary models important? Without clear knowledge of the boundaries of the set of necessary models, it is difficult to generalize from simulations[3]. Only when the set $\boldsymbol{N}$ of necessary models is known to be small is it relatively easy to use simulation to derive necessity. A classic case of simulation—albeit with physical models of metal rods and brass 'atoms', not computer models—resulting in a world-shattering discovery was Watson and Crick's (1953) discovery of the structure of DNA, with their 'stereo-chemical experiments'—simulations. Note that the title of their 1953 paper included the phrase '*a* structure', not *the* structure, flagging sufficiency, not necessity (our emphasis). Experimental results from the work of others had so constrained the degrees of freedom in the space of possible structures that Watson and Crick knew when they had simulated (or solved) the structure correctly. Model-building simulations could not clinch the structure until greater congruence between the model and the observed structure of the actual molecule was shown to exist, as the future Nobel laureates emphasized in their 1953 paper. And any negative results would have meant returning to the drawing board, or in this case the brass balls and metal rods. See Marks (2003) for further discussion of this pioneering simulation. Another momentous discovery of 'sufficient' conditions was Kepler's 1605 theory of elliptical orbits (his first law of planetary motion) to explain Brahe's observations; necessity (that *only* elliptical orbits would result in the observations) would have to await Newton's laws of gravitational attraction of 1687[4].

This discussion—the absence of necessity in simulation—must appear foreign to engineers and computer scientists. Problem-solvers would in general regard having several solutions to choose from as a luxury, especially for difficult problems. The question of necessity—whether this is the entire set of possible solutions—is not in general of concern to them. This difference in emphasis reflects one distinction between social scientists and computer scientists/engineers. Social scientists

[3] Indeed, one criticism of simulation might be called the 'so-what' or 'anything-goes' critique: 'OK, you've found a model A that results in behaviour B. Now what?'.

[4] Newton knew of Kepler's explanation (from Streete, 1661) when he derived his laws. To what extent was his general derivation triggered by the sufficiency of Kepler's model?

in general seek generality of understanding, while engineers in general seek solutions, specific instances, rather than generality. Mirowski (2007) argues further that engineers are eclectic in the methods they use to seek solutions, whereas traditionally economists have sought closed-form solutions, and have eschewed simulation.

3 Analysis in the social sciences

Before reviewing the use of agent-based simulation models in one social science application—market design—we contrast analysis with design, closed-form calculations with simulation in both analysis and design, and non-agent-based simulation with agent-based simulation of analysis and design. Once the designer understands through analysis how the elements of the phenomenon of concern work together, he can ask the question of how to improve its operation: how better to design it.

In the world of analytical, closed-form solutions, there is a certain logic to the progress of research. A real-world phenomenon is observed; a need for explanation and understanding is identified; a model is built, incorporating simplifying assumptions; the model is manipulated to obtain necessary and sufficient results, and perhaps possible improvement in the operation of the system is identified, if it is a human-made system. The former part of the progress is analysis, the latter synthesis, or design, to improve some characteristic of the system or its operation. Successful analyses are published, indexed, and referenced, to be used and modified by future analysts and designers.

A common understanding of this process in general, but particularly the process of model building and deducing the system's behaviour and outcomes, means that, by and large, later researchers can stand on the shoulders of earlier researchers. With today's on-line indexing services, it is even easier to find antecedent papers, to relax an assumption or two, and to attempt to solve the ensuing model, which might (or might not) be a closer approximation to reality, or result in a better design. This process, I believe, is driven in particular directions by the mathematical tractability of particular types of model, and the relative intractability of others. (If this reminds us of the joke about the economist searching for his lost car keys under the street light, instead of in the darkness around his car, it might not be coincidental; Judd, 2006.)

3.1 Simulation and analysis

The advantage of using simulation techniques is that they provide us with light where the analytical techniques cast little or none, in our metaphorical search, so we are no longer restricted to working with models that we hope will prove tractable to our analytical tools. As computing tools (both hardware and software) have grown more powerful and user-friendly, research using simulation techniques has blossomed. Analysis of observed phenomena has not been a driving motivation of the research of computer scientists—yet they have a 50-yr history of design and invention, which continues apace (although they have from time to time looked for analogies to the natural world: neural nets mimic in some sense elements of the brain, simulated annealing mimics the thermodynamic behaviour of cooling material, and Genetic Algorithms (GA) were inspired by natural selection with sexual reproduction; Holland, 1975). Forty years ago it was possible for Donald Knuth to write an encyclopedic study of *The Art of Computer Programming* in three volumes, but such a task would be daunting now.

Moreover, as they attempt to implement automated on-line markets, computer scientists have discovered economists' work on auctions, spurred by applications of game theory to study these traditional market institutions, and to develop new, designer markets, given the opportunities of the modern technology. But the economists' focus on the characterization of equilibrium values (Mirowski, 2007) has dismayed those who need to model out-of-equilibrium interactions and the dynamics associated with real-world markets. Can we surmize that characterizing asymptotic equilibrium is similar to looking for the lost keys under the street light—it is easier to do ('tractable') than the more realistic task of characterizing out-of-equilibrium adjustment? But this characterization is just what those designing markets—especially fully automated markets—need to have. The economists' understanding has proved disappointing[5].

Analysis precedes design: in order to change systems, we must first understand them. Even so, direct design or optimization requires a degree of understanding of the mapping from the design space to the performance space that has rarely been developed. Indeed, given the complexity of market phenomena, direct design might seldom be possible, as Edmonds and Bryson (2003) remind us. Instead, searching the design space will be an iterative process of analyzing the performance of a specific model, modifying the model in the light of this analysis, and analyzing the modified model, until the designer is happy with the performance of the manyfold-modified model against various criteria.

The social sciences are concerned with understanding or explaining social phenomena. Broadly, this process entails building a model of aspects of the real-world phenomenon under review and manipulating the model to see whether the model behaviour matches (at least in a stylized fashion) the observed real-world phenomenon. This gross simplification of the process overlooks the attempts to derive alternative explanations (theories) from the models, in order to test them against the historical data.

The issue of programme verification and empirical validation (or what Midgley *et al.* (2007) call 'assurance') of the model is a lively topic in computational economics, at least. We discuss this further below.

To achieve these ends, the researcher needs three things, at least: historical data; the ability to abstract from reality in building the (necessarily simplified) model of the real-world phenomenon; and the skill to derive implications of the model that can be tested by examining the historical data to determine which possible explanation provides the 'best' fit with the data, suitably defined.

The purpose is not solely understanding for its own sake, but prediction. It is true that accurate prediction does not require clear understanding (Friedman, 1953), and that on the other hand clear understanding does not necessitate accurate predictions, but nonetheless better understanding is more likely to result in better (more accurate) predictions, *ceteris paribus*. This process has several names: analysis, explanation, understanding.

Axtell *et al.* (1996) introduce the term 'docking' when a second team attempts to replicate another team's simulation model. They clarify three decreasing levels of replication: 'numerical identity', 'distributional equivalence' (the results, although not identical, cannot be distinguished statistically), and 'relational equivalence' (the same qualitative relationships).

3.2 Validation

What is a good simulation? The answer to this question must be: a good simulation is one that achieves its aim. But just what the aim or goal of a simulation might be is not obvious. There are several broad possibilities[6]. A simulation might attempt to explain a phenomenon; it might attempt to predict the outcome of a phenomenon; or it might be used to explore a phenomenon,

[5] Mirowski (2007) argues that for a 150 years economists have focused on the agents (buyers and sellers) who exchange, and have ignored the structure and procedures of the market in which the exchange occurs. This would explain the *lacunæ* in the economics literature that confront computer scientists when they seek detailed analysis and explanation of the workings of historical markets in order to implement automated markets of various kinds.

[6] Haefner (2005) lists seven possible goals: usefulness for system control or management, understanding or insights provided, accuracy of predictions, simplicity or elegance, generality (number of systems subsumed by the model), robustness (insensitivity to assumptions), and low cost of building or maintaining the model. Axelrod (2006) also lists seven: prediction, performing tasks, training, entertaining (see those ubiquitous games consoles), educating, existence proofs, and discovery; prediction, existence proofs, and discovery are the main scientific contributions.Rubinstein (1998) lists four purposes: predicting behaviour, as a normative guide for agents or policymakers, sharpening economists' intuitions when dealing with complex situations, and establishing 'linkages' between economic concepts and everyday thinking. Burton (2003) lists three questions: asking 'what is', 'what could be', and 'what should be'.

to play, in order to understand the interactions of elements of the structure that produces the phenomenon. Or it might be used to test a solution to a problem, such as automated allocation of resources in real time.

Consider a retail market for branded coffee at a specific supermarket (Midgley *et al.*, 1997). If the profit-maximizing behaviour of the simulated brand managers, together with some external or internal factors, led to behaviour qualitatively or quantitatively similar to the historical behaviour of the brands—their weekly prices and quantities sold[7]—then we would have obtained one possible explanation. The issue of degrees of similarity is one of closeness of fit, and could be handled using statistical measures. Note, following Durlauf (2005), that by making the assumption of profit maximizing, we are going beyond merely seeking a set of equations exhibiting periodicity similar to the 'rivalrous dance' of the historical brands.

For prediction, sufficiency suffices: there is no need to know which if any alternate conditions will also lead to the observed endogenous behaviours. That is, prediction does not require an understanding of necessity of the underlying exogenous variables. This might explain, as Friedman (1953) argues, that economic actors can behave as though they have solved complex optimization problems, even though they remain ignorant of any formal representation of the problem or its solution.

3.2.1 Formalization of validation

Let set $\mathbf{M}$ be the exhibited outputs of a model of a historical phenomenon in any period. Let set $\mathbf{S}$ be the specific, historical output of the real-world system in any period. Let set $\mathbf{Q}$ be the intersection, if any, between the set $\mathbf{M}$ and the set $\mathbf{S}$, $\mathbf{Q} \equiv \mathbf{M} \cap \mathbf{S}$. We can characterize the model output in five cases[8]:

a. If there is no intersection between $\mathbf{M}$ and $\mathbf{S}$ (i.e. $\mathbf{Q} = \emptyset$), then the model is *useless*.
b. If the intersection $\mathbf{Q}$ is not null, then the model is *useful*, to some degree. In general, the model will correctly exhibit some real-world system behaviours, will not exhibit other behaviours, and will exhibit some behaviours that do not historically occur. That is, the model is both incomplete and inaccurate.
c. If $\mathbf{M}$ is a proper subset of $\mathbf{S}$ (i.e. $\mathbf{M} \subset \mathbf{S}$), then all the model's behaviours are correct (match historical behaviours), but the model does not exhibit all behaviour that historically occurs. The model is accurate but *incomplete*.
d. If $\mathbf{S}$ is a proper subset of $\mathbf{M}$ (i.e. $\mathbf{S} \subset \mathbf{M}$), then all historical behaviour is exhibited, but the model will exhibit some behaviours that do not historically occur. The model is complete but *inaccurate*.
e. If the set $\mathbf{M}$ is equivalent to the set $\mathbf{S}$ (i.e. $\mathbf{M} \Leftrightarrow \mathbf{S}$), then (in your dreams!) the model is complete and accurate.

By *incomplete*, we mean that $\mathbf{S} \backslash \mathbf{Q}$ is non-null, so that the model does not exhibit all observed historical behaviours. By *inaccurate*, we mean that $\mathbf{M} \backslash \mathbf{Q}$ is non-null, so that the model exhibits behaviours that are not observed historically. Haefner (2005) notes that the set boundaries might be fuzzy: not 'in' or 'out', but contours of the probability of belonging to the set.

One goal of the modeller might be to attempt to construct and calibrate the model so that $\mathbf{M} \approx \mathbf{Q} \approx \mathbf{S}$ (case e): there are very few historically observed behaviours that the model does not exhibit, and there are very few exhibited behaviours that do not occur historically. The model is close to being both complete and accurate, for explanation. But this might be overfitting for prediction. In practice, a modeller examining sufficient conditions (existence proofs) for previously unobserved (counterfactual) behaviour might be happier to achieve case d, where the model is complete (and hence provides sufficiency for all observed historical phenomena), but not always accurate (i.e. sometimes produces ahistorical output). Of course, changing the model's parameters will in general change the model's exhibited behaviour (set $\mathbf{M}$). In the calibration stage, we might

[7] Midgley *et al.* (1997) and Marks *et al.* (2006) describe how they calculate each brand's weekly profits, given the combination of marketing actions of all brands that week, and with prior knowledge of the brands' costs.
[8] This conceptual framework was introduced by Mankin *et al.* (1977).

well be happier if we could adjust the model parameters so that $\mathbf{M} \approx \mathbf{S}$, in the belief that the changed set $\mathbf{M}'$ with different parameters might well model a (future) variant of historical reality.

4 Economists' uses of multi-agent systems

In economics, computer simulations have been increasingly popular in analyzing oligopoly behaviour—oligopolies are markets with small numbers of sellers, each of whose actions in general affect the outcomes for other sellers, as well as for itself. This is strategic interaction, and has been analyzed using game theory, the best framework for analyzing interactions where each player's outcome is a function of its own actions and those of other players. Some oligopoly interactions have the character of an iterated Prisoner's Dilemma (IPD), which has attracted computer scientists.

Except for the simplest markets, the interactions of asymmetric agents lead to a degree of complexity that precludes closed-form analyis of markets, let alone direct design of such markets. Instead, in both economics and computer science, multi-agent systems ('agent-based computational systems') have been enlisted to enable analysis (and design) of such markets. That is, unlike top-down models of sets of difference/differential equations, the model is analyzed from the bottom up by simulating the behaviour of interacting agents, buyers and sellers, a more realistic model, *a priori*.

The GA was the first multi-agent system used to simulate and analyze oligopolistic behaviour (Marks, 1989). The GA had been developed and pioneered by computer scientists and engineers who were intent on solving optimization problems exhibiting 'rugged landscapes' (Kauffman, 1995). Although it was at first used only where these were static, where the landscape did not change as the process of genetic 'learning' took place, it also turned out to be well suited to simulating and solving problems where the environment was changing. When the individual agents modelled by the GA are competing against each other, the GA is modelling the process of co-evolution. GAs were originally used as means of seeking optimal solutions to static problems; Marks (1989) and others adapted them to seek solutions of co-evolutionary strategic problems, such as the IPD and oligopolies with asymmetric players, where the fitness of an agent depends on the individual agents' actions, that is, the state of the whole population of agents.

When the interacting parties—the players—face identical payoff sets and choose from identical action sets, a single population is satisfactory, since the GA processes (selection, crossover, and mutation) that model learning among the individuals and between generations of the population are focused to the same end: faced with the same state of the interaction, any of the players would behave identically, and fitness is average (or discounted) profit.

But when modelling asymmetric oligopolistic players who have distinct payoff sets (because of distinct costs, facing distinct demands, and perhaps with distinct action sets), a single population of agents means that the GA processes are faced with a fitness landscape that is not only possibly rugged, but also shifting (as each agent wears a distinct seller's hat, as it were). In this case, separate populations of sellers makes practical sense. It also makes realistic sense. A single population GA was acceptable when the players were not differentiated and when the flow of information from parents to offspring at the genotype level was not an issue (Vriend, 2000), but when the players are modelling heterogeneous actors—in realistic co-evolution, for instance—each player requires a separate population, not least to prevent the modelling of illegally collusive, extra-market transfers of information.

Vriend (2000) draws the distinction between the social learning of the GA (whereby the individuals in the population have learned from their parents, through selection and crossover, and so there is the possibility of good 'genes' spreading through society over several populations) and the individual learning of non-GA agent-based models (with explicit learning incorporated into the structures of the artificial, adaptive agents). Both sorts of models, and both sorts of learning, have been termed 'agent-based' models.

As well as GAs, agent-based models in economics have increasingly included simulations in which agents 'learn' individually, rather than the population learning of the GA, and in which agents are explicitly modelled as boundedly rational, in various ways. Indeed, economists have explored the relationships between so-called zero-intelligence (stochastic-choice) traders (Gode & Sunder, 1993) and the structure and rules of markets. This research has also been undertaken in computer science (Cliff, 2001).

5 Synthesis in the social sciences

We focus on market design, since this is a field of great interest to economists as new markets have been devised and implemented, and also to computer scientists, attempting to implement market allocation mechanisms, which allow interaction between non-human agents. Automated contract design is also a new field, in which computer scientists have made advances (see Jennings *et al.*, 2001).

As engineers say, after analysis comes synthesis—design. Designing markets is a new discipline. At least five examples of designed market can be identified: simulated stock markets; emission markets; auctions for electro-magnetic spectrum; electricity markets; and on-line, e-commerce markets.

First, the markets for new financial instruments, derivatives, that were created and traded after Black, Scholes, and Merton solved the 70-yr-old problem of pricing options. Previously, financial traders understood that options were valuable, but not how to value them exactly. More recently, there has been research into the rules and micro-structure of stock markets, continuous double-auction trading, through the use of simulated markets. See LeBaron (2006) for further discussion of this research.

Second, the markets for pollution emissions, usually sulphur dioxide and carbon dioxide. The realization that the emissions from industrial processes in particular, and the emission of anthropogenic chemicals into the environment in general, were, at least potentially, altering the biosphere for the worse was followed only after a lag with the awareness by policy-makers that market mechanisms could be harnessed to control such emissions, generally more efficiently than could other mechanisms.

Third, the auctions for electro-magnetic spectrum. The simultaneous ascending-bid auctions that have recently been designed for selling bands of local spectrum to be used for new communications technologies did not arise without some hiccups. Perhaps as an offshoot of the privatization of government assets and activities in the 1980s in many countries, the use of auctions to choose the new owners and to value these assets slowly replaced so-called 'beauty contests', in which subject to certain technical requirements licenses were virtually given away. But these new auction mechanisms at first did not allow for the complementary nature of bands in different localities. Only after intensive efforts by economists advising, first, governments, and, second, bidding companies did the successful '3G' auctions occur (Milgrom, 2004).

Fourth, the markets for the exchange of electricity. Again, as a consequence of the twin political aims of privatizing government-owned electricity utilities and of improving the efficiency of electricity generation and distribution systems (perhaps by separating ownership of generators and distributors), while reducing the bureaucratic weight of regulation even on privately owned utilities, there has in many jurisdictions been a move away from centralized engineering-dominated means of allocating electricity load across generators and distribution networks to using market mechanisms of various kinds. Electricity cannot (easily or cheaply) be stored, a characteristic which, with some engineering issues, has meant that previously existing market mechanisms were not appropriate. Instead, several types of new market mechanisms have been introduced (see Marks, 2006).

Fifth, on-line markets. With the growth of the use and extent of the Internet over the past 8 yr, and the dot-com boom, with buying and selling on-line, opportunities for designing on-line markets *de novo*, as opposed to trying to emulate existing face-to-face markets, have arisen. In the last few years these opportunities have given rise to much work by computer scientists, as well as economists. Indeed, there is a productive research intersection of the two disciplines, as many readers would know.

Mirowski (2007) argues that the traditional reduced form of market clearing, and the focus on agents rather than markets, mean that economists are ill-prepared to offer advice on how to design new markets, either to policy-makers or to computer scientists. Indeed, there were salutary tales when spectrum markets were first established, under rules that have since been abandoned (McMillan, 1994).

As remarked by Roth (1991) in an early paper on market design, three approaches are suitable for the iterative process of market design: first, traditional closed-form game-theoretic analysis, as discussed above; second, human-subject experiments (see Brenner, 2006; Duffy, 2006); and, third, computational exploration of different designs. Indeed, if the design criteria are clearly defined, some of the recent techniques of simulation and optimization developed by computer scientists and computational economists can be used to search for optimal market designs, directly and indirectly.

Historical market institutions have in general not simply been imposed from above (so-called top-down design) but have also emerged from the bottom up as a consequence of a multitude of actions and interactions of the myriad traders (McMillan, 2002). Although the omnipotent programmer can experiment with different market forms and different kinds of boundedly rational agents to discover sufficient combinations of each for specific behaviour of the market, evolutionary computation raises the possibility of bottom-up design, or emergence of market design through simulation. As Mirowski (2007) argues, a typology of markets, characterized by their structures and rules, would be very useful for the designers. Unfortunately, the profession is far from accomplishing this, perhaps because of the complexity of markets.

This in turn raises the issue of whether agent-based experiments are being used as a model of human behaviour (where analysis is followed by design, given the behaviour of the agents and the emergent aggregate outcomes)—in which case it is an empirical question as to how boundedly rational the agents should be to best model human agents (Duffy, 2006)—or whether the agent-based experiments are an end in themselves, because on-line it is possible to use agents ('buy-bots', 'sell-bots') to buy and sell, without the errors that human agents are heir to. How has this research proceeded?

5.1 Market design

Design is a process of building directed by the pre-specified design objectives, if not by an explicit how-to plan. Unfortunately, specifying objectives does not always immediately delineate exactly how the model building should occur: these objectives are specified in a performance space (or behaviour space) and the building occurs in a design space. The mapping from the designed structure to the desired performance may not be clear.

In the case of evolution, the design occurs in the genome space, while the behaviour or performance occurs in the phenome space. In the case of designer markets, policy-makers have been using theory, experiments with human subjects, and computer simulations (experiments) to reduce the risk that the mapping from design (structure and rules) to behaviour of the economic actors (the performance of the system) is poorly understood, and so that there are fewer surprises. Where the mapping is sufficiently well understood, and where closed-form analytic solutions are tractable, it should be possible to describe not only sufficiency—if the market has this structure, and the rules of trading are such and such and the traders are given this information, then this performance and behaviour will follow, at least in general form—but also necessity—if you want this performance and behaviour, then this is the only set (or sets) of designs (combinations of structure and rules) that will produce it.

MacKie-Mason and Wellman (2006) present a Marketplace Design Framework, which delineates the three fundamental steps that constitute a transaction: first, the connection (searching for and discovering the opportunity to engage in a market interaction); second, the deal (negotiating and agreeing to terms); and, third, the exchange (executing a transaction). They define a 'marketplace system' as consisting of agents and the market mechanism through which they interact, all embedded in an environment of social institutions (language, laws, etc.). Their market mechanism is the set of 'rules, practices, and social structures of a social choice process, specifying, first, permissible actions' (including messages), and, second, market-based exchange transactions as outcomes of a function of agent messages. If there is some entity, apart from the participating agents, that manages any inter-agent communication and implements the mechanism rules, then the market mechanism is mediated.

MacKie-Mason and Wellman note that, as a consequence of this characterization of a marketplace, there are at least two design decisions: first, the design of the market mechanism, which might be decomposed into the design of mechanisms for, successively, the connection, the deal, and the exchange phases of a transaction; and, second, design of agents to interact with the market mechanism, whether existing or newly designed. They define an agent as an 'autonomous decision-making locus in a system of multiple decision-making entities'; an agent has 'type' attributes, such as preferences, beliefs, intentions, and capabilities. There will be a form of consistency between the agents' behaviour, beliefs, and preferences, consistent with some principle of rationality. Here, the focus is on design of MacKie-Mason and Wellman's market mechanism, specifically, the deal negotiation task. As with most of the existing literature, we here focus on market mechanisms that govern the settlement from allowable actions.

Mechanisms specify, first, the agents' concerns that are recognized, and, second, rules mapping actions into allocation outcomes. A rule might specify which actions are permissible, or the procedure for choosing a settlement of agents' concerns, based on observable actions. For instance, auctions, MacKie-Mason and Wellman point out, include rules governing allowable actions, and rules governing settlement.

To be effective, design of the market mechanism must be measured, and will usually consist of a constrained optimization, even if not explicitly or directly. 'No external subsidies' or 'maintain horizontal equity' are two possible constraints given by MacKie-Mason and Wellman. We explore others below. The general design problem has become designing a market mechanism that includes defining a set of concerns over which agents can interact, specifying rules of permissible actions, and rules for mapping from actions to settlement and outcomes.

5.1.1 Complexity of design

Edmonds and Bryson (2003) speak of the 'syntactic complexity' of design. This is the lack of a clear mapping from design to behaviour (or vice versa): the only way to know the system's outcomes is to run the system and observe the emerging performance. Analysis is in general not able to predict the outcome. They are speaking of multi-agent computer systems, but could be speaking of standard double auctions in continuous time, which have not yet been solved analytically. Simon (1996) put it this way: '… it is typical of many kinds of design problems that the inner system consists of components whose fundamental laws of behavior … are well known. The difficulty of the design problem often resides in predicting how an assemblage of such components will behave'[9]. Nonetheless, Byde (2006) uses evolutionary simulations applied to a space of auction mechanisms to derive novel mechanisms for sealed bid auctions.

One reason why analytical methods of analysis might fail is that the mapping from initial conditions of structure and rules to behaviour and performance is not smooth or continuous, and, as such, is not amenable to calculus-based tools. The rugged nature of this landscape is its complexity, a complexity that is multiplied if it too is changing, perhaps as a function of the strategic complexity that occurs if the design has also to account for the interacting agents' patterns of behaviour changing as a result: the biologist's co-evolution.

5.1.2 Design trade-offs

Where there are several design criteria, the possibility arises of trade-offs among the criteria. For instance, if a firm has market power, it can maximize its seller revenue, but at the cost of market efficiency, as measured by the sum of seller (or producer) surplus and buyer (or consumer) surplus[10]. Or it might be possible to improve the fairness of a market outcome, but at the cost of market

[9] We can see this as referring to elements a_i and model A, respectively, of Section 2 above.

[10] Mirowski (2007) is skeptical about the meaning of this measure when the agents are artificial.

efficiency. To use computer simulation, such trade-offs must be explicit. It might be possible to use a version of Simon's (1996) 'satisficing', whereby so long as the other criteria are met (above some target level), the remaining criterion is used to rank designs. Or different criteria could be weighted to derive a single, scalar maximand.

LeBaron (2006), in examining the use of agent-based models of financial markets, discusses seven basic design questions for his models, which translate across to more general models. First, the economic environment itself needs to be resolved: What will be traded? Second, how are agents' preferences to be modelled? What particular functional forms will be used, such as mean-variance, constant absolute risk aversion, myopic or inter-temporal? Or will specific behavioural rules simply be evaluated directly? Third, market clearing and price formation need to be modelled. Fourth, the fitness of the model must be evaluated. For example, should wealth or utility be used? And should the evolving behavioural rules to which fitness measures are applied be forecasts, demands, or some other type of action? Fifth, how is information to be processed and revealed? Sixth, how does learning occur? Is it social or is it individual? Seventh, how is benchmarking to be undertaken? While these questions relate to the models used to design markets, they may also reflect on the design criteria for the final designer markets.

5.2 Explicit use of agents

It is interesting to note that three of LeBaron's seven design questions above (2, 5, 6) refer to the agents, the rest to the market institution itself. It is possible to design without the use of agents: given a market with demand and supply schedules, economic efficiency is maximized at the output level where marginal value equals the marginal unit cost, no matter how the social surplus is divided between buyers and sellers. But such direct design (optimization) requires a well-defined problem. With several design trade-offs and the possible emergence of unforeseen performance in the system, agent-based analysis and design, in which the market system can be modelled as 'evolving systems of autonomous, interacting agents with learning capabilities' (Koesrindartoto & Tesfatsion, 2004), is increasingly employed.

LeBaron (2006) places some weight on how actual trading occurs: the institutions under which trading is executed. He argues that agent-based models are well suited to examining market design and micro-structure questions because, first, they can produce a large amount of data, and, second, they allow testing of market design in a heterogeneous, adaptive environment. Computer scientists designing automated negotiation and exchange mechanisms (markets) would likely agree.

6 Similarities and differences

Recently, software engineers have been designing systems of exchange, of markets. Their designs—of distributed computing systems, and on-line trading in real time—have begun to borrow from economists' insights into how traditional face-to-face markets have evolved to operate. They have also (Phelps *et al.*, 2002) begun to realize that the equilibrium characterizations of mathematical economics do not always provide the answers they need in designing their on-line markets, which will be in disequilibrium almost always if trading in real time. That is, the adjustments of the operation of the markets to the current equilibrium (or attractor) will almost never happen fast enough to reach equilibrium, especially when the location of the attractor is continuously changing. This means that the economist's characterization of the equilibrium may be of little help in understanding and designing the operation of such markets.

The shortcomings of these results from equilibrium analyses of economic mechanisms have been underlined by Roth (2000, 2002) in two papers that begin to point the way forward for market design, with the economist as engineer. Indeed, Roth makes the point that, as engineers have learned to borrow from the insights of physics, so can the design economist use insights not only from equilibrium mathematical economics, but also from computer science.

When, however, these insights are curtailed, perhaps by the tractability of closed-form analytical methods, both economists and software engineers have been using simulation in analysis, to obtain sufficient, but rarely necessary, conditions. But then a single sufficient condition (solution) may be enough. Simulation has occurred using GAs, numerical solutions, and explicit agent-based models. Iterative analysis has been used as a means of designing systems.

LeBaron (2006), in his conclusion, lists some criticisms of the agent-based approach to modelling financial markets. Some (such as too few assets considered, questions of timing ignored) are more specific to the models he examines, but several are relevant to more general market models: too many parameters; questions about the impacts of the introduction of new trading strategies on the stability of trading; sensitivity to the number of agents trading; over-reliance on inductive models of agents, which respond to past rules and forecasts; and not enough reliance on deductive models of agents, which might learn commonly held beliefs about how markets work.

Epstein (2006) answers the criticisms of some economists that 'anything-goes' with computer simulations by stating that, in fact, the computer algorithms controlling computer simulations can be written as mathematical expressions. His argument is that every agent model is a computer program, and hence is Turing computable. But for every Turing machine, there exists a unique corresponding and equivalent *partial recursive function*. Such functions might be extremely complex and difficult to interpret, but they exist. Hence he argues for considering computer simulations as 'recursive' or 'effectively computable' or 'constructive' or 'generative' (after Chomsky) social science.

7 Borrowings

There has been some cross-fertilization between computer science and the social sciences. Researchers from both sides have considered what attributes simulated agents should have (Gilbert & Troitzsch, 2005). Wooldridge and Jennings (1995) would give computer agents these six properties: first, autonomy: no other agents can control any agent's actions and internal state; second, social ability: agents can interact and communicate with other agents; third, reactive: agents perceive their environment and respond; fourth, pro-active: agents initiate goal-directed actions; fifth (intentionality: agents' metaphors of beliefs, decisions, motives); and sixth (even agent emotions). (I put these last two in parentheses because they are even more removed from traditional thinking in economics than is agent-based computer simulation.)

Epstein (1999) would desire these four qualities in addition: first, heterogeneity: not 'representative' agents, but agents that differ: this is very easy to achieve, and is a natural function of agent-based modelling (Marks, 2010); second, local interactions: agents operate in a defined space; third, boundedly rational (Simon, 1982) agents: they perceive information, possess memory, and some computational capacity: this is explicit in programming agents, of course; and fourth, ability of the models to deal with non-equilibrium dynamics: large-scale transitions, tipping phenomena.

Borrowing from computer science and biology, the first agents in economics models were the members of the populations bred using the GA; they can be characterized as deductive: backwards-looking, stimulus-response automata. Borrowing from models developed in psychology, Arthur (1991, 1993) introduced Reinforcement Learning for individual agents, in which an agent is more (less) likely to choose an action in the future if it has been more (less) successful using it in the past. Arthur is able to calibrate his models using data from human experiments. But these agents only learn by responding to past actions and payoffs, with no attempt to anticipate and reason back inductively.

More recently, agents have been modelled with forward-looking, deductive learning: Gjerstad and Dickhaut (1998) model agents that form beliefs about other agents' likely actions, and so respond to their beliefs: their 'heuristic belief learning'. Computer scientists Wellman *et al.* (2007) highlight this model in their analysis of the successful bidding models that have emerged in the Trading Agent Competition[11]. Duffy (2006) surveys the growing literature on the links between

[11] This study follows earlier tournaments involving competing agents, which have resulted in new insights: political scientist Axelrod (1984), marketers Fader and Hauser (1988), and many computer scientists and social scientists in the past 30 years.

laboratory experiments with human subjects and computer experiments with agents that learn. See also Arifovic and Ledyard (2008).

As further problems are addressed and solved, we might expect more borrowings to appear across the divide, perhaps even from this Special Issue.

8 Conclusion

The practical design of markets—mechanism design—using the tool of agent-based simulation is emerging from its infancy. On the one hand, there are mechanisms, such as monopoly auctions, that have been in use since antiquity (McMillan, 2002, p. 69) without much self-conscious design effort. On the other, recent advances in theory and computation have allowed analysis and design to derive new or better mechanisms. The iterative analysis of electricity markets with agent-based models is <20 yrs old, and the work by computer scientists on automated markets is even newer. Only recently have there been attempts to use such models, after parameterizations of auctions, to directly design markets, including electricity markets. Indeed, direct market-design modelling attempts have only occurred in the last several years. Clearly, we have further to travel down this road, as Roth's (2002) notion of the design economist emerges from the work of many modellers, in economics, engineering, and computer science.

Acknowledgement

The author thanks Peter McBurney and the conveners of the First International Workshop on Market-Based Control of Complex Computational Systems, Liverpool, September 1–2, 2008, when an early version of this paper was first presented.

References

Arifovic, J & Ledyard, J 2008. *Individual Evolutionary Learning and the Voluntary Contributions Mechanism*, Working Paper, California Insititute of Technology.

Arthur, W. B. 1991. Designing economic agents that act like human agents: a behavioral approach to bounded rationality. *American Economic Review* **81**, 353–359.

Arthur, W. B. 1993. On designing economic agents that behave like human agents. *Journal of Evolutionary Economics* **3**, 1–22.

Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books.

Axelrod, R. 2006. Advancing the art of simulation in the social sciences. In *Handbook of Research on Nature-Inspired Computing for Economy and Management*, Rennard, J.-P. (ed.). Idea Group Inc., 90–100.

Axtell, R., Axelrod, R., Epstein, J. & Cohen, M. D. 1996. Aligning simulation models: a case study and results. *Computational and Mathematical Organization Theory* **1**, 123–141.

Brenner, T. 2006. Agent learning representation: advice on modelling economic learning (Ch. 18). In *Handbook of Computational Economics, 2: Agent-Based Modeling*, Tesfatsion, L. & Judd, K. L. (eds). Elsevier Science, 895–947.

Burton, R. M. 2003. Computational laboratories for organizational science: questions, validity and docking. *Computational & Mathematical Organizational Theory* **9**, 91–108.

Byde, A. 2006. Applying evolutionary search to a parametric family of auction mechanisms. *Australian Journal of Management* **31**(1), 1–16. `http://www.agsm.edu.au/bobm/eajm/0606/01-byde.html`

Cliff, D. 2001. *Evolutionary Optimization of Parameter Sets for Adaptive Software-agent Traders in Continuous Double Auction Markets*. Hewlett-Packard Technical Report HPL-2001-99, November.

Duffy, J. 2006. Agent-based models and human subject experiments (Ch. 19). In *Handbook of Computational Economics, 2: Agent-Based Modeling*, Tesfatsion, L. & Judd, K. L. (eds). Elsevier Science, 949–1011.

Durlauf, S. 2005. Complexity and empirical economics. *The Economic Journal* **115** (June), F225–F243.

Edmonds, B. & Bryson, J. J. 2003. *Beyond the Design Stance: the Intention of Agent-based Engineering*. Centre for Policy Modelling, CPM Report no.: CPM-03-126. `http://cfpm.org/papers/btds.pdf`

Epstein, J. M. 1999. Agent-based computational models and generative social science. *Complexity* **4**(5), 41–60.

Epstein, J. M. 2006. Remarks on the foundations of agent-based generative social science (Ch. 34). In *Handbook of Computational Economics, 2: Agent-Based Modeling*, Tesfatsion, L. & Judd, K. L. (eds). Elsevier Science, 1585–1604.

Fader, P. S. & Hauser, J. R. 1988. Implicit coalitions in a generalized prisoner's dilemma. *Journal of Conflict Resolution* **32**, 553–582.

Friedman, M. 1953. *Essays in Positive Economics*. University of Chicago Press.

Gilbert, N. & Troitzsch, K. G. 2005. *Simulation for the Social Scientist*, 2nd edn. Open University Press.

Gjerstad, S. & Dickhaut, J. 1998. Price formation in double auctions. *Games and Economic Behavior* **22**, 1–29.

Gode, D. & Sunder, S. 1993. Allocation efficiency of markets with zero intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy* **101**, 119–137.

Haefner, J. W. 2005. *Modeling Biological Systems: Principles and Applications*, 2nd edn. Springer.

Holland, J. H. 1975. *Adaptation in Natural and Artifical Systems*. University of Michigan Press.

Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S., Sierra, C. & Wooldridge, M. 2001. Automated negotiation: prospects, methods, and challenges. *International Journal of Group Decision and Negotiation* **10**(2), 199–215.

Judd, K. L. 2006. Computationally intensive analyses in economics (Ch. 17). In *Handbook of Computational Economics, 2: Agent-Based Modeling*, Tesfatsion, L. & Judd, K. L. (eds). Elsevier Science, 881–893.

Kauffman, S. A. 1995. *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. Oxford University Press.

Knuth, D. E. (1968–1973). *The Art of Computer Programming*. Addison-Wesley.

Koesrindartoto, D. & Tesfatsion, L. 2004. Testing the reliability of FERC's Wholesale Power Market Platform: an agent-based computational economics approach. In *Energy, Environment and Economics in a New Era, Proceedings of the 24th Annual North American Conference of the United States Association for Energy Economics and the International Association for Energy Economics*, 8–10 July 2004, Washington, DC. IAEE/USAEE.

LeBaron, B. 2006. Agent-based computational finance (Ch. 24). In *Handbook of Computational Economics, 2: Agent-Based Modeling*, Tesfatsion, L. & Judd, K. L. (eds). Elsevier Science, 1187–1233.

MacKie-Mason, J. K. & Wellman, M. P. 2006. Automated markets and trading agents (Ch. 28). In *Handbook of Computational Economics, 2: Agent-Based Modeling*, Tesfatsion, L. & Judd, K. L. (eds). Elsevier Science, 1381–1431.

Mankin, J. B., O'Neill, R. V., Shugart, H. H. & Rust, B. W. 1977. The importance of validation in ecosystem analysis. In *New Directions in the Analysis of Ecological Systems, Part 1, Simulation Council Proceedings Series*, 5, Innis, G. S. (ed.). Simulation Councils, 63–71. Reprinted in Shugart H. H. and O'Neill, R. V. (eds), *Systems Ecology*, (Dowden, Hutchinson and Ross), 1979, 309–317.

Marks, R.E. (1989). Breeding optimal strategies: optimal behavior for oligopolists. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, George Mason University, June 4–7, 1989, Schaffer J.Ď. (ed.). Morgan Kaufmann Publishers, 198–207.

Marks, R. E. 2003. Models rule. *Australian Journal of Management* **28**(1), i–v. http://www.agsm.edu.au/bobm/eajm/0306/edit.html

Marks, R. E. 2006. Market design using agent-based models (Ch. 27). In *Handbook of Computational Economics, 2: Agent-Based Modeling*, Tesfatsion, L. & Judd, K. L. (eds). Elsevier Science, 1339–1380.

Marks, R. E. 2007. Validating simulation models: a general framework and four applied examples. *Computational Economics* **30**(3), 265–290.

Marks, R. E. 2010. *Learning to be Risk-Neutral: CRRA Agents Evolve to be Risk-Neutral, even with the Possibility of Bankruptcy*. Mimeo. http://www.agsm.edu.au/bobm/papers/ralet.pdf

Marks, R. E., Midgley, D. F. & Cooper, L. G. 2006. Co-evolving better strategies in oligopolistic price wars. In *Handbook of Research on Nature-Inspired Computing for Economy and Management*, Rennard (ed.). Idea Group Inc., 806–821.

McMillan, J. 1994. Selling spectrum rights. *Journal of Economic Perspectives* **8**(3), 145–162.

McMillan, J. 2002. *Reinventing the Bazaar: A Natural History of Markets*. Norton.

Midgley, D. F., Marks, R. E. & Cooper, L. G. 1997. Breeding competitive strategies. *Management Science* **43**(3), 257–275.

Midgley, D. F., Marks, R. E. & Kunchamwar, D. 2007. The building and assurance of agent-based models: an example and challenge to the field. *Journal of Business Research* (Special Issue: Complexities in Markets), **60**, 884–893.

Milgrom, P. 2004. *Putting Auction Theory to Work*. Cambridge University Press.

Mirowski, P. 2007. Markets come to bits: evolution, computation and markomata in economic science. *Journal of Economic Behavior & Organization* **63**(2), 209–242.

Newton, I. 1687. *Philosophiæ naturalis principia mathematica*. London: Jussu Societatis Regiae ac Typis J. Streater.

Phelps, S., McBurney, P., Parsons, S. & Sklar, E. 2002. Co-evolutionary auction mechanism design: a preliminary report. In *Lecture Notes In Computer Science: Revised Papers from the Workshop on Agent-Mediated Electronic Commerce IV: Designing Mechanisms and Systems*. Padget, J. A., Shehory, O., Parkes, D. C., Sadeh, N. M. & Walsh, W. E. (eds). Springer 123–142.

Roth, A. E. 1991. Game theory as a part of empirical economics. *Economic Journal* **101**(401), 107–114.
Roth, A. E. 2000. Game theory as a tool for market design. In *Game Practice: Contributions from Applied Game Theory*, Patrone, F., Garcia-Jurado, I. & Tijs, S. (eds). Kluwer 7–18.
Roth, A. E. 2002. The economist as engineer: game theory, experimentation, and computation as tools for design economics. *Econometrica* **70**(4), 1341–1378.
Rubinstein, A. 1998. *Modeling Bounded Rationality*. M.I.T. Press.
Simon, H. A. 1982. *Models of Bounded Rationality*. M.I.T. Press.
Simon, H. A. 1996. *The Sciences of the Artificial*, 3rd edn. M.I.T. Press.
Streete, T. 1661. *Astronomia Carolina: A New Theorie of the Coelestial Motions*. Lodowick Lloyd.
Vriend, N. J. 2000. An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics & Control* **24**, 1–19.
Watson, J. D. & Crick, F. H. C. 1953. Molecular structure of nucleic acids: a structure for deoxyribose nucleic acid. *Nature*, no. 4357 (April 25) 737–738.
Wellman, M. P., Greenwald, A. & Stone, P. 2007. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. M.I.T. Press.
Wooldridge, M. & Jennings, N. R. 1995. Intelligent agents: theory and practice. *Knowledge Engineering Review* **10**, 115–152.