# Analysis and Synthesis: Multi-Agent Systems in the Social Sciences

Robert Marks [1]

[1]*School of Economics, University of New South Wales, Sydney, NSW 2052, Australia*
    E-mail: bobm@agsm.edu.au

**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 summarised 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 characterised.[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, 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 exhaustively that space, 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 generalisations, as well as parameters and initial conditions that must be specified. So model $A$ comprises the conjunction $(a_1 \wedge a_2 \wedge a_3 \cdots \wedge a_n)$, where $\wedge$ means "AND", and the $a_i$ denote the elements (equations, parameters, initial conditions, etc) that constitute the model.

**Sufficiency:**   If model $A$ exhibits the desired target behaviour $B$, then model $A$ is sufficient to obtain exhibited behaviour $B$. That is, $A \Rightarrow B$. 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 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 $A = (a_1 \wedge a_2 \wedge a_3 \cdots \wedge a_n)$, 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 non-linear 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, pers. comm.) argues, however, that necessity is only well defined in a class of

---

[1]This paper draws extensively on Marks (2006) and Marks (2007).
[2]Another is the influence of physics on early classical economics (Mirowski 2007).

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 $N$ of necessary models important? Without clear knowledge of the boundaries of the set of necessary models, it is difficult to generalise from simulations. Only when the set $N$ 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 & 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 "textita 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 emphasised in their 1953 paper. And any negative results would have meant returning to the drawing board, or in this case the brass rods and sheet metal. 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 obervations; necessity would have to await Newton's laws of gravitational attraction of 1686.

This discussion — the absence of necessity in simulation — must appear foreign to engineers and computer scientists. Problem-solvers would in general regard having several solution 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 difference between social scientists and computer scientists/engineers. Social scientists in general seek generality of understanding, while engineers in general seek solutions, specific instances, rather than generality. Mirowski (2007) argued 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 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 behavior 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 which 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 fifty-year 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 characterisation of equilibrium values (Mirkowski 2007) has dismayed those who need to model out-of-equilibrium interactions and the dynamics associated with real-world markets. Can we surmise that characterising asymptotic equilibrium parallels looking for the lost keys under the street-light — it's easier to do ("tractable") than the more realistic task of characterising out-of-equilibrium adjustment? But this characterisation is just what those designing markets — especially fully automated markets — need to have. The economists' understanding has proved disappointing.[3]

Analysis precedes design: in order to change systems, we must first understand them. Even so, direct design or optimisation requires a degree of understanding of the mapping from the design space to the performance space which has rarely been developed. Indeed, given the complexity of market phenomena, direct design might seldom be possible, as Edmonds & 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 stylised fashion) the observed real-world phenomenon. This gross simplification of the process overlooks the attempts to derive alternative explanations (theories) from the models, in oder to test them against the historical data.

The issue of program 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 historical data, and the ability to abstract from reality in building the (necessarily simplified) model of the real-world phenomenon. Then the

---

[3]Mirowski (2007) argued that for a hundred and fifty years economists have focussed 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 lacunæin the economics literature that confronted the computer scientists when they sought detailed analysis and explanation of the workings of historical markets in order to implement automated markets of various kinds.

skill to derive implications of the model that can be tested by examining the historical data to determine which possible explanation is 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 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) introduced the term "docking" when a second team attempts to replicate another team's simulation model. They clarified three decreasing levels of replication: "numerical identity," "distributional equivalence" (the results 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.[4] 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, 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. If the profit-maximising behaviour of the simulated brand managers, together with some external factors or internal factors, led to behaviour qualitatively or quantitatively similar to the behaviour of the brands — prices and quantities sold, 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 maximising, we are going beyond merely seeking a set of equations exhibiting periodicity similar to the "rivalrous dance" of the brands in the figure.

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) argued, that economic actors can behave as though they have solved complex optimisation problems, even though they remain ignorant of any formal representation of the problem or its solution.

### 3.2.1    Formalisation of Validation

Let set $\mathbf{P}$ be the possible range of observed outputs of the real-world system, here the weekly prices, quantities, and profits of the coffee brands sold in a single supermarket each week.[5]

Let set $\mathbf{M}$ be the exhibited outputs of the model in any week. Let set $\mathbf{S}$ be the specific, historical output of the real-world system in any week. Let set $\mathbf{Q}$ be the intersection, if any,

---

[4]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".

[5]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.

between the set $\mathbf{M}$ and the set $\mathbf{S}$, $\mathbf{Q} \equiv \mathbf{M} \cap \mathbf{S}$. We can characterise the model output in five cases.[6]

**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 doesn't 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 accurate. 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 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 variant of historical reality.

## 4    Economists' uses of Multi-Agent Systems

In economics, computer simulations have been increasingly popular in analysing 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 analysed using game theory, the best framework for analysing interactions where each player's outcome is a function of its 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 analysis 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 analysed from the bottom up by simulating the behaviour of interacting agents, buyers and sellers, a more realistic model, a priori.

---

[6]This conceptual framework was introduced by Mankin et al. (1977).

The GA was the first multi-agent system used to simulate and analyse oligopolistic behaviour (Marks 1989). The GA had been developed and pioneered by computer scientists and engineers who were intent on solving optimisation 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 modeling 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 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) which 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 modeling 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 coevolution, 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). Duffy (2006) surveys the growing literature on the links between laboratory experiments with human subjects and computer experiments with agents that learn, using models of learning originally borrowed from the psychological literature (Arthur 1991, 1993). See also Arifovic & Ledyard (2008).

## 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 seventy-year-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 realisation 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 privatisation 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 privatising 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 centralised 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 eight years, 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 revealed in some of the papers to be presented over the next few days.

Mirowski (2007) argued that the traditional reduced form of marekt-clearing, and the focus on agents rather than markets, meant that economiosts were 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 Duffy 2006 and Brenner 2006); and, third, computational exploration of different designs. Indeed, if the design criteria are clearly defined, some of the recent techniques of simulation and optimisation 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 behavior of the market, evolutionary computation raises the possibility of bottom-up design, or emergence of market design through simulation. As Mirowski (2007) argued, a typology of markets, characterised 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 behavior (where analysis is followed by design, given the behavior 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 behavior 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 would occur in the genome space, while the behavior 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 behavior of the economic actors (the performance of the system) is incompletely understood, and so that there are fewer surprises. Where the mapping is sufficiently well understood, and where closed-form analytic solution is 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 behavior will follow, at least in general form — but also necessity — if you want this performance and behavior, then this is the only set (or sets) of designs (combinations of structure and rules) that will produce it.

MacKie-Mason & 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 characterisation 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' behavior, 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 recognised, 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 optimisation, 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 & Bryson (2003) speak of the "syntactic complexity" of design. This is the lack of a clear mapping from design to behavior (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." 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 behavior 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 behavior 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 between the criteria. For instance, if a firm has market power, it can maximise 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.[7] Or it might be possible to improve the fairness of a market outcome, but at the cost of market 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 behavioral rules simply be evaluated directly? Third, market clearing and price formation need

---

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

to be modelled. Fourth, the fitness of the model must be evaluated. For example, should wealth or utility be used? And should the evolving behavioral 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 maximised 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 (optimisation) 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 scietists designing automated negoation 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 realise that the equilibrium characterisations 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 characterisation of the equilibrium may be of little help.

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, the design economist can 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 modeling 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 stability of trading to the introduction of new trading strategies; sensitivity to the number of agents trading; over-reliance on inductive models of agents, which respond to past rules and forecasts; and not enough on deductive models 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-fertilisation between computer science and the social sciences. Researchers from both sides have considered what attributes simulated agents should have (Gilbert & Troitzsch 2005). Wooldridge & Jennings (1995) would give computer agents these properties:

- autonomy: no other agents can control their actions and internal state,
- social ability: agents can interact and communicate with other agents
- reactive: agents perceive their environment and respond
- pro-active: agents initiate goal-directed actions
- (intentionality: metaphors of beliefs, decisions, motives,
- and even emotions)

While most agents used in economics have until recently been deductive (backwards-looking, stimulus-response automata), Epstein (1999) would desire these qualities in addition:

- heterogeneity: not "representative" agents, but may differ
- local interactions: in a defined space
- boundedly rational (Simon 1982): perceive information, possess memory, and some computational capacity
- ability to deal with non-equilibrium dynamics: large-scale transitions, tipping phenomena

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

## 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 now just a dozen years old, and the work by computer scientists on automated markets is even more recent. Only recently have there been attempts to use such models, after parameterisations of auctions, to directly design markets, including electricity markets. Indeed, direct market-design modeling 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.

## References

Arifovic, J., and J. Ledyard (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. (2006) "Advancing the Art of Simulation in the Social Sciences ", in J.-P. Rennard (ed.), *Handbook of Research on Nature-Inspired Computing for Economy and Management*, (Hershey, PA: Idea Group Inc.), pp. 90-100.

Axtell, R., Axelrod, R., Epstein, J., and 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 the *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science) pp. 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/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 the *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, ed. by L. Tesfatsion and K.L. Judd, (Amsterdam: Elsevier Science) pp. 949-1011.

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

Edmonds, B. and J. J. Bryson (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, *Handbook of Computational Economics*, *Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, Amsterdam: Elsevier Science, pp. 1585-1604.

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

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

Gode, D., S. Sunder (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*, New York: Springer, 2nd ed.

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

Jennings, N.R., P. Faratin, A.R. Lomuscio, S. Parsons, C. Sierra and M. Wooldridge (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, *Handbook of Computational Economics*, *Volume 2: Agent-Based Modeling*, ed. by Leigh Tesfatsion & Kenneth L. Judd, Amsterdam: Elsevier Science, pp. 881-893.

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

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

Koesrindartoto, D. and L. Tesfatsion (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, D.C., (IAEE/USAEE, Cleveland).

LeBaron, B. (2006), "Agent-based computational finance ", Ch. 24, *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science) pp. 1187-1233.

MacKie-Mason, J.K. and M.P. Wellman (2006), "Automated markets and trading agents ", Ch. 28, *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science) pp. 1381-1431.

Mankin, J.B., O'Neill, R.V., Shugart, H.H., and Rust, B.W. (1977), "The importance of validation in ecosystem analysis ", in G.S. Innis, ed. *New Directions in the Analysis of Ecological Systems, Part 1, Simulation Council Proceedings Series*, Simulation Councils, La Jolla, California, 5: 63-71. Reprinted

in H.H. Shugart and R.V. O'Neill, eds. *Systems Ecology*, Dowden, Hutchinson and Ross, Stroudsburg, Pennsylvania, 1979, pp. 309-317.

Marks, R.E. (1989) "Breeding optimal strategies: optimal behavior for oligopolists ", *Proc. of the Third International Conference on Genetic Algorithms*, George Mason University, June 4-7, 1989, ed.& by J.D̂. Schaffer, San Mateo, Calif.: Morgan Kaufmann Publishers, pp.198-207.

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

Marks, R.E. (2006), "Market Design Using Agent-Based Models ", Ch. 27 in the *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science) pp. 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., Midgley D.F., Cooper L.G. (2006), "Co-evolving better strategies in oligopolistic price wars ", in Rennard (ed.), *Handbook of Research on Nature-Inspired Computing for Economy and Management*, (Hershey, PA: Idea Group Inc., 2006) pp. 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, New York).

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

Midgley, D.F., Marks R.E., and 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, Cambridge).

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

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

Roth, A.E. (1991), "Game theory as a part of empirical economics ", *Economic Journal*, 101(401): 107-14.

Roth, A.E. (2000), "Game theory as a tool for market design ", in F. Patrone, I. Garcia-Jurado, and S. Tijs, eds., *Game Practice: Contributions from Applied Game Theory*, (Kluwer, Dordrecht), pp. 7-18.

Roth, A.E. (2002), "The economist as engineer: game theory, experimentation, and computation as tools for desgn 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*, M.I.T. Press, 3rd. ed.

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., & F.H.C. Crick (1953) "Molecular structure of nucleic acids: A structure for deoxyribose nucleic acid ", *Nature*, no. 4357, pp. 738-8, April 25.

Wooldridge, M., & N. R. Jennings (1995), "Intelligent Agents: Theory and Practice ", *Knowledge Engineering Review*, 10: 115-152.