Market Design

Robert Marks
Australian Graduate School of Management
bobm@agsm.edu.au

Draft of a chapter in the *Handbook of Computational Economics, Volume 2, Agent-Based Modeling*, edited by Ken Judd and Leigh Tesfatsion, North-Holland, 2005, forthcoming.

1. Introduction[1]

Institutional arrangements for exchange — markets — have emerged and evolved over the millennia since — and perhaps as a consequence of — specialization of labor, which can be intensive (making something "better" than others do, absolutely or relatively) or extensive (taking the risk of fetching an item, not locally available, from afar). "Trade" first meant exchange of foreign-produced goods for domestic goods, a form of barter, which is made more efficient with the emergence of money — numeraire, store of wealth, and medium of exchange, in the textbooks' trio.

Many different market institutions have evolved, well described in John McMillan's recent book, *Reinventing the Bazaar* (2002). The development of economics, in one view, has been the outcome of reflecting on, describing, and analyzing various markets, from the market-town's weekly bazaar to the complex financial markets for exchanging risk. One form of market institution is the auction, and only over the past forty-odd years, with the development of the tools of game theory, has formal analysis of auctions begun.

1.1 Designer Markets

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:

1. 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.
2. 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,

---

1. I acknowledge help in writing this chapter from the editors, Raimo Hämäläinen, Derek Bunn, Peter McBurney, Bob Wilson, Paul Klemperer, Eddie Anderson, Carol McCormack, and my fellow contributors at the Handbook Workshop at the University of Michigan, May, 2004.

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.

3. 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 hick-ups. 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).

4. 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.[2]

5. 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 discussed below.

This chapter will focus on the use of agent-based models in designing the fourth type of market, that for electricity. The first, for emissions abatement, is covered by Janssen and Ostrom. The second is covered by the chapters by Blake LeBaron and Cars Hommes. The fifth, for on-line auctions is covered by Jeff Mackie-Mason and Mike

---

2. Despite the debacle of the California blackouts of 2000, it is increasingly clear that it was not the underlying market design per se at fault, rather it was its implementation and the consequences of lobbying by vested interests: the retail price was regulated, while the unregulated wholesale price sky-rocketed as a consequence of market manipulation, which had the effect of squeezing the retail electricity companies, such as Pacific Gas & Electricity.

Wellman.

The use of game theoretic methods to analyze market design is related to the use of these techniques to analyze another kind of interaction, those governed by contracts. Contract design is another area where agent-based modeling can be used, but analysis of contracts by use of computer simulation and agent-based modeling is in its infancy.

2. Analysis, Design, and Simulation

Before design must come analysis. Simulation allows analysis of systems that are too complex to analyze using traditional, closed-form techniques. Once we understand through analysis how the elements of the phenomenon of concern work together, we can ask the question of how to improve its operation: how better to design it.

2.1 Analysis

In the world of analytical, closed-form solutions, there is a certain logic to the progress of research. A 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, traditionally concerned with existence, uniqueness, and stability of an equilibrium, 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. Successful analyses are published, indexed, and referenced.

A common understanding of this process in general, but particularly the process of model-building and deducing the system's behavior and results, means that, by and large, the research effort builds by aggregation of later research to earlier research: later researchers 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.

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 car keys under the street-light, instead of in the darkness around his car, it might not be coincidental.)

2.2 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 — 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 the brain, and Genetic Algorithms (GA) were inspired by natural selection with sexual reproduction). Over thirty year ago it was possible for Donald Knuth to write an encyclopædic 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.

The focus in this and following sections will be on analysis, rather than design. This is not because we have overlooked design; it is because, as we discuss below, fully blown design requires a degree of understanding of the mapping from the design space to the performance space which has not yet been developed. Indeed, given the complexity of market phenomena, such straightforward design might never be possible, as Edmonds & Bryson (2003) remind us.

2.3 Evolutionary Simulation Techniques

The development of agent-based methods owes its existence to the almost simultaneous emergence in Germany and the U.S.A. of simulation techniques that mimic aspects of natural selection. Holland's Genetic Algorithm (GA) (1976, 1992) was used as a new kind of optimizing tool (for problems intractable to calculus-based tools), in which a population of possible solutions was tested, individual by individual, and then, based on the "fitness" score of each, selection of pairs of "parents" for a new generation of individuals was made and, based on the processes of "cross-over" and "mutation" (analogous to mimicry of existing solutions and to exploration of new regimes of the solution space) the "offspring" generation of possible solutions was derived. Testing, selection, and generation of a new population results in the emergence of never worse best solutions. The GA technique has been widely used as an optimizer, a directed form of trial and error that obviates exhaustive testing of all possibilities.

But used as an optimizer in this way — focusing on the best solution (an individual) — throws away the population's emerged characteristics qua population. A line of research that began with Axelrod's (1987) simulation of individuals playing the repeated Prisoner's Dilemma (RPD) used the population of individuals — stimulus–response automata, where the stimulus was the state of the interaction, and the response was the next action of the player — to consider not only the emergence of new strategic automata, but also to examine the stability of the population against "invasion" by a new strategy.

Axelrod, a political scientist, was interested in combinations of strategies that exhibited the emergence of cooperation, a manifestation of the Folk Theorem of RPD. But since the RPD can be thought of as a simple model of a Bertrand duopoly, his work soon gained the attention of economists, who had found the analytical characterizations of equilibria in oligopolistic competition incomplete, not least in the paucity of out-of-equilibrium characterizations of the dynamics of the interaction. That is, the intermediate behavior of a dynamic interaction, a game, might be more important than its asymptotic properties.[3]

When the players face identical payoff sets and choose from identical action sets,

---

3. Just how to characterize out-of-equilibrium behavior (or bounded rationality, for that matter) remains an open question.

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 on the same end: faced with the same state of the interaction, either of the players would behave identically, and fitness is average (or discounted) profit.

But when modeling 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" (Kauffman 1995) that is not only possibly rugged, but also shifting (as each agent wears a distinct sellers hat, as it were). In this case, separate populations of sellers makes sense.

A further issue is modeling the process. The GA mimics the process of sexual reproduction. It can be thought of as mimicry and exploration, but there has been concern that this is not how people learn, and perhaps — although this is not unequivocal — how organizations learn, if indeed organizational learning is different from individual learning. Agent modeling has expanded from the processes of the GA, with its implicit "learning," to agents whose learning is explicit, as discussed in John Duffy's chapter.

The GA was developed and pioneered by computer scientists and engineers who were intent on solving optimization problems exhibiting rugged landscapes. 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 modeled by the GA are competing against each other, the GA is modeling the process of co-evolution. This process was mistakenly called boot-strapping by Marks (1989), in the first published research into co-evolution of rivals' strategies in oligopolies.

2.4 Design

As Roth remarked in an early paper on Market Design (1991), market design is a suitable case for using three approaches: first, traditional closed-form game-theoretic analysis, as discussed above; second, experimental results from economics laboratories; 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.

Market performance may depend on the degree of "intelligence" or "rationality" of the agents buying and selling, which has led to computer experiments in which trading occurs between artificial agents of limited or bounded rationality. As Walia et al. (2002) remark, if a market design with agents of low degree of "intelligence" is found to be sufficient for a specific level of market performance, then we can have reasonable confidence that agents with a higher level of intelligence, or agents whose rationality is less bounded, will, through their decisions to buy and sell, inadvertently create for themselves a market that is working efficiently.

Of course, historical market institutions have in general not been imposed from above (so-called top-down design) but have emerged from the bottom up as a consequence of a multitude of actions and interactions of the myriad traders. 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 computational devices raise the possibility of bottom-up design, or emergence of market design through simulation.

This 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 — 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.

These alternatives raise two issues (Tesfatsion 2002, p.19): first, to what extent are the learning processes of human participants in real-world markets mal-adapted to market institutions? Perhaps the use of agent-based optimization tools could improve human market behavior, as is already seen, for instance, in eBay auctions, when bidders use software to enhance their chances of being the high bidder at the deadline.

Second: to what extent have existing market protocols (or market designs) evolved or been designed to avoid the need for any great rationality on the part of market participants? Gode & Sunder (1993) and others seek to answer this question for financial markets, but their results may, under certain conditions, be valid for other markets. These issues are explored at greater length in the chapters by Blake LeBaron and Cars Hommes.

When there are several criteria by which the desirability of a designer market might be judged, trade-offs are necessary, and in the case of the GA, which needs one measure of each agent's fitness, such trade-offs must be explicit beforehand.

3. Market Design

Design is a process of building, where the ends are specified first, so the process of building is directed by the design objectives, if not by an explicit plan. Unfortunately, specifying objectives does not always immediately delineate exactly how the 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 not well 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 of designs (combinations of structure and rules) that will produce it.

Without a closed-form analytical solution, but instead with human experiments or with computer simulations, necessity is in general out of reach, and we must make do with sufficiency. (Note that this is not always the case: James Watson and Francis Crick (1953) used a form of simulation to determine the structure of DNA, with their metal rods and brass atoms, but the experimental results from the work of others had so constrained the degrees of freedom in the space of possible structures that they knew when they'd solved the structure: see the Appendix.)

3.1 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: the only way to know the system's outcomes is to run the system, and analysis is 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 (1981) 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."

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.*

3.2 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 maximize its seller revenue, but at the cost of market efficiency, as measured by the sum of sellers' (or producers') surplus and buyers' (or consumers') surplus. Or it might be possible to improve the fairness of a market outcome, but at the cost of market efficiency. As we shall see below, to use computer simulation, such trade-offs must be explicit. It might be possible to use a version of Simon's (1981) satisficing, whereby so long as the other criteria are met (above some target level), the remaining criterion is used to rank designs.

As Phelps et al. (2002) point out, possible criteria for judging how good a designed auction market is might include: maximizing seller revenue, maximizing market efficiency, discouraging collusion, discouraging predatory behavior, discouraging entry-deterring behavior, or other criteria. Wilson (2002) and Cramton (2003) discuss issues of electricity market design.

3.3 Moving from Closed-Form Equilibria

Traditionally for the last sixty years, economists have sought closed-form solutions to understanding the performance of economic institutions. Economic actors have been assumed to be perfectly rational, with the means to solve for equilibria outcomes in complex situations. Economists have sought to characterize the equilibria of economic interactions in terms of their existence, uniqueness, and stability, under this assumption. When the interactions among economic actors are strategic, the equilibria become Nash equilibria.

But in an operating, real-time actual market, it turns out that we are not interested just in equilibrium characterization: continual shocks might never allow the system to approach, let alone reach, the equilibrium. And, moreover, it turns out in a repeated interaction that almost any individually rational outcome for each player can be supported as an equilibrium. (The Folk Theorem of repeated games.) This is particularly so for interactions which have the general character of the repeated Prisoner's Dilemma (RPD).

Consequently, there are at least two reasons why market design has moved away from traditional closed-form solutions: first, because of tractability: it has been very difficult, despite advances made in recent years, to obtain solutions to the design of some markets, such as continuous double auctions. Second, we should like to characterize out-of-equilibrium behavior, and especially the dynamic behavior of an operating market with fluctuating demand, and perhaps varying numbers of sellers, with unpredictable, varying costs.

A third reason for considering other techniques of analysis is that the assumption of perfect rationality and unlimited computational ability on the part of human traders is unrealistic, and not borne out by laboratory experiments with human subjects. Instead, using computer models of trading agents, we should like to model economic actors in markets as "boundedly rational." This might mean bounded computational ability, or bounded memory, or bounded perception (Marks 1998).

3.4 Learning

There is a fourth reason for wanting to move from closed-form solutions, even where they are available: to model learning. There are two reasons to include learning in any models used to design markets: first, individuals and organizations learn. Human players learn (perhaps with the added incentive of the prospect of bankruptcy if they do not learn from their mistakes), which means that a model without learning is not as realistic as one incorporating learning. Bunn & Oliveira (2003) note that many researchers (including Roth & Erev, 1998) have shown that learning models predict better than do the Nash equilibrium how people behave.

Second, learning can help to eliminate many otherwise legitimate Nash equilibria from further contention. Indeed, evolutionary (or learning) game theory has been seen as a solution to the multiplicity of Nash equilibria that occur in closed-form game-theoretic solutions: a priori, all are possible, but to see which are likely in reality, see how players learn and choose amongst them.

There are four types of models of reinforcement learning, in which how an actor chooses to behave later is a function of the outcomes he has experienced earlier, in part as a consequence of his earlier choices (the Thorndike effect, Thorndike 1911). They are the Roth-Erev model (Roth & Erev, 1999), *Q*-learning, which optimizes long-term payoffs, rather than immediate (Watkins & Dayan, 1992), multi-agent *q*-learning (Hu & Wellman 1998), and Adaptive Play (Young, 1998). John Duffy's chapter discusses these models of learning and others. We discuss some papers which use these below, and discuss two kinds of learning in the context of simulations using artificial agents.

3.5 Explicit Use of Agents

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 RPD and oligopolies with asymmetric players, where the fitness of an agent depends on the state of the whole population of agents: state-dependent fitness (Riechmann 2001). Sargent (1993) surveys studies using adaptive algorithms (including the GA) to model macro-economic phenomena with learning agents, but not explicitly agent-based models.

The use of GAs has been criticized on two levels: Chattoe (1998) is critical of evolutionary programming in general, and of the GA in particular, as models of social evolution; this, however, is not the purpose of GAs used to analyze markets. Dawid (1999) has argued that, despite its foundation in computer science, the GA is good at modeling the ways in which populations of economic actors can learn. Vriend (2000) draws the distinction between the social learning of the GA (whereby the individuals in the population have learnt 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. There is some evidence that the number of GA-based models is falling and the number of agent-based models is growing (from the IEEE’s Xposure on-line index of the Congress on Evolutionary Computation annual proceedings), although terminology shifts, which might explain any perceived shifts in usage.

Design of markets might occur with simultaneous “design” of trading agents, a line of research pursued with GA learning at Hewlett-Packard by Cliff (2001a, 2001b, 2002, 2003a) on continuous double auctions and by Byde (2002) on sealed-bid auctions.[4] This might be of interest when the designer market will not be a venue for human traders (or their organizations), but rather will be a venue for the designer trading agents (the “buy-bots” and “sell-bots”). This situation has become a possibility with the growth of the Internet. The use of artificial trading agents in business-to-business wholesale trading and in allocations internal to the company or organization is where one might expect such agents to appear most naturally.

Byde (2002) at Hewlett-Packard described an evolution-based method for evaluating sealed-bid auction mechanisms. He applied it to a space of mechanisms including the standard first- and second-price sealed-bid auctions, and was able, first, to replicate results known already in the Auction Theory literature regarding the suitability of different mechanisms for different bidder environments. Using standard GA learning, he then found that “under several classes of non-pathological conditions (e.g. bidders

4. On his web page, Dave Cliff (2003b) explains how he came to develop computer traders — his ZIP (Zero Intelligence Plus) traders — that researchers at IBM found outperformed human traders (Das et al. 2001). “The wonderful results in the IBM paper, and the success of using the GA to get better ZIPs, led me to think about using a GA to design new marketplaces that are specialized for trading agents.” (See Cliff 2001b, et seq.)

were risk-averse, and unaware of how many players they would face in a given auction), there existed exotic sealed-bid mechanisms expected to return significantly higher revenue to the auctioneer than either the first- or second-price sealed-bid mechanisms," specifically a (0.3, 0.7)-price auction. (See the chapter by Mike Wellman and Jeff Mackie-Mason.)

3.6 The Design Economist

So the software programmers are *designing* systems of exchange, of markets. Recently, 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.

The shortcomings of the 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 learnt 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. Simulation has occurred using GAs, numerical solutions, and explicit agent-based models. Iterative analysis has been used as a means of designing systems.

4. Electricity Market Design

In 1998 the U.S. Federal Energy Regulatory Commission Chairman, James Hoecker (1998), said: "Arguably, a well-constructed computer model could improve the accuracy of our competitive analysis in at least two ways: by explicitly representing economic interactions between suppliers and loads at various locations on the transmission network, and by accounting for the actual transmission flows that result from power transactions." He warned, however, that: "Consistency of data sources and consistent application of those data is an attraction, but such techniques require time, education, and consistent refinement. Moreover, adequate data may not be available. I hope the benefits will be worth our trouble and investment. Our economists are trying to get a handle on precisely that equation."

Other economists and engineers and computer scientists had already been at work on this issue for some years, when Mr Hoecker spoke. Applications of agent-based modeling to electricity market analysis and design occurred independently in several research centers. The application of genetic algorithms to, first, oligopolies, and then to macro-economic models, has more recently been followed by its use in analyzing the behavior of new markets for electricity generation and transmission, most recently as a means of designing electricity markets.

4.1 Academic Engineers

In 1992, a pioneering paper by Verkama et al. (1992) at the Helsinki University of Technology argued that the two disparate areas of oligopoly theory and distributed artificial intelligence (DAI) could learn from each other, since each was concerned with modeling the interaction of autonomous, self-interested, interacting agents. Using object-oriented programming, they had developed a test-bed for examining agents' interactions under various initial conditions. They acknowledged that "very general results are difficult to come by with simulations and computer experiments" (p. 157), but argued that such approaches allow the exploration of market evolution, with entry and exit, learning, and reputation effects. They even suggested that the market itself could be modeled as an agent, the first suggestion in the literature that the design of markets could be modeled and analyzed, necessary antecedents for market design using agents.

Verkama et al. (1992) do not cite any works in evolutionary computation, but two years later, after presentation at a workshop in computational organization theory (Hämäläinen 1992), they (Verkama et al. 1994) cited Arthur (1991, 1993), Holland and Miller (1991), and Lane (1993a, 1993b). The linkages between two previously independent lines of research had been made.[5] In the 1994 paper, as well as object-oriented programming, they mention inter alia genetic algorithms and learning automata, and the need for agents to mimic human behavior in simulation models of strategic interaction (their "reactive behavior"). The test-bed itself had evolved: in their Multi-Agent Reactions Testbed agents can inherit properties from previous generations and add new features, in order to explore the interactions of different decision rules, and the market structure and rules of engagement.

In 1994 Räsänen et al. (1994) introduced an object-oriented model of electricity demand-side load, the first application of such techniques to electricity market modeling, although the use of inherited characteristics was not to allow the objects to evolve or learn, but rather to aid the programmer in modeling changed load. A year later, however, Hämäläinen and Parantainen (1995) introduced a new "agent-based modeling framework" for analyzing electricity markets by using agents to model the demand-side load, and two years later Hämäläinen et al. (1997) went much further, with agents representing both sides of the electricity market — consumers and producers — with

---

5. In a private communication Raimo Hämäläinen (2004) explains: "The origins of my interest go very far back. We had been working on game theory, coordination and resource economics, and to me as an engineer it was a natural idea to see what could be achieved by a computational analysis of economic systems. One of the first computational analyses was [a 1978] paper on the role of information in decentralized macro-economic stabilization. Later, coordination ideas grew in my head when I was working on fishery models [in 1986 and 1990]. This was followed by incentive and coordination work: Verkama et al. (1992). At the time of the emergence of our interest in energy economics the Finnish market had not yet been deregulated, but this took place during our research project on real-time pricing of electricity. For a period this kind of research was not considered interesting as markets were the hot topic."

bounded reasoning capabilities and bounded reactions. They have not adopted GAs or other computer science techniques referred to in the 1994 paper. This has been left to others.

Meanwhile, at Carnegie Mellon University, Talukdar & Ramesh (1992) suggested modeling the software to manage electricity generation when the operating environment (market) could change rapidly with an asynchronous “A-team,” one of the first examples of a multi-agent system in the electricity literature. Krishna & Ramesh (1998) extend the idea to developing “intelligent software agents” to help generators to negotiate with potential coalition partners; they point to the possibility of such agents replacing human players in computerized electricity exchanges. Talukdar (2002) continues to use artificial agents as members of his asynchronous teams, sometimes borrowing from the GA models, most recently to simulate and verify the trades that occur in repeated markets, such as electricity markets, as part of the market design process.

Meanwhile, at Iowa State University, a group of electrical engineers led by Sheblé had started in 1994 to examine the operation and design of electricity markets. Maifeld & Sheblé (1996) use a GA for solving the unit-commitment scheduling problem in electricity markets. They referred to no earlier work by economists, but Richter & Sheblé (1998) referred to unpublished work by LeBaron and by Tesfatsion, and used a GA to learn (evolve) bidding strategies in an electricity market as generators and distributors buy and sell power via double auctions; amongst other things things this model can be used to explore how bidding behavior affects overall market performance. Richter et al. (1999) extended their previous work on bidding strategies in double auctions for trading electricity competitively. They used adaptive automaton strategies: tested in an auction simulator, the automata learn using a GA. The paper examined high-profit strategies and also modeled certain types of trading behaviors. Lane et al. (2000) broadened the scope of the research: they modeled the traders in an electricity market as adaptive agents learning with the help of a GA in a double auction, and, perhaps influenced by Tesfatsion’s economics research, calculated the degrees of market power for various combinations of relative capacity and production costs.

With the increased use of markets to help allocate the generation and distribution of electricity in several countries, this concern with using models of electricity markets to examine the exercise of market power is an obvious extension of the simulations, and reflects the shift from analysis of the traders’ actions to analysis of the markets’ performance, a necessary step for market design.

Meanwhile, engineers at the University of New South Wales (MacGill & Kaye 1999) were exploring a decentralized coordination framework to maximize the efficiency of the power-system operation, not through the operation of Smith’s invisible hand as each resource competes to maximize its own return, but with a decentralized framework in which each resource is operated to achieve overall system objectives. The authors believed that evolutionary algorithms were not well suited to their problem, and instead developed a so-called “dual evolutionary approach,” which uses a version of the GA, but not explicitly with autonomous, self-interested agents. Cau & Anderson (2002) used GAs to examine co-evolutionary behavior of agents in markets for electricity, where such agents were modeled as autonomous, self-interested players (see also Cau, 2003). In particular they were interested in exploring the conditions of the players and of the market under which tacit collusion occurred. Since collusion leads to inefficiencies, from a policy-maker's viewpoint a market structure which discourages the emergence of learned tacit collusion is a good design. Engineers, mathematicians, and economists in Sydney (Outhred, MacGill, Anderson, and Marks) continue to use evolutionary, agent-based modeling to explore electricity spot markets (MacGill, 2004).

The number of engineering studies of electricity supply and distribution networks that employ agent-based (or "multi-agent") simulations of some sort or other continues to grow, as reflected in published papers in the IEEE journals, transactions, and proceedings.

4.2 Economists

In 1997 an economist at University College London, Curzon Price (1997), presented simulation models of simple electricity pools, in which he used the GA as a means of simulating the repetition of two rival sellers. He saw competition in electricity markets, often across jurisdictional borders, as a field in which the "underlying economic models are often quite simple," but the real-world phenomena "complicated and richly detailed in important ways" (p. 220), and hence suitable for simulation. This was a clear progression from the work that Marks (1992) and others had undertaken to use the GA to examine oligopolistic behavior, following Axelrod's (1987) work using the GA to examine strategies in the repeated Prisoner's Dilemma. Curzon Price suggested that plausible behavioral elements could be included in the simulations.

Curzon Price's work was directly descended from Axelrod's (1987) work with GAs and RPDs, Marks' (1992) work on oligopolistic behavior, and other economists' use of GAs, such as Andreoni & Miller's (1995) exploration of auctions using the GA to model the co-evolution of artificial adaptive agents. They found that their model of adaptive learning was consistent with the main results from laboratory experiments, and that — significantly for us — various auction designs ("institutions") display very different adaptive dynamics.

Iowa State University has been a fertile place for cross-disciplinary research in agent-based modeling of electricity markets. As well as Sheblé in engineering, it is home to Tesfatsion in economics. Two of the most widely cited papers on the application have emerged from her research group.

Nicolaisen et al. (2000) used a GA agent-based model of a double auction electricity market to examine the exercise of market power (as deviations from competitive equilibrium values of prices and quantities). They found no evidence that the market power of buyers is negatively related to their relative capacity or that the market power of buyers is positively related to their relative capacity. But in this model, traders were quite boundedly rational, with little memory. Moreover, the social learning of the GA learning process meant that any comparative advantage in strategies soon spread to the rest of the population of players and became dissipated, as Vriend (2000) discussed. The paper cites earlier work by Lane and by Richter, both at Iowa State.

In an attempt to obtain results on market power that were closer to those from standard theory, Nicolaisen et al. (2001) used reinforcement learning (a modification of Roth-Erev, 1998) instead of GA learning to allow individual learning and to prevent any comparative advantage in strategies being dissipated among the artificial agents. Otherwise the paper's model was similar to the earlier work (Nicolaisen et al. 2000): a double auction with discriminatory pricing. A high market efficiency was generally obtained, and the relative market powers of buyers and sellers were clearly related to the micro-structure of the market, independent of the learning parameters. In asking whether the market design ensured efficient, fair, and orderly market outcomes over time despite repeated attempts by traders to game the design for their own personal advantage, Nicolaisen et al. were clearly focused on market design. The paper cited Bower & Bunn (2001) and Lane et al. (2000).

One of the most successful academic economists to use agent-based techniques to analyze electricity markets is Bunn with his associates at the London Business School. As well as publishing in the economics literature, he has also published in the energy and regulatory literature, and his models have been calibrated against historical data. In Bunn & Oliveira (2001), we read: "The development of a detailed simulation platform representing the agents, the markets, and the market-clearing mechanisms, together with reinforcement learning to facilitate profit-seeking behavior by the agents, can, in principle, provide a computational framework to overcome the limitations of the analytical approaches." That is, such a platform could be used to design a market.[6]

Following the deregulation and privatization of the electricity generation sector in Britain, Bunn & Day (1998) proposed using agent-based simulation of electricity power pools to analyze the short- and longer-term behavior of the generators, as they learnt, partly to see whether high prices might be the result of implicit collusion.

Bower & Bunn (2000, 2001) developed a simulation model of the wholesale electricity market in England and Wales as a means of systematically testing the potential impact of alternative trading arrangements on market prices, specifically uniform- versus discriminatory-price auctions, thus undertaking a form of market design. Generators were represented as autonomous, adaptive, computer-generated agents, which progressively learned better profit-maximizing bidding behavior, by developing their own trading strategies, in order to explore and exploit the capacity and technical constraints of plant, market demand, and different market-clearing and settlement arrangements. Their agents used simple internal decision rules that allowed them to discover and learn strategic solutions which satisfied their profit and market-share objectives over time. These rules constituted what is essentially a naïve reinforcement-learning algorithm, and the behavior of the simulated market is thus almost entirely emergent. The agents knew everything about their own portfolio of plants, bids, output levels, and profits, but nothing about other agents or the state of the market. Their ability to capture and retain data was limited, they had no powers of strategic reasoning, and hence they exhibited a high degree of bounded rationality. The agents were modeled as data arrays in Excel 97 and manipulated with Visual Basic. They concluded that the discriminatory auction results in higher market prices than does the uniform-price auction. The papers did not cite any earlier work on agent-based modeling.

---

6. In a private communication, Derek Bunn (2004) remembered that his interest in using agent-based models followed from a new Ph.D. candidate with a computer science background who suggested using Object-Oriented Programming (Gamma et al. 1995), such as Java, as a better platform for simulating the electricity market than Systems Dynamics (Forrester 1961). As we see below, OOP leads to agent-based models relatively easily.

This research did not capture the interaction between the bilateral trading and the balancing market, nor did it incorporate any sophistication in the agents' learning abilities. Bunn & Oliveira (2001), however, describe a model with agents whose learning was inspired by the fitness function and selection mechanisms used in GAs. They argue that, by keeping the probabilities of exploration and exploitation independent of the expected reward from following a particular bidding strategy, their GA model should be trapped at local equilibria less often than would agents using a naïve reinforcement-learning algorithm, such as Roth-Erev (1998), especially in non-stationary environments.. Their new simulation platform was a much more detailed representation: it actively modeled the demand side and the interactions between two different markets, as well as the settlement process; and it took into accounts the daily dynamic constraints and different marginal costs for each generation technology. It referenced two earlier works from the GA simulation literature: LeBaron et al. (1999) and Nicolaisen et al. (2000).

Bower et al. (2001) applied a similar agent-based model to the German electricity market, specifically examining the effects on peak prices of consolidation, and the potential for the exercise of market power by the dominant generators. The references in this paper include Hämäläinen (1996) and Curzon Price (1997).

Bunn & Oliveira (2003) use agent-based simulation in a coordination game to analyze the possibility of market power abuse in a competitive electricity market. The model builds on the work in Bunn & Oliveira (2001), but does not allow the agents to learn as they did in the earlier, GA-based model, in order to retain more transparency in understanding their actions. Instead, the model uses reinforcement learning. The aims of the authors was not to evaluate the market structure but rather to see whether market conditions were sufficient to allow the exercise of market power by a certain player. The paper referenced Nicolaisen et al. (2001).

4.3 Computer Scientists

As mentioned above, inspired by natural phenomena, computer scientists invented various forms of evolutionary programs, such as as Holland's GA. They had for some time also been interested in DAI and object-oriented programs, which allow parallel processing to speed solution of the simulation models. This use of multi-agent systems resulted in a special issue of the *Journal of Artificial Intelligence*, edited by Boutilier et al. (1997), on the Economic principles of multi-agent systems, which attempted to introduce computer scientists to the work of economists and game theorists in modeling the interactions of few and many economic actors in markets.

Note that, as they design computerized trading systems, computer scientists have also become interested in the means by which explicit communication between agents might facilitate the operation of these virtual markets. Economists analyzing oligopolistic markets and auctions using agent-based models have denied their agents the possibility of explicit communication: under the various antitrust regimes such communication would be illegal. Instead, any communication must be arm's-length signaling by means of prices chosen in previous rounds, if common knowledge.

As well as developing algorithms to pursue simulations of market interactions, computer scientists have also been pioneers in the task of parameterizing auction design space (Wurman et al. 2001). This achieves two things: it allows a standard way to describe auction rules, for human beings or for software agents; and, more importantly for us, parameterization of the design space of auctions is necessary to allow agent-based design of markets in general and auctions in particular to proceed. A further motivation is to aid the development of auctioneer programs, perhaps on-line.

At IBM, Walsh et al. (2002) used replicator dynamics (Weibull, 1995) to model learning in a multi-agent system to analyze the dynamics and equilibria of two market types for which a full game-theoretic analysis is intractable: automated dynamic pricing, where sellers compete, and automated bidding in the continuous double auction. Unlike GA learning, replicator dynamics cannot generate new strategies or rules: it can only alter the likelihoods of strategies and rules existing at the start of the simulation. The authors are explicit about the need to obtain clear understanding of the workings of such mechanisms through analysis before design is possible: efficiency and stability are two design criteria mentioned.

At the University of Liverpool, Phelps et al. (2002), in a study of mechanism (market) design, sought to co-evolve the buyers, the sellers, and the auctioneer. That is, they viewed the market as the outcome of some evolutionary process involving these three types of actors. They took Nicolaisen et al. (2001) and identified two possible techniques for computer-aided auction design based on evolutionary computing: genetic programming (GP) and the modified Roth-Erev (MRE) reinforced learning algorithm (Roth & Erev, 1998). They were able to replicate Nicolaisen et al.'s 2001 results (although with greater variance); they then compared Nicolaisen et al.'s MRE learning with a model that used GP co-evolutionary learning (Koza, 1993) — they found that efficiencies with the GP model were reasonably high. Finally, they used GP to evolve auction rules (designs), and obtained relatively high-efficiency outcomes. Unfortunately, the auction rules so evolved were anything but transparent to human eyes: twenty lines of Lisp code. They cited Curzon Price (1997). The team at Liverpool continues to these issues, as part of a project into market-based control of complex computational systems.

4.4 Recent Non-Academic Research Centers

It is the mark of a successful research method that its use has spread beyond the academy into government agencies (as foreshadowed eight years ago by the head of the FERC) and commercial research organizations and companies. The agent-based analysis and design of electricity markets is a successful research method. We briefly mention the latest centers of research into electricity market design using agent-based models: the Electric Power Research Institute (EPRI) and the Lawrence Berkeley National Laboratory, Argonne National Laboratory, and Hewlett-Packard. (Tesfatsion (2003) discusses other centers.)

The Argonne National Laboratory has developed the Electricity Markets Complex Adaptive Systems (EMCAS) model, which incorporates agent learning and adaptation based on performance and changing conditions (North et al. 2001, 2002). There are user-specified market rules affecting the behavior of individual agents as well as the system. Earlier work at Argonne (North 2000) was based on the SWARM agent-based modeling platform (Burkhart et al., n.d.). EMCAS, however, is based on the RePast open-source agent-based simulation platform (Collier & Sallach, 2001) and uses GA learning for certain agents. EMCAS is designed to determine the state or states to which the market will gravitate, and the transients involved in getting there. Customer agents represent electricity users and company agents represent electricity suppliers. In EMCAS, each company agent seeks to maximize its individual corporate utility, not overall social utility, as it interacts with other agents and with the Independent System Operator (ISO) or Regional Transmission Organization (RTO) agent. EMCAS operates at six interdependent time scales: from real-time dispatch, to planning day-ahead, week-ahead, month-ahead, year-ahead, and in the medium-to-long term (2−10 years). The authors are aware that as well as allowing alternative company strategies to be simulated, EMCAS allows market rules to be tested: iterative market design.

Meanwhile, Harp et al. (2000) developed a proof-of-concept software tool, SEPIA (for simulator for electric power industry agents), an agent-based simulation platform for modeling and exploring a complex adaptive system, the electric power industry. It used two kinds of learning algorithms: *Q*-learning (Watkins & Dayan, 1992), a version of reinforcement learning, and genetic classifier systems. SEPIA was hosted at Honeywell, and was under-written by EPRI.

EPRI has used agent-based models to explore market design: Entriken & Wan (2003) describe experiments using computer-based agents to simulate the impact of the California Independent System Operator's proposed Automatic Mitigation Procedure (AMP) on market behavior. These computer agents play the role of market participants seeking to maximize their profits as they formulate bids under a number of scenarios over a simple, two-node market at various levels of demand and transfer capability and with and without the AMP in force. The study demonstrates that agent-based simulation is a useful tool for analyzing existing and proposed design features of electricity markets. One aim was to eliminate the need for human laboratory subjects, and they configured the computer agents in an attempt to eliminate experimental bias. The researchers modeled demand players as price takers: they always bidded their willingness-to-pay. Suppliers used an identical strategy of aggressive profit maximization. By comparing their bid prices with the market-clearing price, suppliers could determine whether they were marginal, in which case they used a very simple naïve rule for rent capture: they tested the margin by raising their bid prices. Agents were given the opportunity to learn, although the exact learning algorithm is not described.

5. Conclusion

Design of markets, mechanism design, using the tool of agent-based simulation is in its infancy. The iterative analysis of electricity markets with agent-based models is now just ten years old, and only recently have there been attempts to use such models, after parameterizations of auctions, to directly design markets, including electricity markets, as we have seen. Indeed, direct market-design modeling attempts have only occurred in the last two 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 modelers, in economics, engineering, and computer science.

Agent-based market models have used two kinds of learning: evolutionary learning algorithms, such as Holland's GAs or Koza's GP; and versions of reinforcement learning, such as the Roth-Erev model and modifications. On the one hand, Nicolaisen et al. (2001) believe that the social learning implicit in the GA militates against the expected emergence of the exercise of market power in agent-based models, while a version of Roth-Erev is sufficient for its emergence. On the other hand, Phelps et al. (2002) believe that a GP model of learning in electricity markets is a better model in which to design the auction by including the auction rules in the search space of the GP algorithm, as well as including the buyers' and sellers' strategies. It remains a challenge to reconcile the power of evolutionary algorithms in searching a complex design space for agents' strategies and auction rules with the greater realism (but less effective exploration and exploitation of the design space) of models using individual reinforcement learning.

In this chapter, we have discussed the meaning of market design, its challenges, and the use of agent-based models, first, to analyze electricity markets, and, second, to attempt to design such markets directly. It has been impractical to mention all or even most modeling efforts in the literature, and we have focused on the pioneering efforts and the most successful efforts so far. Nonetheless, the future development of the field of agent-based market design will flourish, as evidenced by the large numbers of researchers in different disciplines across the Internet now involved in advancing our knowledge and understanding.

6. Appendix: Models Rule

As a simulator for sixteen years, I have tried in the past to anticipate criticisms of my chosen methodology, as well, of course, as marshaling the advantages of simulation. One of the most telling criticisms is that simulations can only determine sufficiency: if you set these parameters so, then you'll observe the following outcome. With a closed-form mathematical characterization, it is possible, at least in principle, to determine necessity, as well as sufficiency: to observe a specific outcome requires (necessitates) one of the following combinations of parameters — no other combinations will do. Such a conclusion is not generally available to the simulator.

As a result, simulation — although in cases of intractable mathematical formulations the only way to get any results, even if these results are merely sufficient conditions, a subset of the underivable necessary conditions — is accepted, but hardly acclaimed.

Recently I was reading about discovery of the structure of DNA by Watson and Crick fifty years ago (Richards 2003). I hadn't properly registered the fact that, following Linus Pauling, they were building physical models of the mysterious molecule. Pauling had rushed into publication with his own model, a three-chain helix (Pauling & Corey 1953). This model, however, had an elementary error: chemically it could not be an acid — remember, deoxyribonucleic *acid*!

Crick and Watson had already been tinkering with models cobbled together out of sheet metal plates and brass rods, all propped up by retort stands and clamps. That is, they were simulating the molecule's structure, given whatever information was available: the chemical composition of DNA, the relative sizes and charges of the atoms, the chemical properties, and the potential biological properties of the molecule. With X-ray photos from Rosalind Franklin at King's College, they redoubled their efforts at cracking the structure.

On 28 February 1953, regulars at the old Cambridge pub, The Eagle, at the end of a cobbled courtyard off Bene't Street, became the first to learn the news that the secret of the procreation of life had been cracked. Using simulation! Models rule!

Of course, as Pauling learnt to his embarrassment, these were models of the unknown structure with few degrees of freedom: physics, chemistry, and biology each imposed restrictions on the arrangement of the atoms and sub-molecules in the DNA structure. Pauling's triple helix had earlier been considered and then abandoned by Crick and Watson, after advice from Franklin at King's.

Model-building ("stereo-chemical arguments" in Watson & Crick's 1953 phrase) 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 rods and sheet metal.

In general, simulation to explore emerging behaviour of systems can be seen as a process of induction — inferring general principles from the observation of many particular instances, as opposed to the process of deduction — deriving particular properties from more general principles; or asking what the necessary conditions are to obtain particular properties of a system.

Induction was first claimed as a means of scientific advance by Francis Bacon, in 1620. To work, induction must be properly applied. Specifically, conclusions, Bacon argued, had to be grounded in relevant observation. One negative, or false instance, would always undermine a host of positives: the particular is stronger than the general.

For the simulator, with many degrees of freedom, this focus underlines the importance of the kind of sensitivity analysis known as Monte Carlo (after the casino in Monaco), whereby the simulation is performed many times with different parameters. As Judd (1998) discusses at length in his Chapters 8 and 9, we cannot obtain truly random samples to initiate our simulations with, but at best pseudo-random numbers. Instead, Judd suggests that so-called quasi-Monte Carlo methods (that do not rely on probabilistic ideas and pseudo-random sequences for constructing an initial sample and analyzing the outcome) might be used and, suitably constructed, even outperform true Monte Carlo methods.

Whatever the details of the simulations, or models, it is useful to remember that fifty years ago a simple physical model was the key to the most important discovery in biology for the past century and a half.

(A revised version of Marks, 2003.)

7. References.

[1] Andreoni, J., and J.H. Miller (1995), Auctions with artificial adaptive agents, *Games and Economic Behavior,* 10: 38–64.

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

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

[4] Axelrod, R., (1987), The evolution of strategies in the iterated Prisoner's Dilemma, in *Genetic Algorithms and Simulated Annealing,* L. Davis (ed.), London: Pittman.

[5] Bacon, F. (2000), *Novum Organum*, ed. by L. Jardine and M. Silverthorne, Cambridge University Press, Cambridge.

[6] Boutilier C., Y. Shoham and M.P. Wellman (1997), Editorial, Economic Principles of Multi-Agent Systems, *Journal of Artificial Intelligence,* 94(1-2): 1–6.

[7] Bower J., and D.W. Bunn (2000), Model-based comparison of pool and bilateral markets for electricity, *Energy Journal,* 21(3): 1–29.

[8] Bower J., and D.W. Bunn (2001), Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the England and Wales electricity market, *Journal of Economic Dynamics and Control,* 25(3–4): xx

[9] Bower, J., D.W. Bunn, and C. Wattendrup (2001), A model-based analysis of strategic consolidation in the German electricity industry, *Energy Policy,* 29: 987–1005.

[10] Bunn, D.W. (2004), personal communication.

[11] Bunn, D.W., and C.J. Day (1998), Agent-based simulation of electric power pools, *Proc. 19th International Association of Energy Economists North American Conference.*

[12] Bunn, D.W., and F.S. Oliveira (2001), Agent-Based Simulation: An Application to the New Electricity Trading Arrangements of England and Wales, *IEEE Transactions on Evolutionary Computation*, 5(5): 493-503, October.

[13] Bunn, D.W., and F.S. Oliveira (2003), Evaluating individual market power in electricity markets via agent-based simulation, *Annals of Operations Research,* 121: 57–77.

[14] Burkhart, R., M. Askenazi, and N. Minar (n.d.), Swarm Release Documentation. Available as http://www.santafe.edu/projects/swarm/swarmdocs/set/set.html.

[15] Byde, A. (2002), Applying Evolutionary Game Theory to Auction Mechanism Design, Hewlett-Packard Company, Report HPL-2002-321.

[16] Cau, T.D.H. (2003), Analyzing tacit collusion in oligopolistic electricity markets using a co-evolutionary approach, PhD dissertation, Australian Graduate School of Management, University of New South Wales.

[17] Cau, T.D.H., and E.J. Anderson (2002), A co-evolutionary approach to modeling the behavior of participants in competitive electricity markets, *Proc. IEEE Power Engineering Society, Summer Meeting,* 3: 1534–1540.

[18] Chattoe, E. (1998), Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes? *Journal of Artificial Societies and Social Simulation,* 1(3). <http://www.soc.surrey.ac.uk/JASSS/1/3/2.html>

[19] Cliff, D., (2001a) Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double auction markets, TR HPL-2001-99, HP Labs, 2001.

[20] Cliff, D., (2001b) Evolution of market mechanism through a continuous space of auction-types, HPL-2001-326

[21] Cliff, D. (2002), Evolution of market mechanism through a continuous space of auction-types II: Two-sided auction mechanisms evolve in response to market shocks, HPL-2002-128

[22] Cliff, D., (2003a), Explorations in Evolutionary Design of Online Auction Market Mechanisms, *Journal of Electronic Commerce Research and Applications,* 2(2): 162–175.

[23] Cliff, D. (2003b), Artificial trading agents for online auction marketplaces, http://www.hpl.hp.com/personal/dave_cliff/traders.htm Accessed 15 July 2004.

[24] Collier, N. and D. Sallach (2001), RePast. Available at http://repast.sourceforge.net/.

[25] Cramton, P. (2003), Electricity Market Design: The Good, the Bad, and the Ugly, *Proceedings of the Hawaii International Conference on System Sciences, January, 2003.*

[26] Curzon Price, T. (1997) Using co-evolutionary programming to simulate strategic behavior in markets, *Journal of Evolutionary Economics,* 7(3): 219–254.

[27] Das, R., J.E. Hanson, J.O. Kephart and G. Tesauro (2001) Agent-Human Interactions in the Continuous Double Auction, *The Proceedings of the 17th International Joint Conferences on Artificial Intelligence (IJCAI), Seattle,* ed. by B. Nebel, S.F.: Morgan Kaufmann, pp. 1169–1187.

[28] Dawid, H., (1999), *Adaptive Learning By Genetic Algorithms: Analytical Results and Applications to Economic Models,* Berlin: Springer, 2nd ed.

[29] 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

[30] Entriken, R., and S. Wan (2003), Agent-Based Simulation of an Automatic Mitigation Procedure, Palo Alto: EPRI.

[31] Forrester, J.W. (1961), *Industrial Dynamics,* Cambridge, Mass.: M.I.T. Press.

[32] Gamma, E., R. Helm, R. Johnson, and J. Vlissides (1995), *Design Patterns: Elements of Reusable Object-Oriented Software,* Reading, Mass.: Addison-Wesley.

[33] Gode, D., and 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.

[34] Hämäläinen, R.P. (1992), Multi-agent interaction processes: from oligopoly theory to decentralized artificial intelligence, presented at the Mathematical Organization Workshop, Orlando, April 25–26.

[35] Hämäläinen, R.P. (1996), Agent-Based Modeling of the Electricity Distribution System, *Proc. of the 15th International Association of Science and Technology for Development (IASTED) International Conference Modelling, Identification and Control,* Innsbruck, Austria, February 19–21, 1996, pp. 344–346.

[36] Hämäläinen, R.P. (2004), personal communication.

[37] Hämäläinen, R.P., E. Kettunen, and H. Ehtamo (1997), Game modeling and coordination processes for two-level multi-agent systems, *Proc. of the 16th IASTED International Conference: Modelling, Identification, and Control,* Feb. 17–19, Innsbruck.

[38] Hämäläinen, R.P., and J. Parantainen (1995), Load Analysis by Agent-Based Simulation of the Electricity Distribution System, *Proc. of the 2nd IFAC Symposium on Control of Power Plants and Power Systems,* SIPOWER´95, Cancun, Mexico, December 6-8, 1995, 1: 213–217.

[39] Harp, S.A., A. Brignone, B.F. Wollenberg, and T. Samad (2000), SEPIA: a simulator for Electric Power Industry Agents, *IEEE Control Systems Magazine,* 20(4): 53–69.

[40] Hoecker, J. (1998), Keeping Electric Restructuring Moving Forward (Feb. 3, 1998) (11th Annual Utility M&A Symposium, New York).

[41] Holland J.H., and J.H. Miller (1991), Artificial Adaptive Agents in Economic Theory. *American Economic Review,* 81(2): 365–370.

[42] Holland, J.H. (1992), *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence,* Camb.: MIT Press (1976), 2nd ed.

[43] Hu, J., and M.P. Wellman (1998), Multiagent reinforcement learning: theoretical framework and an algorithm. In *Proceedings of the Fifteenth International Conference on Machine Learning,* 242−250. Morgan Kaufmann.

[44] Judd, K.L., (1998), *Numerical Methods in Economics*, M.I.T. Press, Cambridge.

[45] Kauffman, S.A., (1995), *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity,* NY: OUP.

[46] Knuth, D. E. (1968−1973), *The art of computer programming,* Reading, Mass.: Addison-Wesley.

[47] Koza, J. R., (1993), *Genetic programming: On the programming of computers by means of natural selection,* Cambridge: MIT Press.

[48] Krishna, V., and V.C. Ramesh (1998), Intelligent agents in negotiations in market games, Part 2, Application, *IEEE Trans. Power Systems,* 13(3): 1109−1114.

[49] Lane, D.A., (1993a), Artificial worlds and economics, part I, *Journal of Evolutionary Economics,* 3: 89−107.

[50] Lane, D.A., (1993b), Artificial worlds and economics, part II, *Journal of Evolutionary Economics,* 3: 177−197.

[51] Lane, D., A. Kroujiline, V. Petrov, and G. Sheblé (2000), Electricity market power: marginal cost and relative capacity effects, *Proc. of the 2000 Congress on Evolutionary Computation,* Vol. II, ed. by A. Alzala, Piscataway, N.J., IEEE Society Press, pp. 1048−1054.

[52] LeBaron, B., W. B. Arthur, and R. Palmer (1999) Time series properties of an artificial stock market, *Journal of Economic Dynamics & Control,* 23(9−10): 1487−1516.

[53] "Linus Pauling and the Race for DNA: A Documentary History," Special Collections, Oregon State University, http://osulibrary.orst.edu/specialcollections/coll/pauling/dna/index.html Accessed on 2003/05/21

[54] Maifeld, T., and G. Sheblé (1996), Genetic-based unit commitment, *IEEE Trans. on Power Systems,* 11(3): 1359.

[55] MacGill, I.F., (2004), Exploring spot electricity market operation through agent-based simulation and evolutionary programming, Canberra: CSIRO Agent-Based Modeling Seminar, February.

[56] MacGill, I.F., and R.J. Kaye (1999), Decentralized coordination of power system operation using dual evolutionary programming, *IEEE Trans. Power Systems,* 14(1): 112−119.

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

[58] Marks, R.E., (1992), "Breeding hybrid strategies: optimal behaviour for oligopolists," *Journal of Evolutionary Economics*, **2**: 17−38.

[59] Marks, R.E., (1998), "Evolved perception and behaviour in oligopolies," *Journal of Economic Dynamics and Control,* **22**(8−9): 1209−1233.

[60] Marks, R.E., (2003), Models rule, Editorial, *Australian Journal of Management,* 28(1): i−ii.

[61] McMillan, J. (2002), *Reinventing the Bazaar: A Natural History of Markets,* NY: Norton.

[62] Milgrom P. (2004), *Putting Auction Theory to Work,* Cambridge University Press.

[63] Nicolaisen, J., M. Smith, V. Petrov, and L. Tesfatsion (2000), Concentration and capacity effects on electricity market power, *Proc. 2000 Congress on Evolutionary Computation,* ed. by A. Alzala, Piscataway, N.J., IEEE Society Press, pp. 1041−1047.

[64] Nicolaisen, J., V. Petrov, and L. Tesfatsion (2001), Market Power and Efficiency in a Computational Electricity Market with Discriminatory Double-Auction Pricing, Economic Report No. 52, Iowa State University, revised August. The final version of this paper is published in *IEEE Transactions on Evolutionary Computation,* 5(5): 504−523.

[65] North, M.J., (2000), SMART II: The Spot Market Agent Research Tool Version 2.0, *Proceedings of SwarmFest 2000,* pp. 9-13, Swarm Development Group, Logan, UT.

[66] North, M.J., C. Macal, R. Cirillo, G. Conzelmann, V. Koritarov, P. Thimmapuram, and T. Veselka, (2001), Multi-Agent Social and Organizational Modeling of Electric Power and Natural Gas Markets, *Computational & Mathematical Organization Theory,* 7(4): 331−337.

[67] North, M., G. Conzelmann, V. Koritarov, C. Macal, P. Thimmapuram, and T. Veselka (2002), E-Laboratories: agent-based modeling of electricity markets, for submission to 2002 American Power Conference, Chicago.

[68] Pauling, L. and R.B. Corey (1953), A proposed structure for the nucleic acids, *Proceedings of the National Academy of Sciences*, 39(2): 84−97.

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

[70] Räsänen, M., R.P. Hämäläinen, and J. Ruusunen (1994), Visual interactive modeling in electricity load analysis, ul Proc. of the 13th International Association of Science and Technology for Development (IASTED) International Conference Modelling, Identification and Control, Grindelwald, Switzerland, Feb. 21−23, pp. 339−342.

[71] Richards, P. (2003), Life class, *Cam: Cambridge Alumni Magazine,* 38: 10–13, Lent Term.

[72] Richter, C.W., and G. Sheblé (1998), Genetic algorithm evolution of utility bidding strategies for the competitive marketplace, *IEEE Trans. Power Systems,* 13(1): 256–261.

[73] Richter, C.W., Jr., G.B. Sheblé, and D. Ashlock (1999), Comprehensive bidding strategies with genetic programming/finite state automata, *IEEE Transactions on Power Systems,* 14(4): 1207–1212.

[74] Riechmann, T., (2001), Genetic algorithm learning and evolutionary games, *Journal of Economic Dynamics & Control,* 25: 1019–1037.

[75] Roth, A.E., (1991), Game Theory as a Part of Empirical Economics, *Economic Journal,* 101(401): 107–14.

[76] Roth, A.E., (2000), Game theory as a tool for market design, In F. Patrone, I. García-Jurado,S. Tijs (eds.) *Game Practice: Contributions from Applied Game Theory,* Kluwer, pp. 7–18.

[77] Roth, A.E., (2002), The economist as engineer: game theory, experimentation, and computation as tools for design economics, *Econometrica,* 70(4): 1341–1378.

[78] Roth, A.E., and I. Erev (1998), Predicting how people play games: reinforcement learning in experimental games with unique mixed strategy equilibria, *American Economic Review,* 88(4): 848–881.

[79] Sargent, T.J., (1993), *Bounded Rationality in Macroeconomics,* NY: OUP.

[80] Simon, H., (1981), *The Sciences of the Artificial,* 2nd ed., Cambridge: The MIT Press.

[81] Talukdar, S., (2002), Agent-based market testing, DOE Transmission Reliability Research Review, Washington DC, December 10.

[82] Talukdar, S., and V. C. Ramesh (1992), A-teams for real-time operations, *International Journal of Electrical Power & Energy Systems,* 14(2–3): 138–143.

[83] Tesfatsion, L., (n.d.), Electricity market design: an agent-based computational economics approach,

[84] Tesfatsion, L., (2002), Agent-Based Computational Economics: Growing Economies from the Bottom Up, Economics Working Paper No. 1, Iowa State University, March. The final published version of this survey appears in *Artificial Life*, 8(1): 55–82, published by the MIT Press October 2001, pp. 504-523.

[85] Thorndike, E., (1911), *Animal Intelligence: Experimental Studies,* NY: Macmillan.

[86] Verkama, M., R.P. Hämäläinen, and H. Ehtamo (1992), Multi-agent interaction processes: from oligopoly theory to decentralized artificial intelligence, *Group Decision and Negotiation,* 1(2): 137–159.

[87] Verkama, M., R.P. Hämäläinen, and H. Ehtamo (1994), Modeling and computational analysis of reactive behavior in organizations, *Computational Organization Theory,* ed. by K.M. Carley and M.J. Prietula, Hillsdale, N.J.: Lawrence Erlbaum Assoc., 161−177.

[88] Vriend, N., (2000), An Illustration of the Essential Difference Between Individual and Social Learning and its Consequences for Computational Analyses, *Journal of Economic Dynamics and Control,* 24: 1−19.

[89] Walia, V., A. Byde, and D. Cliff (2002), Evolving Market Design in Zero-Intelligence Trader Markets, HPL-2002-290 http://www.hpl.hp.com/techreports/2002/HPL-2002-290.html

[90] Walsh W.E., R. Das, G. Tesauro, and J.O. Kephart (2002), Analyzing complex strategic interactions in multi-agent systems, *Proceedings of the 4th Workshop on Game Theoretic and Decision Theoretic Agents,* ed. by P. Gymtrasiwicz and S. Parsons.

[91] Watkins, C.J.C.H., and P. Dayan (1992), Q-learning, *Machine Learning,* 8: 279–292.

[92] Watson, J.D., and F.H.C. Crick (1953), Molecular structure of nucleic acids: A structure of deoxyribose nucleic acid, *Nature,* 4356: 737−8, April 25. http://osulibrary.orst.edu/specialcollections/coll/pauling/dna/papers/msna.html Accessed 2003/05/27

[93] Weibull, J.W., (1995), *Evolutionary Game Theory,* Cambridge: The MIT Press.

[94] Wilson, R., (2002), Architecture of power markets, *Econometrica,* 70(4): 1299−1340.

[95] Wurman, P.R., M. P. Wellman, and W. E. Walsh (2001) A Parameterization of the Auction Design Space, *Games and Economic Behavior,* 35: 304–338.

[96] Young, H.P., (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions,* Princeton: P.U.P.