Market Design Using Agent-Based Models

Robert Marks

May 17, 2005

Australian Graduate School of Management
The Universities of Sydney and New South Wales
Sydney, NSW 2052, Australia
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.

Contents

Abstract

This chapter explores the state of the emerging practice of designing markets by the use of agent-based modeling, with special reference to electricity markets and computerized (on-line) markets, perhaps including real-life electronic agents as well as human traders. The paper first reviews the use of evolutionary and agent-based techniques of analyzing market behaviors and market mechanisms, and economic models of learning, comparing genetic algorithms with reinforcement learning. Ideal design would be direct optimization of an objective function, but in practice the complexity of markets and traders' behavior prevents this, except in special circumstances. Instead, iterative analysis, subject to design criteria trade-offs, using autonomous self-interested agents, mimics the bottom-up evolution of historical market mechanisms by trial and error. The chapter highlights ten papers that exemplify recent progress in agent-based evolutionary analysis and design of markets in silico, using electricity markets and on-line double auctions as illustrations. A monopoly sealed-bid auction is examined in the tenth paper, and a new auction mechanism is evolved and analyzed. The chapter concludes that, as modeling the learning and behavior of traders improves, and as the software and hardware available for modeling and analysis improves, the techniques will provide ever greater insights into improving the designs of existing markets, and facilitating the design of new markets.

Keywords: market analysis design auctions learning electricity on-line

JEL codes: D440, D490, C150, C630, C790

1 Introduction

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 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. See LeBaron (2005) for further discussion of this research.

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, 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 hiccups. Perhaps as an offshoot of the privatization of government assets and activities in the 1980s in many countries, the use of auctions to choose the new owners and to value these assets slowly replaced so-called "beauty contests," in which subject to certain technical requirements licenses were virtually given away. But these new auction mechanisms at first did not allow for the complementary nature of bands in different localities. Only after intensive efforts by economists advising, first, governments, and, second, bidding companies did the successful "3G" auctions occur [Milgrom (2004)].

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.[1]

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.

[1]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 [Sweeney (2002)].

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 might be used, but negotiation and design of contracts by use of computer simulation and agent-based modeling is only now emerging from its infancy.[2]

As examples of the use of agent-based models in market design, this chapter will examine the use of such models in designing the fourth type of market, that for electricity, and the fifth, for on-line trading, which is also examined in MacKie-Mason and Wellman (2005). The first, for emissions abatement, is covered by Janssen and Ostrom (2005).[3] The second is covered by the chapter by LeBaron (2005), and referred to further below. The chapter by Duffy (2005) provides evidence of validation of artificial ("designed") agents and the behavior of human subjects in experiments, as discussed below.

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, non-agent-based simulation with agent-based simulation of analysis and design, and finally different models of learning and adaptation in agent-based simulation models.

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

[2] A good starting point is Jennings et al. (2001).

[3] Agent-based models have also been used in other environmental settings: Hailu and Schilizzi (2004).

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 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 — 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 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 encyclopedic study of *The Art of Computer Programming* in three volumes, but such a task would be daunting now.[4]

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 section will be on analysis, rather than design. This is because, as we discuss in section 3.1 below, direct design or optimization

[4]Apparently, Knuth has been undertaking a fourth volume, since TeX and METAFONT were put to bed [Knuth (1979)].

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, direct design might never be possible, as Edmonds and Bryson (2003) remind us. Instead, searching the design space will be an iterative process of analyzing the performance of a specific model, modifying the model in the light of this analysis, and analyzing the modified model, until the designer is happy with the performance of the multi-modified model against various criteria.

2.3 Evolutionary Simulation Techniques

To the evolutionary biologist, the design is the genotype, and the performance is the phenotype. Evolution can be characterized as a dynamic search in a population for genotypes that result in better phenotypes, where that mapping too is ill-defined. It might not be surprising, therefore, that the development of agent-based methods of optimization and simulation began with 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. The GA tests and scores individual solutions in a population of possible solutions, and, based on the "fitness" score of each, selects pairs of "parents" for a new "offspring" generation of possible solutions. This artificial reproduction uses the genetic operations of "crossover" and "mutation" (analogous to mimicry of existing solutions and to exploration of new regimes of the solution space) on the parents. Testing, selection, and generation of a new population results in the emergence of never-worse best solutions. GA has been widely used as an optimizer, a directed form of trial and error that obviates exhaustive testing of all possibilities.

But using the GA as an optimizer in this way — focusing on the single best solution (an individual) — throws away the population's emerged characteristics qua population. A line of research then began with Axelrod's (1987) simulation of individuals playing the Iterated Prisoner's Dilemma (IPD). It 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 [see Axelrod (2005)], a manifestation of the Folk Theorem of repeated games [Fudenberg and Maskin (1986)]. But since the IPD can be thought of as a simple model of a repeated 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.[5]

When 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 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.

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.[6] 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: 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.

Chattoe (1998) argues that GA applications in economics confuse the role of the GA as instrumental in searching the solution space and its role as a description of firms' decision-making and individual learning. 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. Indeed, Curzon Price (1997) spoke of the GA as providing a stream

[5]Just how to characterize out-of-equilibrium behavior (or bounded rationality, for that matter) remains an open question. See Arthur (2005).

[6]This process was mistakenly called boot-strapping by Marks (1989), in the first published research into co-evolution of rivals' strategies in oligopolies.

of hypothetical scenarios within the firm, even if not all are acted upon. Duffy (2005) provides an extensive review of the descriptive role of GAs in economic models, and concludes that the findings from many studies "provide some support for the reasonableness of GAs as models of adaptive learning by populations of heterogeneous agents."

When applied to economic systems, the GA operators have been interpreted several ways. Each individual string can represent either an individual agent or one possible decision of a single agent. The selection operator ensures that past performance is reflected in future choices: well (badly) performing decisions are more (less) likely to be chosen in the future. Each new generations of strings might be new individual decision-makers, or it might be new ideas or heuristics among long-lived players.

With few exceptions, the models of analysis and design that we discuss below are evolutionary in nature — "dynamic models in which successful agents and activities gradually increase their share of the economy at the expense of less successful agents and activities" [Conlisk (1996)] — whether explicitly so (as with GAs) or implicitly.

2.4 Learning

The populations in the first applications of GAs were seen as trial solutions to arguments that would optimize the function in question (usually highly non-linear and discontinuous). Later applications, however, treated the populations as comprised of agents rather than numbers. Individual agents were immutable, but in each generation the population of agents would change, under selective pressure. This is implicit learning and adaptation.[7] Just how learning and adaptation are modeled can clearly affect the model's behavior.

Agent-based modeling has since modeled learning as explicit. Arthur (1991, 1993) was the first economist to support modeling agent behavior using reinforcement-learning (RL) algorithms and to calibrate the parameters of such learning models using data from human-subject experiments.[8] In RL models, 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

[7] "Implicit" in that the individual agents do not change at all, but succeeding populations embody improvements ("learning") in the manner of response. Wood (2005) points out that psychological experiments have shown that for human subjects learning can be adaptive, but that adaptation does not necessarily imply learning, the long-term rewriting of memory.

[8] Brenner (2005, Section 2.1) recounts how Arthur generalized the Bush and Mosteller (1955) model, also used by Cross (1973, 1983).

(1911)].[9] At first, Arthur was interested in calibrating *individual* learning to experimental data, but later he and his associates [Arthur et al., (1997)] "model calibrations that yield aggregate data that are similar to relevant field data" [Duffy (2005)].

Roth and Erev (1995) and Erev and Roth (1998) ask how well RL algorithms track experimental data across various multi-player games. Their general RL model, which improves the fit of the model to human-subject experimental data, includes Arthur's earlier model as a subset, as seen below.

The general Roth-Erev model of reinforcement learning can be characterized as follows: Suppose there are N actions/pure strategies. In round t player i has a propensity $q_{ij}(t)$ to play the j^{th} pure strategy, where propensities are equivalent to strengths in Arthur's model. Initial propensities are equal, $q_{ij}(1) = q_{ik}(1)$ for all available strategies j, k, and $\sum_j q_{ij}(1) = S_i(1)$, where $S_i(1)$ is an initial strength parameter, equal to a constant that is the same for all players, $S_i(1) = S(1)$; the rate of learning is proportional to the size of $S(1)$:

$$(1) \quad \sum_j q_{ij}(1) = S_i(1) = S(1) \text{ for all } i.$$

The probability that agent i plays strategy j in period t is made according to the linear choice rule:

$$(2) \quad p_{ij}(t) = \frac{q_{ij}(t)}{\sum_{k=1}^{N} q_{ik}(t)}$$

Suppose that, in round t, player i plays strategy k and receives payoff of x. Let $R(x) = x - x_{\min}$, where $x_{\min}$ is the smallest possible payoff. Then player i updates his propensity to play action j according to the rule:

$$(3) \quad q_{ij}(t+1) = (1-\phi)q_{ij}(t) + E_k(j, R(x)),$$

$$(4) \quad \text{where } E_k(j, R(x)) = \begin{cases} (1-\epsilon)R(x), & \text{if } j = k; \\ \frac{\epsilon}{N-1} R(x), & \text{otherwise.} \end{cases}$$

This is a three-parameter model, where the parameters are: the initial-strength parameter, $S(1)$; a recency parameter ϕ that gradually reduces the power of past experiences to influence future actions; and an

[9] Recent psychological research is questioning Thorndike's Law of Effect: the more specific and immediate the feedback, the greater the effect on learning. The Law is a reasonable description of human behavior in a simple world (of decision-making), but is not so good in a complex, stochastic world (of exploration and problem-solving) [Wood (2005)].

experimentation parameter ϵ, which can be localized for similar strategies, or be made more intrepid.

If $\phi = \epsilon = 0$ then the model becomes a version of Arthur's model, but without re-normalization of the sum of propensities in every period. The model without re-normalization reflects a learning curve that flattens with experience over time.

Duffy (2005) and Brenner (2005) discuss, among others, four types of RL models: the Arthur-Roth-Erev model mentioned above; Q-learning, which optimizes long-term payoffs, rather than immediate payoffs [Watkins and Dayan (1992)]; multi-agent Q-learning [Hu and Wellman (1998)]; and Adaptive Play [Young (1998)]. Below we discuss several papers that use these models, including a useful modification of the Roth-Erev model summarized above in equations (1) through (4).

Selten [Selten and Stoecker (1986), Selten (1998)] has devised a much simpler learning mechanism, directed learning. This is based on the notion that ex-post rationality is the strongest influence in adaptive behavior. It requires an ordering over the set of possible actions, and models players learning to do better by probabilistically altering their actions in the direction that would have led to higher payoffs had these actions been chosen earlier, and never altering their actions in a direction that would have lowered their payoffs [Brenner (2005)]. For instance, Hailu and Schilizzi (2004) model bidders' learning in auctions: if a bidder won the previous auction, then choose an equi-probable mixed action of the same bid or one ten percent higher for the next auction; if the bidder did not win in the previous auction, then choose an equi-probable mixed action of the same bid or one ten percent lower, with prior upper and lower limits to legitimate bids. They find that the efficiency benefits of one-shot auctions dissipate with repetition and learning.

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).[10] Both sorts of models, and both sorts of learning, have been termed "agent-based" models.

The learning in reinforcement-based models and in the extant GA models is inductive: that is, future actions are based on past experience, with no attempt to anticipate and reason back, in a truly deductive, strategic fashion. Belief-based learning, however, incorporates recognition by the

[10] Strictly speaking, individual learning can also be modeled using classifier systems, closely related to the GA [Holland (1992)].

players that they are interacting with other players. They thus form beliefs about the likely actions of these other players. "Their choice of strategy is then a best response to their beliefs", [Duffy (2005), Section 3.2]. "By contrast, reinforcement learners do not form beliefs about other players, and need not even realize that they are playing a game or participating in a market with others." Almost all the research we review models inductive learning, but two papers which use anticipatory, belief-based learning are reviewed in Section 3.4 below.

2.5 From Analysis to Design

As remarked by Roth (1991) in an earlier 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; 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 discussed further below. As Walia et al. (2003) 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 might expect 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.

But this is not necessarily the case: for instance, a market design could have a loophole — obscure to stupid agents — that makes it completely degenerate. Even without loopholes, smarter agents might find strategic ploys that reduce efficiency, or might spend more effort (wasted, from a social efficiency perspective) on counter-speculation.[11] This is confirmed by Arifovic (2001), who finds that more complicated agents do not necessarily do better in her simulated market environment.

Of course, 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

[11] I thank an anonymous referee for pointing this out.

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.

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 (2005)] — 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 and 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 LeBaron (2005) and Duffy (2005).

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. See Section 3.2 below.

3 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.

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 had simulated the structure correctly. Model-building ("stereo-chemical arguments" in Watson and 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. See Marks (2003) for further discussion of this pioneering simulation.)

MacKie-Mason and Wellman (2005) present a Marketplace Design Framework, which delineates the three fundamental steps that constitute a transaction: first, the connection (searching for and discovering the opportunity to engage in a market interaction); second, the deal (negotiating and agreeing to terms); and, third, the exchange (executing a transaction). They define a "marketplace system" as consisting of agents and the market mechanism through which they interact, all embedded in an environment of social institutions (language, laws, etc.). Their market mechanism is the set of "rules, practices, and social structures of a social choice process, specifying, first, permissible actions" (including messages), and, second, market-based exchange transactions as outcomes of a function of agent messages. If there is some entity, apart from the participating agents, that manages any inter-agent communication and implements the mechanism rules, then the market mechanism is mediated.

MacKie-Mason and Wellman note that, as a consequence of this characterization of a marketplace, there are at least two design decisions:

first, the design of the market mechanism, which might be decomposed into the design of mechanisms for, successively, the connection, the deal, and the exchange phases of a transaction; and, second, design of agents to interact with the market mechanism, whether existing or newly designed. They define an agent as an "autonomous decision-making locus in a system of multiple decision-making entities"; an agent has "type" attributes, such as preferences, beliefs, intentions, and capabilities. There will be a form of consistency between the agents' behavior, beliefs, and preferences, consistent with some principle of rationality. In this chapter, 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, this chapter focuses on market mechanisms that govern the settlement from allowable actions.

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

To be effective, design of the market mechanism must be measured, and will usually consist of a constrained optimization, even if not explicitly or directly. "No external subsidies" or "maintain horizontal equity" are two possible constraints given by MacKie-Mason and Wellman. We explore others below.

The general design problem has become designing a market mechanism that includes defining a set of concerns over which agents can interact, specifying rules of permissible actions, and rules for mapping from actions to settlement and outcomes.

3.1 Complexity of Design

Edmonds and Bryson (2003) speak of the syntactic complexity of design. This is the lack of a clear mapping from design to behavior: the only way to know the system's outcomes is to run the system and observe the emerging performance. 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.*

It is partly because of these complexities that direct design of markets is hardly ever attempted. Another reason is the possibility of conflicts among several design trade-offs.

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 seller (or producer) surplus and buyer (or consumer) 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. Or different criteria could be weighted to derive a single, scalar maximand.

Possible criteria for judging the design of a single-auction market might include [Phelps et al. (2002a), (2005)]: first, maximizing seller revenue: this has been one of the main criteria in the design of the spectrum auctions, most famously the 3G auctions [Milgrom (2004)]; second, maximizing market allocative efficiency: from a policy viewpoint and not a seller viewpoint this is a desirable attribute of a marketplace system; third, discouraging collusion, as a means to attaining the first and second criteria; fourth, discouraging predatory behavior, which will also help to maximize efficiency; fifth, discouraging entry-deterring behavior, again as a means of maximizing seller revenue (in a single (selling) auction the greater the number of potential bidders, the greater the seller revenue); sixth, budget balance: no third-party payments or subsidies for a deal to be reached; seventh, individual rationality: the expected net benefit to each participant from the market mechanism should be no less than the best alternative; and eighth, strategy-proofness: participants should not be able to gain from non-truth-telling behavior.

Myerson and Satterthwaite (1983) derived an impossibility result that demonstrates that no double-sided auction mechanism with discriminatory pricing[12] can be simultaneously efficient, budget-balanced, and individually rational.

Talukdar (2002) emphasizes that before the market can be designed (solved), the design problem must be well posed, that is, complete, feasible (all constraints can be satisfied), and rich (allows for innovative and desirable solutions). To be complete, the design problem must contain: first, the attributes to be used in characterizing behavior of the market; second, the decision variables to be used to characterize the structure; third, the goals to be attained (desired behaviors, laws, regulations); and, fourth, a computable mapping of decision variables into goals (does each point in decision space meet the goals?). This fourth requirement is achieved for complex design problems by iterative analysis, which can be achieved using agent-based simulation tools and agent-based verification tools, since such tools are open and modular.

LeBaron (2005), 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 modeled? 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 to be modeled. 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.

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

[12] In discriminatory-price auctions (sometimes known as "pay-as-bid" auctions), distinct trades in the same auction round occur at distinct prices; in uniform-price auctions, all trades in any given auction round occur at the same price.

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 IPD.

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 (CDAs); and, 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)].[13]

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 and Oliveira (2003) note that many researchers [including Erev and Roth (1998)] have shown that learning models predict people's behavior better than do Nash equilibria.

Moreover, 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

[13]Rubinstein (1998) elaborates on some of these bounds. Conlisk (1996) gives four reasons for incorporating bounded rationality into economic models: empirical evidence of limits to human cognition; successful performance of economic models embodying bounded rationality (including some surveyed here); sometimes unconvincing arguments in favor of unbounded rationality; and the costs of deliberation.

possible, but to see which are likely in reality, see how players learn and choose amongst them.

3.4 Explicit Use of Agents

It is possible to design without the use of agents: given a market with demand and supply schedules, economic efficiency is maximized at the output level where marginal value equals the marginal unit cost, no matter how the social surplus is divided between buyers and sellers. But such direct design (optimization) requires a well defined problem. With several design trade-offs and the possible emergence of unforeseen performance in the system, agent-based analysis and design, in which the market system can be modeled as "evolving systems of autonomous, interacting agents with learning capabilities" [Koesrindartoto and Tesfatsion (2004)], is increasingly employed.

LeBaron (2005) 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.

Audet et al. (2002) report an agent-based study of micro-structure (order books v. dealers), while Bottazzi et al. (2003) examine tick sizes (and unexpectedly determines that smaller tick sizes do not necessarily improve the market's efficiency). Chan and Shelton (2001) examine how a model behaves with different RL mechanisms, all of which enable the optimum policy function for a market-making broker to be found.

Belief-based learning has been used to study market design: Gjerstad and Dickhaut (1998) propose heuristic rules by which buyers and sellers in a double auction will assess and update their probabilities that their bids (offers to buy) and asks (offers to sell) will be accepted, given market history. "Using these beliefs together with private information on valuations and costs, individual buyers or sellers propose bids or asks that maximize their (myopic) expected surplus" [Duffy (2005)]. The main parameter of their model is the length of memory that players use in calculating probabilities. Their model, with stricter convergence criteria than Gode and Sunder (1993) adopt, more reliably converges to competitive equilibrium, and the anticipatory, belief-based learning model provides a better fit to *aggregate* human-subject data as well. Gjerstad (2004) coins the phrase "heuristic belief learning" to describe this version of belief learning, and shows that what he calls "pace," the timing of the bid, is pivotal.

GA strings can be used to encode decisions that agents make (e.g., how much to consume, what price to charge, etc.) and the GA works to find the optimal decision, given feasibility and other constraints. This is how Marks et al. (1995), Arifovic (1994), Midgley et al. (1997) modeled the interactions. Duffy (2005) calls this learning-how-to-optimize.

An alternative is to use the strings as encoding beliefs about how prices will change from period to period. This learning-how-to-forecast model [Duffy (2005)] was first used by Bullard and Duffy (1999). It was introduced in order to calibrate the GA model with human-subject experiments of overlapping-generation decision models. Duffy (2005) explains that subjects found it easier to forecast future prices than to decide on inter-temporal consumption/saving decisions. Given the price forecasts, the GA algorithm solved that individual's optimal consumption/savings allocations and determined the market-clearing prices at future dates.

3.5 The Design Economist

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. (2002a)] 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 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. Simulation has occurred using GAs, numerical solutions, and explicit agent-based models. Iterative analysis has been used as a means of designing systems.

LeBaron (2005), 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. These are issues that have been addressed in the two areas of market design that we now consider: electricity markets and automated markets.

4 Electricity Market Design

In 1998 the U.S. Federal Energy Regulatory Commission (FERC) 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, 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 [Marks (1989)], and then to macro-economic models [Arifovic (1994)], 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 Electricity Market Design Trade-offs

As a consequence of the California blackouts of 2000, market efficiency has been joined by several other criteria for the design of electricity markets. The FERC (2003) White Paper discusses four primary objectives for wholesale electricity market design: reliable service (no blackouts or brownouts); fair and open access to the transmission grid at reasonable prices; effective price signals to provide incentives for appropriate investment in generation and transmission capacity; and effective procedures for market oversight and mitigation of exercise of market power.[14] Koesrindartoto and Tesfatsion (2004) speak of "efficient, orderly, and fair" market outcomes.

Cramton (2003) discusses issues of electricity market design, in general, and the mitigation of market power in particular. He also emphasizes that the market designer must understand the preferences and constraints of the market participants, in order to keep the design as simple as possible, but not too simple.[15] The greater the number of dimensions for measuring the performance of market designs, the greater the relative attractiveness of simulation as a design tool: as discussed above, closed-form analysis — with its promise of the derivation of necessary conditions — becomes ever more elusive.

4.2 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, 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.[16] In

[14] Nicolaisen et al. (2001) distinguish the exercise of structural market power that occurs when the buyers and sellers ask and bid their true reservation values, from the exercise of strategic market power that occurs when opportunistic bids or asks are made.

[15] Wilson (2002) surveys the experiences of electricity market designs in the U.S. and beyond at a much greater level of detail than has yet been seen even in simulation studies.

[16] In a private communication 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

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.

4.2.1 Hämäläinen et al. (1997) model both sides of the market.

Two years later Hämäläinen et al. (1997) went much further, with agents representing both sides of the electricity market — consumers and producers — with bounded reasoning capabilities and bounded reactions. Specifically, they use a two-hierarchy, multi-agent system to model a von Stackelberg market, where the leader (the seller) anticipates and reasons back to set a price for electricity which maximizes the overall market efficiency, given the responses of the followers (the buyers, who use electricity for space-heating). Agents can be modeled as: sufficiently rational to determine their best response dynamics; or as boundedly rational (and so not always succeeding in determining the best response, perhaps because of limited comparisons of possible actions); or as constrained in their reactions from one period to the next; or with asynchronous reactions.

The electricity price can vary hourly, and the electricity producer, announcing prices 24 hours ahead, can attempt to control consumption in order to smoothly costly load peaks. Each consumer takes account of the heat-storage capacity of its dwelling and the outside temperature. The consumer's payoff is the difference between the utility from consumption

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."

and the cost of the energy. The producer's cost function reflects the increasing (quadratic) marginal cost of production.

Using their Power Agents software for simulation of electricity markets, the authors gain insight into a market in which consumers could have their homes heated by a computerized agent-based heating system which would respond to changing electricity tariffs in order to maximize the system goals. It is not clear which form of bounded rationality they use. 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 and Ramesh (1992) suggested software to manage electricity generation when the operating environment (market) could change rapidly. Their asynchronous and autonomous agents represent one of the first examples of a multi-agent system in the electricity literature. Krishna and 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.

4.2.2 Talukdar (2002) models customers holding down the wholesale price.

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. His focus is on centralized auctions, without electricity storage facilities, where sellers have multiple blocks of energy to sell, and customers can adjust their demands and can automatically learn. He asks: What structures (load-adjustment facilities) do customers need so they can use automatic learning to hold the monopolistic price to reasonable levels?

His agents are simple: they are not intended to simulate human behavior; rather, the dynamics of the repeated markets are probed using the emergent behaviors (which can be quite complex) of simple agents. He finds that sellers with multiple generating units can learn which units to withhold and that total profits rise by a fifth over a thousand periods, with prices and profits almost at centralized (monopolistic) levels, under several (static) demand characterizations. He then allows buyers to learn too. They aim, first, to minimize cost, and, second, to minimize energy deviation. Over 1400 periods they learn to reduce the price to less than a third of the monopolistic price. But with the same quantity sold, the sellers' profits fall below 30% of the monopolist's profits.

Meanwhile, at Iowa State University, a group of electrical engineers led by

Gerald Sheblé had started in 1994 to examine the operation and design of electricity markets. Maifeld and Sheblé (1996) use a GA for solving the unit-commitment scheduling problem in electricity markets.[17] They referred to no earlier work by economists, but Richter and 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.

4.2.3 Lane et al. (2000) use GAs for double auctions.

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 discriminatory-price k-double auction [Satterthwaite and Williams (1989, 1993)], and, perhaps influenced by Tesfatsion's economics research, calculated the degrees of market power for various combinations of relative capacity and production costs.

They use the EPRI[18] Market Simulator, which simulates a double auction between buyers and sellers on a graph, where the edges are the capacity-constrained transmission lines and the buyers and sellers are at the nodes. The auction is performed in rounds or generations; buyers and sellers are matched in each round and the price of their contract is given by $kb + (1 - k)a$, where bid $b \geq$ ask a, and $0 \leq k \leq 1$; here $k = 0.5$. Learning is via a GA with binary strings. The GA treats buyers and sellers separately, and takes the risk-neutral traders' profits as their fitnesses.

The benchmarking simulation assumes price-taking agents. The buyers' profits and sellers' profits in a competitive equilibrium and in the auction simulation are determined (respectively, $PBCE$, $PSCE$, PBA, PSA), and an index of market power (MPI) is calculated:

$$(5) \quad MPI = \frac{(PBA + PSA) - (PBCE + PSCE)}{PBCE + PSCE}.$$

The simulated market has three sellers and three buyers (homogeneous, and unconstrained in their total demand, in order to allow sellers to play

[17] Unit commitment is the problem of determining the optimal set of generating units within a power system to be used up to a week ahead.

[18] The Electric Power Research Institute, Palo Alto, CA.

strategically). There are three scenarios for capacity constraints on sellers, increasingly unequal, and three scenarios for relative costs of the largest producer. Will this market converge to a near-competitive equilibrium? Will the largest seller exhibit price leadership and the exercise of market power?

Analysis of MPI averaged over 100 runs for each of the 3×3 treatments indicates that neither marginal cost nor market share has a significant effect on the large seller's exercise of market power, and on average total available surplus is evenly distributed among buyers and sellers. The authors identify an anomaly: relatively more expensive sellers gain market power, the opposite of what theory would suggest.

Why? The authors suggest four aspects of their model: limited numbers of buyers and sellers; the GA allows indirect exchange of information among sellers and among buyers through the genetic operators; a seller's (buyer's) ask (offer) price is its marginal cost (revenue) plus (minus) a real number derived from its bit-string; and calculation of profits in the auction (relative the averaged transaction price) is different from such calculation in competitive equilibrium (relative to the uniform market price). There is a rapid loss of genetic diversity, and each of the three sellers (buyers) will be tacitly working to solve the same maximization problem.

The authors argue that a GA is inappropriate when there are few agents, and conclude that another search method which incorporates memory and self-learning would be better. They do not mention the possibility of a GA with separate populations for the six distinct agents. Would such a model give results closer to those suggested by theory? I believe so.

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.

4.2.4 MacGill and Kaye (1999) simulate for system efficiency.

Meanwhile, engineers at the University of New South Wales [MacGill and Kaye (1999), MacGill (2004)] were exploring a decentralized coordination framework to maximize the market 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 via a decentralized framework in which each resource is operated to achieve overall system objectives. The authors use a a so-called "dual evolutionary approach," which uses a (non-binary coding) version of the GA, but not explicitly with autonomous, self-interested agents. Their model contained explicit intertemporal links for actions and payoffs (such as energy storage in, for example, pumped storage dams) across periods, and the dual evolutionary programming model, rather than optimizing on trajectories (the "primal" approach), uses as its variables the incremental benefit-to-go functions (the "dual"), which means that detailed knowledge of the resources models is not required.

They are able to solve for two-storage networks, with ramp rates, leakage, and stochastic supply (with photo-voltaics and hydro). A surprising result, when they allow strategic actions by players, is that the surplus of the sole strategic player is lower than its surplus without strategic actions. Is this a consequence of their modeling agents' fitness not as their individual surpluses but as the system's goals? They also find that the total system surplus falls with strategic actions, and the surplus on the load side falls most, as one might expect from theory.

Cau and 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, even if discouraging the exercise of market power is not an explicit goal of market design.

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.

4.3 Economists

4.3.1 Curzon Price (1997) models electricity markets.

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" (1997, p. 220), and hence suitable for simulation.

Curzon Price derived two models, both a simplification of the England and Wales electricity market, where the pool price is equal to the bid of the last producer required to satisfy demand, a uniform-price auction. The first model assumes that neither producer can supply the whole market, but that together their capacity exceeds demand. With price as the only choice variable, this model has two pure- and one mixed-strategy Nash equilibria. He was able to derive the pure-strategy equilibria, but not clearly the mixed-strategy equilibrium, even when he set the initial population proportions to the mixed-strategy proportions. For levels of GA crossover above 6% he found that the equilibrium mix could not be sustained. He concluded that, with an underlying situation of residual monopoly, the electricity pool rules would not lead to competitive prices, a finding of great significance to the market design.

His second model included the two producers' choice of capacities as well as prices. The first model was modified: if either producer could satisfy the entire market, then the lowest bid would be chosen. Players offered both price and quantity bids, the quantity offered incurring a cost whether or not the capacity was used. His analysis yielded three regimes: one where the high bidder is a residual monopolist; one where the low bidder can satisfy the demand; and one where there is excess demand because the higher bid is too high. The equilibrium strategies found by the GA can be characterized as similar to the first model without capacity as a strategic variable: one producer offering the lowest capacity possible and bidding it at the maximum price, and the other producer offering the highest residual quantity at a low price. The firms evolve their capacities to avoid Bertrand (marginal cost) outcomes.

Curzon Price's work was directly descended from Axelrod's (1987) work with GAs and IPDs, Marks' (1992) work on oligopolistic behavior, and other economists' use of GAs, such as Andreoni and Miller's (1995) exploration of auctions using the GA to model the co-evolution of artificial adaptive agents. Andreoni and Miller found that their model of adaptive learning was consistent with the main results from laboratory experiments, and that — significantly for the purpose at hand — various auction designs ("institutions") display very different adaptive dynamics. Curzon Price suggested that plausible behavioral elements could be included in the simulations.

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. These we now discuss.

4.3.2 Nicolaisen et al. (2000) search for market power.

Nicolaisen et al. (2000) used a GA agent-based model of a discriminatory-price clearinghouse[19] k-double auction electricity market [Klemperer (2002)] to examine the exercise of market power (as deviations from competitive equilibrium values of prices and quantities). They used the EPRI Power Market [see Lane et al. (2000), above], where each agent simultaneously submitted a single price-quantity bid or ask. Buyers and sellers are matched to maximize total profit, using $k = 0.5$ again. Each agent's fitness is proportional to its profit in the last round: only the last round's bid or ask is remembered. The linear revenue and cost functions ensures that bids and asks are at the capacity quantities. Bids (asks) are bound between [marginal revenue – \$40, marginal revenue] (marginal cost). Two definitions: first, the relative concentration of sellers NS to buyers NB, $RCON = \frac{NS}{NB}$; and, second, the relative capacity of buyers to sellers, $RCAP = \frac{NB}{NS} \times \frac{CB}{CS}$, where CB (CS) is the maximum quantity of electrical energy that each buyer (seller) can resell (generate) in a retail market. Six buyers and six sellers compete, with 3×3 treatments of three values of $RCON$ and three values of $RCAP$.

The authors derived sellers' market power, $MPS = \frac{PSA-PSCE}{PSCE}$, and buyers' market power, $MPB = \frac{PBA-PBCE}{PBCE}$. They found no evidence that MPB is negatively related to $RCAP$, or that MPS is positively related to $RCAP$, either in aggregate or individually, contrary to expectations from theory.

How could this be explained? As Tesfatsion (2005) notes, the measures of concentration and capacity ($RCON$ and $RCAP$) are structural characteristics of the market. As is standard in the industrial organization literature, they are calculated before any experiments have been run, and hence before the analyst knows which traders are inframarginal (and so will actually engage in trade) and which are extramarginal (and so will not engage in any trades). Because they do not trade, the bids/asks of extramarginal traders will have no affect on market power outcomes. As a result, by varying the numbers and capacities of the extramarginal traders, the concentration and capacity measures can be made arbitrarily large or small while keeping the market power measure constant. Consequently, so long as the extramarginal/inframarginal decision for each trader is endogenous [as in Nicolaisen et al. (2000)], no systematic relationship among $RCON$, $RCAP$, and market power outcomes will be seen.

In Nicolaisen et al. (2000), trading agents were quite boundedly rational, with only one round of memory. Moreover, the GA was given only two

[19] A clearinghouse (or call) market is one in which all traders place offers before the market is cleared; they can have discriminatory or uniform prices. A continuous market is one in which trades are executed as new offers arrive; prices are thus discriminatory.

populations (one for buyers and one for sellers), whereas the treatments meant that agents with different marginal costs and revenues faced different concentrations and capacities: the GA was not modeling this heterogeneity. Furthermore, the social learning process (mimicry of other buyers or other sellers) of the GA meant that any comparative advantages in strategies (as a consequence of different firm structures) soon spread to the rest of the population of players and became dissipated, as Vriend (2000) discussed. Moreover, social learning means that firms that would rightly decline to trade (the extramarginals) may now engage in opportunistic trades (and become inframarginal), thus potentially lowering market efficiency. The paper cites earlier work by Lane and by Richter, both at Iowa State.

4.3.3 Nicolaisen, Petrov, and Tesfatsion (2001) use reinforcement learning.

Following from their 2000 study (see above), Nicolaisen et al. (2001), henceforth referred to as NPT, altered their model by using a form of learning that, unlike the GA, did not impose strategic homogeneity on structurally distinct buyers and sellers. As well as mimicry, individual learning would be permitted. The model used the EPRI Power Model again, suitably modified, with the same 3×3 treatments of $RCON$ and $RCAP$, the same six buyers and sellers, as characterized by their (private) marginal revenues and costs, respectively.

But in an attempt to obtain results on market power that were closer to those from standard theory, NPT used reinforcement learning [a modification of Erev and Roth (1998)] instead of GA learning to allow individual learning and to prevent any comparative advantage in strategies being dissipated among the artificial agents. They point out that there are two shortcomings of the Roth-Erev model (see equations (1)–(4) above). First, there might be degeneracy of its parameters: when the experimentation parameter $\epsilon = \frac{N-1}{N}$, there is no updating of the choice parameter. Second, if there are zero profits, then the choice probabilities are not upgraded, because a trader's current propensity values are reduced proportionately. Lack of probability updating in response to zero profits can result in a substantial loss of market efficiency as traders struggle to learn how to make profitable price offers.

NPT present a simple modification of the Roth-Erev RL algorithm that addresses both of these issues while maintaining consistency with the learning principles in the original formulation. The update function $E(.)$ in equation (4) was replaced by the following modified function:

$$(6) \qquad E_k(j, R(x)) \;=\; \begin{cases} (1-\epsilon)R(x), & \text{if } j \;=\; k; \\ \frac{\epsilon}{N-1}\, q_{jk}, & \text{otherwise.} \end{cases}$$

In effect, this modification introduces a differentiated value for the recency parameter ϕ for selected versus non-selected actions, while also omitting the profit term in the updating equation for propensities corresponding to non-selected actions. The recency parameter for non-selected actions falls from ϕ to $\phi^* = \phi - \frac{\epsilon}{N-1}$. As NPT put it, "The choice probabilities corresponding to action choices resulting in zero-profit outcomes tend to decrease relative to other choice probabilities while the choice probabilities corresponding to action choices resulting in positive profit outcomes tend to increase." Otherwise the paper's model was similar to the earlier work [Nicolaisen et al. (2000)]: a clearinghouse k-double auction with discriminatory pricing, and $k = 0.5$.

The nine treatments were each tested three times, using different settings for the three parameters of the modified Roth-Erev (MRE) model of equations (1)–(3) and (6): the scaling parameter $S(1)$, a recency parameter ϕ, and an experimentation parameter ϵ. For the first two tests, the parameter values were chosen to facilitate the emergence for each trader of a dominant price offer with a relatively large choice probability, by the final auction round in each run. The third test used the parameter values obtained by Erev and Roth (1998) by best overall fit of their RL algorithm (equations 1–4) to experimental data from twelve distinct types of games run with human subjects: $S(1) = 9.00$, $\phi = 0.10$, $\epsilon = 0.20$.

Under all treatments, the presence of active buyers and sellers reduces the ability of structurally disadvantaged traders to exercise strategic market power, that is, to use strategic pricing to overcome the structural market-power biases inherent in the discriminatory-pricing protocol. Moreover, traders' ability to exercise strategic market power is further limited by the threat of entry of extramarginal traders, as discussed in Section 4.3.2 above.

NPT (2001) obtained generally high market efficiency (defined as $EA = \frac{PBA+PSA}{PBCE+PSCE}$) under all treatments. Notably, as seen in the earlier study (above) by Nicolaisen et al. (2000), market efficiency was relatively low when the traders used the inappropriate form of social mimicry embodied in GA learning. The later results from NPT (2001) suggest that the market efficiency of double auctions operating under a discriminatory pricing rule is reliably high when buyers and sellers refrain from inappropriate learning behavior or bad judgment [Tesfatsion (2005)]. These results confirm Vriend's (2000) argument that market efficiency is not robust with respect to a switch from individual learning (here MRE) to social learning (here GA).

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, NPT were clearly focused on market design. The paper cited Bower and 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 and 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.[20]

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

Bower and 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 RL 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

[20]In a private communication, 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.

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. Bower and Bun 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.

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 and 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 RL algorithm, such as Erev and Roth (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).

4.3.4 Bunn and Oliveira (2003) help design a new wholesale market.

Bunn and 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 and 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 were 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 NPT (2001).

The authors used agent-based simulation in a coordination game to analyze the possibility of market power (structural or strategic, as measured by higher prices and profitability than competitive outcomes) being exercised in a competitive electricity market: the policy issue was to help answer the question of whether two specific generators could influence wholesale electricity prices.

They extended Bun and Oliveira's (2001) New Electricity Trading Arrangements simulation platform. Their agents can be modeled as having the capacity to learn, and represented generating companies (possibly owning several plants with different generation philosophies) and buyers in the wholesale market who then supply end-use consumers. Agents use a RL algorithm to improve their performance: each agent evaluates the profit earned, and then derives new policies to bid or offer, given its strategic objectives of profit maximization and market exposure.

The authors were not interested in whether a particular market design, or structure, resulted in a competitive equilibrium; rather, whether a particular player, by its conduct, finds it profitable to act (choosing its offer price and strategically withholding capacity) in order to increase wholesale electricity prices.

They derive a simplified analytical model of the market: two generators in a stylized discriminatory-price Bertrand game with capacity constraints, from which they derive several propositions, which are then tested in the simulation of a more realistic model of the electricity industry. They used the eight largest generators in the England and Wales electricity market in 2000, splitting each generator's capacity into three categories, based on the degree of flexibility and running times of each technology (nuclear, large coal and combined-cycle gas turbines, and the rest). The simulated industry had 80 gensets, owned by 24 generators, who sell power to 13 suppliers. Four daily demand profiles were used. After initial learning by the agents, they found that the evolution of prices settled by about 50 iterations (trading days), and results were averaged over the last 10 days (of 50).

They simulated six different strategies for one and (or) both of the generators whose behavior was under scrutiny, under six different scenarios, each of which was repeated twice, with small differences. Average prices of the six strategies (under the 12 simulations) were higher than the marginal costs (even with full capacity available). This indicated structural market power caused by the industry structure, exacerbated by strategic market power (such as deliberately withholding capacity).

In order to evaluate the capacity of the two generators to manipulate market prices through capacity withholding, they compared different simulations using t-statistics (for pooled samples), a result of the complexities introduced by multiple equilibria and the effects of agents' learning. The two can act as price makers, but only when they both simultaneously withdraw capacity from the market can they profit from price manipulation.

They argued that the agent-based simulation technique enabled substantial insights to be gained before the new wholesale electricity market was introduced, and enabled the modeling of complex adaptive behavior in an environment with possible multiple equilibria, with heterogeneous agents and price uncertainty.

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 seven 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: EPRI and the Lawrence Berkeley National Laboratory; Argonne National Laboratory; and Hewlett-Packard. [Koesrindartoto and Tesfatsion (2004) discuss 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. (2000)]. Although EMCAS is based on the RePast open-source agent-based simulation platform [Collier and Sallach (2001)] and uses GA learning for certain agents, it is a proprietary system. 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 (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 and Dayan (1992)], a version of reinforcement learning; and genetic classifier systems. SEPIA was hosted at Honeywell, and was under-written by EPRI. [See Amin (2002) for further discussion.]

EPRI has used agent-based models to explore market design: Entriken and Wan (2005) 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, 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 bid 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 Computer Trading and On-Line Markets

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 probably 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 the purpose at hand, parameterization of the design space of auctions is necessary to allow direct 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 CDA. 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 [Duffy (2005)]. 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.

5.0.1 "Evolutionary mechanism design" at Liverpool.

A group at the University of Liverpool have been developing techniques of what they dub "evolutionary mechanism design" to examine not just buyer and seller behavior, but auctioneer behavior too, that is, how the transaction price is (or might be) derived in double auctions. Specifically, they took the wholesale electricity market of NPT (2001) almost intact, with one change: they moved from a clearinghouse double auction to a CDA, using the open-source "4-heap" algorithm [Wurman et al. (1998)]. As a CDA, there was discriminatory pricing, and Myerson and Satterthwaite's (1983) impossibility theorem holds.

In the first of a series of papers, Phelps et al. (2002a) 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 identified two possible techniques for computer-aided auction design based on evolutionary computing: Koza's (1993) genetic programming (GP) and the MRE RL algorithm as formalized in equations (1)–(3) and (6) above.

The authors first used the same best-fit MRE parameters and the same 3 $\times$ 3 treatment of $RCON$ and $RCAP$ as in NPT (2001). They were able to replicate NPT's results for market power and for mean market efficiency (close to 100%). But market efficiency was more volatile than in NPT (2001), perhaps because of the change from clearinghouse to CDA.

The authors then switched to assuming that each trader used GP instead of MRE reinforcement learning to search for a pricing strategy. Each agent's fitness was a function of its profits. Separate populations allowed the emergence of collusive strategies between self-interested traders. Could high-efficiency outcomes be sustained in this model? The answer was no: After 2000 generations, market efficiency stabilized at the relatively low level of 74%.

The final section of the paper added a seventh population, that of auctioneers, again using GP to search a space of pricing rules that included both uniform-pricing and discriminatory-pricing versions of the k-double auction. The auctioneer's fitness was proportional to the total profits earned in the market.

The simulation results showed that the adaptive auction was able to significantly improve its mean EA: to 94.5% and stability after only 500 generations, with the same 3 $\times$ 3 treatment of $RCON$ and $RCAP$ as above. In each of the 9 cases the evolved pricing rule was a linear function of either b or a, the two prices, but not both. When $NS = NB$, the price is determined by a, suggesting that sellers control the market whatever the values of $RCAP$. They cited Curzon Price (1997).

In a succeeding paper, Phelps et al. (2002b) use an objective function which is a weighted sum of MPB, MPS, and EA, each suitably normalized. They restrict search of the mechanism design space to the question: What is the best k-double-auction rule? Are there alternatives that perform as well or better when agents play strategies derived from a cognitive model of strategic interacting: the MRE?

They first simulated the same wholesale electricity market for a range of k values, using stochastic sampling, and found that $k \approx 0.5$ gave good performances. Then they used GP to search the larger space of arbitrary pricing rules, from b and a prices in the CDA. They derived several pages of "completely impenetrable" Lisp-based arithmetical expressions, which only became clear when plotted: effectively the discriminatory-price k-CDA with $k = 0.5$, apart from a small variation when a is small, or $a = b$. So $k = 0.5$ is reasonable.

A third paper [Phelps et al. (2003)] extended the earlier work to examine the strategy-proofness of k. It found that $k = 0.5$ is close to strategy-proof. A fourth paper [Phelps et al. (2005)] uses a "heuristic-strategy" approach and replicator dynamics [Duffy (2005)] to compare the clearinghouse double auction with the CDA, in terms of strategy-proofness and EA efficiency. It concluded that although the CDA is, on average, slightly less efficient, it can handle higher flows of transactions.

To summarize the significance of these papers: Agent-based market models have used two kinds of learning: social evolutionary learning algorithms, such as Holland's GAs or Koza's GP; and versions of individual reinforcement learning, such as the Roth-Erev model and modifications. On the one hand, NPT (2001) argue that the social learning implicit in the GA together with the endogenous extramarginal/inframarginal decision militates against the emergence high market efficiency in agent-based models, while a version of Roth-Erev is sufficient for its emergence. On the other hand, Phelps et al. (2002a) 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.

Design of markets might occur with simultaneous "design" of trading agents, a line of research pursued with GA learning at Hewlett-Packard by Cliff[21] (2001, 2002a, 2002b, 2003a) on CDAs and by Byde (2002) on sealed-bid auctions. Two weakness of Cliff (2001) are that, one, it uses a single population for many heterogeneous agents, and, two, the fitness function selects only for globally desirable outcomes, not individually desirable ones. 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.

[21] On his web page, 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 (2002a), et seq.] See the chapter by Duffy (2005) for an extensive discussion of Zero-Intelligence traders.

5.0.2 Byde (2002) evolves a new form of sealed-bid single auction.

The emphasis of the mechanism-design research in this chapter has been almost exclusively on double auctions. Yet, the single (or monopolist) auction is also of great interest, especially the new, spectrum auction. Byde (2002) examines the design of the sealed-bid single auction, using automated agents as bidders. The agents learn via a GA, and the objective is to maximize seller revenue, while not ignoring buyers' von-Neumann-Morgenstern utilities under different designs. Each bidding agent's valuation of the item for sale is some function of the signals received by all bidders.

Byde defines a w-price auction as a generalization of first- and second-price auctions: let $w = (w_1, w_2, ..., w_n)$ be a vector of n non-negative real numbers. A w-price auction is a sealed-bid auction in which the highest bidder wins the item, and pays

$$\frac{\sum_{j=1}^{N} w_j bid_j}{\sum_{j=1}^{N} w_j}, \tag{7}$$

where N is the minimum of n and the number of bidders, and bid_1, bid_2, ... are the bids ordered from highest to lowest. Byde used the GA to examine a one-dimensional sub-space of w-price auctions: those of the type where the vector $w = (1 - w_2, w_2)$. When $w_2 = 0$, this is a standard first-price auction; when $w_2 = 1$, this is a second-price (Vickrey) auction; and when $0 < w_2 < 1$, the payout is $(1 - w_2)bid_1 + w_2 bid_2$, a non-standard sealed-bid auction.

The space of agent preferences and environmental variables searched allowed Byde to examine exceptions to the Revenue Equivalence Theorem [Milgrom (2004)]: variable numbers of bidders, risk preferences, correlated signals, and degrees of commonality of values. Using a GA, he simulated a population of bidding agents which bid as a function of the signal each received, and played the game many times with stochastic sampling. He noted that each agent's fitness is relative to other agents (although, with a single population, he was not strictly co-evolving agents), which can lead to strategic behavior, such as bidding above one's signal if low, in order to reduce the winner's surplus. The game was repeated, not once-off, modeling bidders who come to know each others' behaviors.

With risk-neutral bidders, independent signals, and a fixed number of bidders, Byde benchmarked the Revenue Equivalence Theorem: there is no seller revenue advantage to any particular w_2. With risk-averse agents cet. par., first-price ($w_2 = 0$) gave highest seller revenue; with correlated signals cet, par., second-price ($w_2 = 1$) gave highest. He then found that "under several classes of non-pathological conditions (e.g. bidders were risk-averse, and unaware of how many players they would face in a given auction), there existed sealed-bid mechanisms expected to return significantly higher revenue to the auctioneer than either the first- or second-price sealed-bid mechanisms," specifically a payout where $w_2 = 0.3$, or $= 0.7$ under other conditions. He noted that since agents' average expected utility seems insensitive to w_2, sellers could design sealed-bid auctions to maximize their revenue without much buyer resistance. Byde's paper directs the market engineer to a new family of designs for sealed-bid auctions, and a new way to examine their performance in silico, before committing to real-world construction.

6 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 ten years old, and the work on automated markets is even more recent. 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 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 modelers, in economics, engineering, and computer science.

In this chapter, we have discussed the meaning of market design, its challenges, and the use of agent-based simulation models to achieve it, examining in detail published research in two of the five designer markets we introduced in Section 1 above, as examples of design by simulation.

We have discussed, first, analyzing electricity markets; second, attempting to design such markets directly; and, third, designing new markets for on-line and automated transactions. We have also mentioned in passing design issues in financial markets. 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.

7 Acknowledgments

I acknowledge help in writing this chapter from the editors, and from Raimo Hämäläinen, Derek Bunn, Peter McBurney, Bob Wilson, Paul Klemperer, Simon Parsons, Enrico Gerding, Eddie Anderson, Thai Cau, Steve Phelps, Carol McCormack, Robert Wood, and my fellow contributors at the *Handbook* Workshop at the University of Michigan, May, 2004. Three anonymous referees were very helpful with their comments.

References

Amin, M. (2002), "Restructuring the electric enterprise: simulating the evolution of the electric power industry with intelligent adaptive agents," Chapter 3 in: A. Faruqui and K. Eakin, eds., *Market Analysis and Resource Management,* (Kluwer, Dordrecht).

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

Arifovic, J. (1994), "Genetic algorithm learning and the cobweb model," *Journal of Economic Dynamics and Control* 18: 3–28.

Arifovic, J. (2001), "Performance of rational and boundedly rational agents in a model with persistent exchange rate volatility," *Macroeconomic Dynamics* 5: 204–224.

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.

Arthur, W.B. (2005), "Out-of-equilibrium economics and agent-based modeling," this *Handbook.*

Arthur, W. B., J. Holland, B. LeBaron, R. Palmer and P. Tayler (1997), "Asset pricing under endogenous expectations in an artificial stock market," in: W. B. Arthur, S. Durlauf and D. Lane, eds., *The Economy as an Evolving Complex System II* (Addison-Wesley, Reading, Mass.) 15–44.

Audet, N., T. Gravelle and J. Yang (2002), "Alternative trading systems: does one shoe fit all?" Working Paper 2002–33 (Bank of Canada, Ottawa).

Axelrod, R. (1987), "The evolution of strategies in the iterated Prisoner's Dilemma," in: L. Davis, ed., *Genetic Algorithms and Simulated Annealing* (Pittman, London) 32–41.

Axelrod, R. (2005), "Agent-based modeling as a bridge between disciplines," this *Handbook.*

Bottazzi, G., G. Dosi and I. Rebesco (2003), "Institutional architectures and behavioural ecologies in the dynamics of financial markets: a preliminary investigation," Technical Report, Laboratory of Economics and Management, Sant'Anna School of Advanced Studies, Pisa, Italy.

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.

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

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): 561–592.

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.

Brenner, T. (2005), "Agent learning representation," this *Handbook.*

Bullard, J. and J. Duffy (1999), "Using genetic algorithms to model the evolution of heterogeneous beliefs," *Computational Economics* 13(1): 41–60.

Bunn, D.W. (2004), personal communication.

Bunn, D.W. and C.J. Day (1998), "Agent-based simulation of electric power pools: a comparison with the supply function equilibrium approach," in: *Technology's Critical Role in Energy and Environmental Markets*, Proceedings of the 19th Annual North American Conference of the United States Association for Energy Economics and the International Association for Energy Economics, 18–21 October 1998, Albuquerque, New Mexico, (IAEE/USAEE, Cleveland).

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.

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.

Burkhart, R., M. Askenazi and N. Minar (2000), "Swarm Release Documentation." Available as http://www.santafe.edu/projects/swarm/swarmdocs/set/set.html. Accessed 25 November 2004.

Bush, R.R. and F. Mosteller (1955), *Stochastic Models for Learning* (Wiley, New York).

Byde, A. (2002), "Applying evolutionary game theory to auction mechanism design," Hewlett-Packard Technical Report HPL-2002-321.

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.

Cau, T.D.H. and E.J. Anderson (2002), "A co-evolutionary approach to modeling the behavior of participants in competitive electricity markets," in: *Proceedings of the Power Engineering Society Summer Meeting* (IEEE Society Press, Piscataway, N.J.) 1534–1540.

Chan, N.T. and C. Shelton (2001), "An electronic market-maker," Artificial Intelligence Lab, M.I.T., AI Memo 2001-005, April.

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

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.

Cliff, D. (2002a) "Evolution of market mechanism through a continuous space of auction-types," in: *Proceedings of the 2002 Congress on Evolutionary Computation, (CEC '02)* Honolulu, (IEEE Society Press, Piscataway, N.J.) 2029–2034.

Cliff, D. (2002b), "Evolution of market mechanism through a continuous space of auction-types II: Two-sided auction mechanisms evolve in response to market shocks," Hewlett-Packard Technical Report HPL-2002-128.

Cliff, D. (2003a), "Explorations in evolutionary design of online auction market mechanisms," *Electronic Commerce Research and Applications* 2(2): 162–175.

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

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

Conlisk, J. (1996) "Why bounded rationality?" *Journal of Economic Literature* 34: 669–700.

Cramton, P. (2003), "Electricity market design: the good, the bad, and the ugly," in: *Proceedings of the 36th Hawaii International Conference on System Sciences* (IEEE Society Press, Piscataway, N.J.).

Cross, J.G. (1973), "A stochastic learning modle of economic behavior," *Quarterly Journal of Economics* 87: 239–266.

Cross, J.G. (1983), *A Theory of Adaptive Economic Behavior* (Cambridge University Press, Cambridge).

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

Das, R., J.E. Hanson, J.O. Kephart and G. Tesauro (2001), "Agent-human interactions in the continuous double auction," in: B. Nebel, ed., *Proceedings of the 17th International Joint Conferences on Artificial Intelligence (IJCAI), Seattle* (Morgan Kaufmann, San Francisco) 1169–1187.

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

Duffy, J. (2005), "Agent-based models and human-subject experiments," this *Handbook.*

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

Entriken, R. and S. Wan (2005), "Agent-based simulation of an Automatic Mitigation Procedure," in: *Proceedings of the 38th Hawaii International Conference on System Sciences* (IEEE Society Press, Piscataway, N.J.).

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

FERC (2003), "Notice of White Paper," U.S. Federal Energy Regulatory Commission Docket No. RM01-12-000, April 28.

Forrester, J.W. (1961), *Industrial Dynamics* (M.I.T. Press, Cambridge).

Fudenberg, D. and E. Maskin (1986), "The Folk Theorem in repeated games with discounting or incomplete information," *Econometrica* 54: 533–554.

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

Gjerstad, S. (2004), "The impact of bargaining pace in double auction dynamics," Department of Economics, University of Arizona.

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

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.

Hailu, A. and S. Schilizzi (2004), "Are auctions more efficient than fixed price schemes when bidders learn?" *Australian Journal of Management* 29: 147–168.

Hämäläinen R.P. (1996), "Agent-based modeling of the electricity distribution system," in: M.H. Hamza, ed., *Modelling, Identification and Control,* Proceedings the 15th International Association of Science and Technology for Development (IASTED) International Conference, February 19–21, Innsbruck, (ACTA Press, Calgary) 344–346.

Hämäläinen R.P. (2004), personal communication.

Hämäläinen R.P., E. Kettunen and H. Ehtamo (1997), "Game modelling and coordination processes for two-level multi-agent systems," in: M.H. Hamza, ed., *Modelling, Identification and Control,* Proceedings of the 16th IASTED International Conference, February 17–19, Innsbruck, (ACTA Press, Calgary) 234–240.

Hämäläinen R.P. and J. Parantainen (1995), "Load analysis by agent-based simulation of the electricity distribution system," in: *Proceedings of the 2nd International Federation of Automatic Control (IFAC) Symposium on Control of Power Plants and Power Systems* SIPOWER95, Cancun, Mexico, December 6-8, 1995 (Elsevier, Oxford) 213–217.

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.

Hoecker, J. (1998), "Keeping electric restructuring moving forward" (Feb. 3, 1998) (11th Annual Utility M&A Symposium, New York), quoted in: E.P. Kahn, "Numerical techniques for analyzing market power in electricity," *The Electricity Journal* 34–43, July.

Holland, J.H. and J.H. Miller (1991), "Artificial adaptive agents in economic theory," *American Economic Review* 81(2): 365–370.

Holland, J.H. (1992), *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (M.I.T. Press, Cambridge), 2nd ed.

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* (Morgan Kaufmann, San Francisco) 242–250.

Janssen, M.A. and E. Ostrom (2005), "Governing socio-ecological systems," this *Handbook.*

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.

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

Klemperer, P. (2002), "What really matters in auction design," *Journal of Economic Perspectives* 16(1): 169–189.

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

Knuth, D. E. (1979), *TeX and METAFONT: New Directions in Typesetting* (Digital Press, Bedford, 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).

Koza, J. R. (1993), *Genetic Programming: On the Programming of Computers by Means of Natural Selection* (M.I.T. Press, Cambridge).

Krishna, V. and V.C. Ramesh (1998), "Intelligent agents in negotiations in market games, Part 2, Application," *IEEE Transactions on Power Systems* 13(3): 1109–1114.

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

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

Lane, D., A. Kroujiline, V. Petrov and G. Sheblé (2000), "Electricity market power: marginal cost and relative capacity effects," in: A. Alzala, ed., *Proceedings of the 2000 Congress on Evolutionary Computation* (IEEE Society Press, Piscataway, N.J.) 1048–1054.

LeBaron, B. (2005), "Agent-based computational finance," this *Handbook.*

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

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

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

MacGill, I.F. and R.J. Kaye (1999), "Decentralized coordination of power system operation using dual evolutionary programming," *IEEE Transactions on Power Systems* 14(1): 112–119.

MacKie-Mason, J.K. and M.P. Wellman (2005), "Automated markets and trading agents," this *Handbook.*

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

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

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

Marks, R.E. (2003), "Models rule," Editorial, *Australian Journal of Management* 28(1): i–ii.

Marks, R.E., D.F. Midgley and L.G. Cooper (1995), "Adaptive behavior in an oligopoly," in: J. Biethahn and V. Nissen, eds., *Evolutionary Algorithms in Management Applications* (Springer, Berlin) 225–239.

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

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

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

Myerson, R.B. and M.A. Satterthwaite (1983), "Efficient mechanisms for bilateral trading," *Journal of Economic Theory* 29: 265–281.

Nicolaisen, J., M. Smith, V. Petrov and L. Tesfatsion (2000), "Concentration and capacity effects on electricity market power," in: A. Alzala, ed., *Proceedings of the 2000 Congress on Evolutionary Computation* (IEEE Society Press, Piscataway, N.J.) 1041–1047.

Nicolaisen, J., V. Petrov and L. Tesfatsion (2001), "Market power and efficiency in a computational electricity market with discriminatory double-auction pricing," *IEEE Transactions on Evolutionary Computation* 5(5): 504–523.

North, M.J. (2000), "SMART II: The Spot Market Agent Research Tool Version 2.0," in: *Proceedings of SwarmFest 2000* (Swarm Development Group, Logan, Utah) 9–13.

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.

North, M., G. Conzelmann, V. Koritarov, C. Macal, P. Thimmapuram and T. Veselka (2002), "E-Laboratories: agent-based modeling of electricity markets," in: *Proceedings of the 2002 American Power Conference* (PennWell, Tulsa, Okla.).

Phelps, S., P. McBurney, S. Parsons and E. Sklar (2002a), "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.

Phelps, S., S. Parsons, E. Sklar and P. McBurney (2002b), "Applying multi-objective evolutionary computing to auction mechanism design," University of Liverpool Computer Science Technical Report ULCS-02-031.

Phelps, S., P. McBurney, E. Sklar and S. Parsons (2003), "Using genetic programming to optimise pricing rules for a double auction market," in: *Proceedings of the Workshop on Agents for Electronic Commerce*, Pittsburgh, PA.

Phelps, S., S. Parsons and P. McBurney (2005), "Automated trading agents versus virtual humans: an evolutionary game-theoretic comparison of two double-auction market designs," in: P. Faratin and J.A. Rodriguez-Aguilar, eds., *Agent-Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems,* Lecture Notes in Computer Science, (Springer, Berlin).

Räsänen M., R.P. Hämäläinen and J. Ruusunen (1994), "Visual interactive modelling in electricity load analysis," in: M.H. Hamza, ed., *Modelling, Identification and Control,* Proceedings the 13th International Association of Science and Technology for Development (IASTED) International Conference, Grindelwald, Switzerland, Feb. 21–23, (ACTA Press, Calgary) 339–342.

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

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.

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

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. García-Jurado and S. Tijs, eds., *Game Practice: Contributions from Applied Game Theory* (Kluwer, Dordrecht) 7–18.

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

Roth, A.E. and I. Erev (1995), "Learning in extensive form games: experimental data and simple dynamic models in the intermediate term," *Games and Economic Behavior* 8: 848–881.

Rubinstein, A. (1997) *Modeling Bounded Rationality* (M.I.T. Press, Cambridge).

Sargent, T.J. (1993), *Bounded Rationality in Macroeconomics* (Oxford University Press, New York).

Satterthwaite, M.A. and S.R. Williams (1989), "Bilateral trade with the sealed bid k-double auction: existence and efficiency," *Journal of Economic Theory* 48: 107–133.

Satterthwaite, M.A. and S.R. Williams (1993), "The Bayesian theory of the k-double auction," in: D. Friedman and J. Rust, eds., *The Double Auction Market: Institutions, Theories, and Evidence* (Addison-Wesley, Reading, Mass.) 99–123.

Selten, R. (1998), "Features of experimentally observed bounded rationality," *European Economic Review* 42: 413–436.

Selten, R. and R. Stoecker (1986), "End behavior in sequences of finite Prisoner's Dilemma supergames," *Journal of Economic Behavior and Organization* 7: 47–70.

Simon, H. (1981), *The Sciences of the Artificial* (M.I.T. Press, Cambridge) 2nd ed.

Sweeney, J.L. (2002), *The California Electricity Crisis* (Hoover Institution Press, Stanford).

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

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.

Tesfatsion, L. (2002), "Agent-based computational economics: growing economies from the bottom up," *Artificial Life* 8(1): 55–82.

Tesfatsion, L. (2005), personal communication.

Thorndike, E.L. (1911), *Animal Intelligence: Experimental Studies* (Macmillan, New York).

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.

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

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.

Walia, V., A. Byde and D. Cliff (2003), "Evolving market design in zero-intelligence trader markets," in: *Proceedings of the IEEE International Conference on E-Commerce, 2003, (CEC '03)* (IEEE Society Press, Piscataway, N.J.) 157–164.

Walsh W.E., R. Das, G. Tesauro and J.O. Kephart (2002), "Analyzing complex strategic interactions in multi-agent systems," in: P.J. Gmytrasiwicz and S. Parsons, eds., *Game Theoretic and Decision Theoretic Agents*, American Association for Artificial Intelligence Technical Report WS-02-06 (AAAI Press, Menlo Park, Calif.) 109–118.

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

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.

Weibull, J.W. (1995), *Evolutionary Game Theory* (M.I.T. Press, Cambridge).

Wilson, R. (2002), "Architecture of power markets," *Econometrica* 70(4): 1299–1340.

Wood, R.E. (2005), personal communication.

Wurman, P.R., W.E. Walsh and M.P. Wellman (1998), "Flexible double auctions for electronic commerce: theory and implementation," *Decision Support Systems* 24: 17–27.

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.

Young, H.P. (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton University Press, Princeton).