

Market Power and Efficiency in a Computational Electricity Market With Discriminatory Double-Auction Pricing

James Nicolaisen, Valentin Petrov, and Leigh Tesfatsion

Abstract—This study reports experimental market power and efficiency outcomes for a computational wholesale electricity market operating in the short run under systematically varied concentration and capacity conditions. The pricing of electricity is determined by means of a clearinghouse double auction with discriminatory midpoint pricing. Buyers and sellers use a modified Roth–Erev individual reinforcement learning algorithm to determine their price and quantity offers in each auction round. It is shown that high market efficiency is generally attained and that market microstructure is strongly predictive for the relative market power of buyers and sellers, independently of the values set for the reinforcement learning parameters. Results are briefly compared against results from an earlier study in which buyers and sellers instead engage in social mimicry learning via genetic algorithms.

Index Terms—Agent-based computational economics, capacity, concentration, efficiency, genetic algorithm social learning, individual reinforcement learning, market power, repeated double auction, restructuring, wholesale electricity market.

I. INTRODUCTION

ANY ELECTRIC power industry must carry out three basic functions, regardless of its structure [1], [2]. First, it must produce electricity from existing capacity. Second, it must distribute this electricity to final consumers. Third, it must engage in longer run planning and investment for the production of new capacity.

Until recently, most electricity has been supplied by vertically integrated statutory monopolies operating either as public utilities or as regulated investor-owned utilities [3, Ch. 6]. The regulatory compact has generally decreed that utilities should provide enough generation capacity to ensure an adequate supply of electricity for all users in return for receiving a fair rate of return on their capacity.

In recent years, however, this regulatory compact has been widely eroded due to the growing irrationality of the rate structure and the public perception that current industry practices fail to provide the proper incentives to ensure market efficiency. In

consequence, calls have increasingly been made to restructure the electric power industry into a more competitive industry in which prices would better reflect true marginal costs. At the same time, concerns have been expressed that restructuring efforts might actually result in efficiency losses. Moreover, equity concerns have arisen regarding possible increased opportunities for some market participants to exercise market power, i.e., to secure a higher share of profits for themselves than would be possible under competitive market conditions [4].

To date, restructuring proposals for the electric power industry have focused primarily on the wholesale electricity market. In this market, electricity is produced by generating companies (“generators”) from existing capacity and sold either to other generators or to some form of energy service provider. The energy service providers subsequently resell the electricity to household, industry, or commercial users in a retail market.

Short-run production efficiency in a wholesale electricity market requires that current demand be met using the least costly mix of existing capacity. Any market mechanism proposed for the short-run efficient determination of trades in a wholesale electricity market must therefore address four tasks [1, pp. 11–12]. First, the buyers and sellers who would benefit from trade must be identified. Second, these buyers and sellers must be matched so as to maximize total gains to trade. Third, a specific price and quantity level must be determined for each matched buyer-seller pair. Fourth, trades between matched buyers and sellers must be carried out within the constraints of the electric power transmission grid.

One market mechanism currently under intense consideration for wholesale electricity markets in many parts of the world is a *clearinghouse double auction* [5]–[7]. In such an auction, wholesale buyers and sellers of electricity participate repeatedly in auction rounds. At the beginning of each auction round, the buyers and sellers submit price and quantity offers simultaneously to a clearinghouse. At the end of each auction round, the clearinghouse matches the price and quantity offers received during the round in accordance with publicly known protocols, subject to transmission grid constraints, and reports these matches back to the buyers and sellers. The prices set by the clearinghouse are either *discriminatory* (set individually for each matched buyer-seller pair) or *uniform* (set equal across all matched buyer-seller pairs). The practical advantage of such an auction mechanism is that its trading and settlement rules do not require the clearinghouse to know in advance the number of participants, their private costs, or their privately held beliefs and preferences.

Manuscript received August 27, 2000; revised April 28, 2001. This work was supported by the National Science Foundation under Grant NSF-900589 and by the Electric Power Research Institute and the Department of Defense under Grant EPRI/DOD-W0-833-01.

J. Nicolaisen is with the Electrical Engineering Department, Iowa State University, Ames, IA 50011 USA (e-mail: jamesnic@iastate.edu).

V. Petrov is with the Computer Engineering Department, Iowa State University, Ames, IA 50011 USA (e-mail: vpetrov@iastate.edu).

L. Tesfatsion is with the Economics Department, Iowa State University, Ames, IA 50011 USA (e-mail: tesfatsi@iastate.edu).

Publisher Item Identifier S 1089-778X(01)09041-5.

Ideally, the performance of any proposed auction mechanism should be understood prior to its actual implementation. The difficulty for wholesale electricity markets is that these markets generally comprise small numbers of buyers and sellers with differentiated costs and capacities who interact repeatedly over time. The buyers and sellers may, thus, have an incentive to “game” an auction mechanism, i.e., to behave opportunistically within the limits set by the auction protocol in an attempt to increase their individual gains to trade. In particular, buyers and sellers may have an incentive to submit price offers that misrepresent their true willingness to pay or their true marginal costs and to submit quantity offers that misrepresent their true capacities.

For example, it is well known that implicit collusion problems can arise in uniform-price auctions for multiple units of a homogeneous good such as electricity [8], [9]. In uniform-price auctions, the marginally matched buyer and seller determine the price for all units, and auction participants may be able to collude tacitly to move this price in their favor. Klemperer [9, p. 4] notes that it was partly to avoid such problems that electricity regulators in the U.K. recently proposed a set of New Electricity Trading Arrangements for the U.K. Under these arrangements, an exchange market followed by a discriminatory-price auction would replace the existing uniform-price auction.

Implicit collusion is more difficult in discriminatory-price auctions. Nevertheless, auction participants may still have an incentive to engage in opportunistic behavior with regard to their price and quantity offers. Moreover, as found in [10] in the context of a one-sided auction for electricity generators, a discriminatory auction may permit large generators with many generating plants to have informational advantages over smaller generators.

Consequently, while it is highly desirable to predict in advance the market power and efficiency implications of proposed new auction protocols for wholesale electricity markets, the complexity of these markets makes it difficult to do so using standard analytical tools or human-subject laboratory experiments.¹ Empirical study is also difficult since relevant data is scarce. This suggests a potentially useful role for computational experiments.

This study constructs an agent-based computational model of a wholesale electricity market that can be used as a laboratory for systematic experimentation.² We use this laboratory to investigate market power and efficiency outcomes for a short-run wholesale electricity market with double-auction pricing and with buyers and sellers who continually update their price offers on the basis of past profit experiences. We consider how the relative market power of the buyers and sellers varies in response to changes in concentration and capacity when auction

prices are determined by means of a discriminatory pricing rule. We also consider the implications of this discriminatory pricing rule for short-run market efficiency.

One special concern of this study is the development of conceptual tools that permit market power effects due to market structure to be distinguished from market power effects due to buyer and seller learning. We focus particularly on the degree to which the discriminatory pricing rule induces structural versus behavioral market power effects.

Another special concern of this study is the testing of an empirically based representation for individual learning. Our electricity buyers and sellers are assumed to learn in accordance with a modified version³ of a reinforcement learning algorithm due to Roth and Erev [16], [17]. The attractiveness of the Roth–Erev learning algorithm is that its form embodies the most salient regularities observed in the decision-making behavior of human subjects across a wide variety of multiagent experimental games.

Our main experimental finding is that structural biases are inherent in discriminatory pricing rules and that these biases are highly predictive for relative market power outcomes. Although high market efficiency is generally attained, the buyers or sellers who are less favored in terms of market power under the discriminatory pricing rule are not able to overcome this structural market power bias through learning. This is the case even though the less favored agents have the potential to gain positive market power through appropriate strategic pricing. The symmetric nature of the double auction, which requires both sides of the market to submit price offers simultaneously, appears to prevent either buyers or sellers from successfully learning to gain a relative market power advantage through strategic price offers.

Section II presents our computational electricity market framework, including a detailed description of the auction mechanism, the learning algorithm used by traders to determine their price offers, and the calculation of benchmark competitive market outcomes. In Section III, we explain the experimental design of our study in terms of both tested hypotheses and tested parameter values. Section IV reports our basic experimental findings and Section V provides a detailed discussion of these findings. The concluding Section VI summarizes our key findings and discusses the relationship of these findings to work by other authors, in particular [18] on the relative efficiency effects of market structure versus learned behavior.

II. ELECTRICITY MARKET FRAMEWORK

A. Overview

Our computational electricity market incorporates several features anticipated to be key aspects of short-run wholesale electricity markets operating under restructuring. Small numbers of buyers and sellers submit price offers repeatedly to a clearinghouse double auction in an attempt to maximize their

¹As discussed in [7] and [11]–[13], researchers studying auctions by means of analytical tools and human-subject laboratory experiments have focused largely on simpler auction contexts in which the scope for opportunistic behavior is limited, e.g., single-round single-unit auctions in which the participants have extensive common knowledge. Rust *et al.* [14] is an important exception.

²Other researchers who have undertaken agent-based computational studies of wholesale electricity markets with double-auction pricing include [6] and [15]. Various resources on agent-based computational economics (ACE) in general, including surveys, an annotated syllabus of readings, software, and pointers to individual researchers and research groups, can be found at the ACE website at <http://www.econ.iastate.edu/tesfatsi/ace.htm>.

³In preliminary versions of this study (ISU Economic Report no. 52, August 2000, revised February 2001), it is stated that the original Roth–Erev algorithm is used. It was subsequently determined, however, that our code implementation actually included a “small” modification of this learning algorithm. As will be clarified below, the resulting modified learning algorithm turns out to have important advantages over the original Roth–Erev learning algorithm in the current double-auction setting.

profits. The buyers and sellers have multi-unit capacities and differentiated revenues and costs, private information that other traders cannot observe directly. Moreover, each buyer and seller continuously updates its price offers on the basis of its past profit experiences in a manner that permits the exploitation of profit opportunities arising from the pricing behavior of other traders.

More precisely, our computational electricity market models the short-run wholesale trading of electricity by traders attached to an electric power transmission grid. The transmission grid is a fully connected graph with traders as the nodes and transmission lines as the edges. Each trader is assigned a maximum amount of electricity (*capacity*) that it can buy or sell in each auction round as well as a certain available transmission capability (ATC) with respect to each other trader. Traders with electricity to buy are referred to as *buyers* and traders with electricity to sell are referred to as *sellers*.

The following parameter values are specified for each buyer: capacity in megawatt-hours, (constant) marginal revenue per megawatt-hour purchased and resold in a secondary retail market, and fixed costs. Also, the following parameter values are specified for each seller: capacity in megawatt-hours, (constant) marginal cost per megawatt-hour generated, and fixed costs. These parameter values are private to each trader.

The buyers and sellers trade electricity repeatedly in a discriminatory-price double auction run by an independent clearinghouse, henceforth referred to as a *discriminatory auction*. The goal of each buyer and seller is to maximize its own profits.

The discriminatory auction is performed in rounds. In each round, the buyers and sellers simultaneously submit *bids* (offers to buy) and *asks* (offers to sell) to the clearinghouse. Each bid and ask consists of a single price-quantity pair. The linearity assumed for the traders' revenue and cost functions, together with the discriminatory auction protocol, ensures that the profit-maximizing quantity offered by each trader is simply its capacity quantity. As detailed more fully in Section II-B, the clearinghouse matches these bids and asks, using as its criterion the maximization of perceived total profit, and communicates these matches back to the buyers and sellers.

At the end of the auction round, the matched buyers and sellers carry out their assigned trades and record their profit outcomes. They then use these profit outcomes to determine their price offers for the next auction round.

B. Auction Round Implementation

A single auction round proceeds as follows. First, each trader selects a feasible price offer in accordance with the trader's current "choice probabilities"—the determination of these choice probabilities is explained in Section II-E. Each trader submits this price offer to the clearinghouse along with a quantity offer equal to the trader's capacity.⁴ The clearinghouse then separately sorts the buyers and sellers by their price offers in descending and ascending order, respectively.⁵

⁴Recall from Section II-A that the profit-maximizing quantity offer for each trader in each auction round is simply its capacity quantity.

⁵Before each (bubble) sort, the ordering of the traders is randomized to avoid unintended incumbency effects in cases where some buyers or sellers make identical price offers.

TABLE I
BUYER-SELLER MATCHING ILLUSTRATION

Sellers	Buyers
\$4 / 20 MWh;	\$9 / 10 MWh
\$5 / 10 MWh;	\$8 / 10 MWh
\$6 / 10 MWh;	\$7 / 10 MWh

Matches: (1-1) for 10 MWh at Unit Price \$6.50/MWh;
(1-2) for 10 MWh at Unit Price \$6/MWh;
(2-3) for 10 MWh at Unit Price \$6/MWh;
Seller 3 Not Matched.

The buyer with the highest bid price is first matched with the seller with the lowest ask price. The unit price for the contract is set at the midpoint of the bid-ask spread. If there is nonzero ATC between the buyer and the seller, then the buyer is matched with the seller for an amount of electricity calculated as the minimum of three amounts: the ATC, the capacity of the buyer, and the capacity of the seller. Thus, if the ATC is 5 MWh, the buyer's capacity is 10 MWh and the seller's capacity is 20 MWh, then the contract is made for 5 MWh because this is the maximum quantity that the power grid can support. The carryover amount to buy or sell is then calculated and the next pair is matched in similar fashion.

Table I gives an example of a matching outcome for a market comprising three buyers and three sellers in which the ATC between any paired buyer and seller is assumed to be 10 MWh.

At the end of the auction round, each trader implements its auction-assigned trades and obtains a profit outcome. Each trader then uses this profit outcome to calculate updated choice probabilities for selecting among its feasible price offers in the next auction round. This updating is accomplished by means of a modified version of a reinforcement learning algorithm due to Roth and Erev [16], [17]. The latter algorithm will be motivated and illustrated prior to introducing the modified version.

C. Motivation for the Roth-Erev Algorithm

In a series of studies, Roth and Erev [16], [17] have sought to understand how people learn individually to behave in games with multiple strategically interacting players. To this end, they have developed a three-parameter reinforcement learning algorithm, hereafter referred to as the *RE algorithm*.

The basic intuition underlying any reinforcement learning algorithm is that the tendency to implement an action should be strengthened (reinforced) if it produces favorable results and weakened if it produces unfavorable results [19]. Roth and Erev take this *law of effect* principle, widely accepted in the psychological learning literature, as the basic starting point in their search for a robust model of individual learning. In addition, they argue for an additional learning principle, also adhered to widely in the psychological learning literature, which they refer to as the *power law of practice*. The latter principle asserts that learning curves tend to be initially steep, after which they flatten out.

Psychologists generally have focused on individual learning in "games against nature" for which there is only one decision maker. In contrast, Roth and Erev are interested in individual learning in strategic environments with multiple decision makers. Roth and Erev argue that, in such contexts, the law of

effect and the power law of practice fail to account sufficiently for the observed responsiveness of decision makers to other decision makers in their choice environments.

Based on extensive observations of individual learning in multiagent games, Roth and Erev argue for two additional learning principles that help to capture learning responsiveness, which they refer to as *experimentation* and the *recency* (or *forgetting*) *effect*. The former principle asserts that, not only are choices that were successful in the past more likely to be employed in the future, but similar choices will be employed more often as well. The latter principle asserts that recent experience generally plays a larger role than past experience in determining behavior.

The RE algorithm incorporates each of these four learning principles to some degree. Roth and Erev show that this algorithm is able to track successfully the observed intermediate-term behavior of human subjects over a wide variety of multiagent repeated games with unique equilibria achievable using stage-game strategies.

D. Form of the Roth–Erev Algorithm

The three parameters characterizing the RE algorithm are a *scaling parameter* $s(1)$, a *recency parameter* r , and an *experimentation parameter* e . The implementation of the RE algorithm will now be illustrated for a group of buyers and sellers participating in a double auction. For simplicity, each buyer and seller is assumed to learn in accordance with an RE algorithm characterized by the same three values for these parameters.

The feasible price offer domain for each buyer and seller is approximated by a discrete grid consisting of K feasible actions (bid or ask prices) k , where K is the same for each trader. At the beginning of the first auction round 1, each trader j assigns an equal *propensity* $q_{jk}(1)$ to each of its feasible actions k , given by $q_{jk}(1) = s(1)X/K$, where X is the average profit that buyers and sellers can achieve in any given auction round.

Moreover, each trader j assigns an equal *choice probability* $p_{jk}(1)$ to each of its feasible actions k , given by $p_{jk}(1) = 1/K$. Each trader j then probabilistically selects a feasible action k' to submit to the clearinghouse in accordance with its current choice probabilities. On the basis of all received bids and asks, the clearinghouse determines buyer-seller matches. It then communicates these matches back to the traders along with a quantity amount and a midpoint price for each match. Each trader j then implements its assigned trades and records the total profits $R(j, k', 1)$ that it gained from this trading activity.

Now suppose that trader j is at the end of the n th auction round for arbitrary positive n and that in the n th auction round trader j has submitted a feasible action k' to the clearinghouse and achieved total profits $R(j, k', n)$ from its resulting auction-directed trading activity. Trader j then updates its existing action propensities $q_{jk}(n)$ on the basis of its newly earned profits, as follows. Given any feasible action k , the propensity $q_{jk}(n+1)$ for choosing k in the next auction round $n+1$ is determined as

$$q_{jk}(n+1) = (1-r)q_{jk}(n) + E(j, k, k', n, K, e)$$

where

- r value of the recency parameter;
- e value of the experimentation parameter;
- $E(\cdot)$ update function reflecting the experience gained from past trading activity.

The recency parameter r slowly reduces the importance of past experience, thus implementing the recency effect. The update function $E(\cdot)$ takes the form

$$E(j, k, k', n, K, e) = \begin{cases} R(j, k', n)(1-e), & k = k' \\ R(j, k', n) \frac{e}{K-1}, & k \neq k' \end{cases}$$

The selected action k' is, thus, reinforced or discouraged on the basis of the profits $R(j, k', n)$ earned subsequent to this selection, but some propensity to experiment among all other feasible actions k is also retained. Thus, $E(\cdot)$ is an implementation of the experimentation principle.

Given the updated propensities $q_{jk}(n+1)$ for auction round $n+1$, trader j 's updated choice probabilities $p_{jk}(n+1)$ for selecting among its feasible actions k in auction round $n+1$ take the form

$$p_{jk}(n+1) = \frac{q_{jk}(n+1)}{\sum_{m=1}^K q_{jm}(n+1)}.$$

In summary, Roth–Erev traders solve a myopic stimulus-response problem of the following form. Given this profit outcome, what price should I next choose? They do not engage in any explicit look-ahead reasoning, e.g., if I choose this price now, how will this affect the price choices of my rivals in the future?

E. Modified Roth–Erev Algorithm

The RE algorithm outlined in Section II-D has two drawbacks: 1) parameter degeneracy and 2) no probability updating in response to zero profits.⁶

First, the updating of the choice probabilities is slow if e is set close to $[K-1]/K$ and ceases entirely if e is set equal to $[K-1]/K$. Consequently, care must be taken in specifying values for e and K .

Second, a much more substantial difficulty in a double-auction context is that each trader only updates its choice probabilities in response to *nonzero* profit outcomes. A zero-profit outcome leaves a trader's choice probabilities unchanged because each of the trader's current propensity values is shrunk to the same degree. In a double auction, traders must learn to make price offers for which bids exceed asks in order for matching (hence, positive profits) to occur at all. An absence of probability updating in response to zero profits can therefore result in a substantial loss of market efficiency as traders struggle to learn how to make profitable price offers.

A simple modification of the RE algorithm addresses both of these problems while still maintaining consistency with the learning principles embodied in the original RE algorithm.

⁶Kroujiline pointed out the parameter degeneracy problem to us. Koesrindartoto alerted us about the zero-profit updating problem in a particularly compelling way. In his own double-auction experiments with Roth–Erev learners participating in 1000 auction rounds, persistent market inefficiency arose for certain parameter specifications because the choice probabilities associated with various zero-profit (nonmatching) price offers remained at persistently positive levels.

Specifically, we replace the update function $E(\cdot)$ in the original RE algorithm with the following modified update function:

$$ME(j, k, k', n, K, e) = \begin{cases} R(j, k', n)(1 - e), & k = k' \\ q_{jk}(n) \frac{e}{K-1}, & k \neq k' \end{cases}$$

This modification essentially introduces different values for the recency parameter r for selected versus nonselected actions while at the same time omitting the profit term in the updating equation for propensities corresponding to nonselected actions. In particular, the effect is to reduce the magnitude of the recency parameter for nonselected actions from r to $r^* = (r - e/[K - 1])$. Clearly, degeneracy no longer occurs for $e = [K - 1]/K$, but how does this modification also ameliorate the zero-profit updating problem?

Note that the shrinkage induced by $[1 - r]$ in the propensity value for the selected action is now larger than the shrinkage induced by $[1 - r^*]$ in the propensity values for nonselected actions. Consider, then, what happens when a zero-profit outcome results from a selected action k' . All propensities are shrunk, but the propensity corresponding to k' undergoes the most shrinkage. Consequently, in the next auction round, the choice probabilities for the nonselected actions will increase relative to the choice probability for k' , encouraging the trader to move away from the action that resulted in zero profits.

On the other hand, suppose the selected action k' results in a positive profit outcome. Then the positive profit reinforcement in the propensity updating equation for k' will tend to outweigh the larger shrinkage and hence to induce a relative increase in the updated choice probability for this action in the next auction round.

In summary, when the update function $E(\cdot)$ in the RE algorithm is replaced with the modified update function $ME(\cdot)$, the zero-profit updating problem is ameliorated. 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.

In the current study, it will be assumed that electricity buyers and sellers adaptively update their price offers in accordance with this *modified* RE algorithm, hereafter referred to as the *MRE algorithm*.

F. Competitive Equilibrium Calculation

The main objective of this study is to determine market power and efficiency outcomes by comparing the profits that buyers and sellers of electricity obtain in a discriminatory auction against the profits they would obtain under competitive equilibrium. This section explains the meaning and calculation of competitive equilibrium.

A competitive equilibrium in a market for a positively valued good is a (positive) unit price P , a total quantity supplied $Q_S(P)$, and a total quantity demanded $Q_D(P)$ such that $Q_S(P) = Q_D(P)$. That is, the total quantity supplied must equal the total quantity demanded.

The notation $Q_S(P)$ and $Q_D(P)$ indicates that these supply and demand quantities depend on the price P of the good. How is this dependence determined?

The total supply at each given price is simply the sum of the quantities of good that each seller plans to sell at that price. Letting $q_i(P)$ denote how much of the good seller i plans to sell at each price P

$$Q_S(P) = \sum_i q_i(P).$$

Similarly, the total demand at each given price is the sum of the quantities of good desired by each buyer of the good at that price. Letting $q_j(P)$ denote how much of the good buyer j plans to buy at each price P

$$Q_D(P) = \sum_j q_j(P).$$

Note that the supplies and demands of the individual sellers and buyers are represented as functions of the market price P . This dependence comes from the assumption that these individual supplies and demands are the solutions of competitive profit maximization problems, i.e., profit maximization problems in which the traders are assumed to take the market price P as given.

Specifically, for the electricity model at hand, the competitive profit maximization problem for each seller i takes the following form:

$$\text{Max}_{q_i} Pq_i - \alpha_i q_i \text{ s.t. } 0 \leq q_i \leq CS_i.$$

The marginal cost parameter α_i denotes how much it costs seller i to generate each megawatt-hour of electricity and the capacity parameter CS_i denotes an upper bound on the amount that seller i can generate in any one auction round. The solution to this maximization problem is

$$q_i(P) = \begin{cases} CS_i, & \text{if } P > \alpha_i \\ [0, CS_i], & \text{if } P = \alpha_i \\ 0, & \text{if } P < \alpha_i \end{cases}.$$

Note that seller i is assumed to take the market price P as given and, hence, as exogenous to its profit maximization problem. The assumption in the competitive model is that the sellers believe that their quantity choices have no effect on the market price. Consequently, if they try to sell electricity at a price above the market price, they will sell nothing; buyers will buy electricity from the sellers with the lowest price. If they try to sell electricity at a price below the market price, they will succeed in selling all they can generate, but they could also do so at the market price and make higher profit. Thus, there is no incentive for the sellers to sell at any price other than the market price.

Similarly, for the electricity model at hand, the competitive profit maximization problem for each buyer j takes the following form:

$$\text{Max}_{q_j} r_j q_j - Pq_j \text{ s.t. } 0 \leq q_j \leq CB_j.$$

Here, r_j represents the marginal revenue received by buyer j for each megawatt-hour of electricity that buyer j resells in a secondary retail electricity market and CB_j is an upper bound

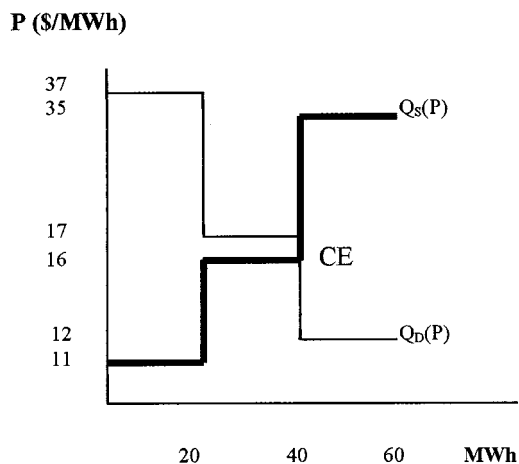


Fig. 1. Competitive equilibrium for a three-buyer three-seller model. Each trader has the same capacity (20 MWh).

on how much electricity buyer j can resell in any one auction round. The solution to this profit maximization problem is

$$q_j(P) = \begin{cases} 0 & \text{if } P > r_j \\ [0, CB_j] & \text{if } P = r_j \\ CB_j & \text{if } P < r_j \end{cases}.$$

The buyers are assumed to believe that their quantity choices have no effect on the market price P , so this price is taken as an exogenous parameter in their profit maximization problems.

A competitive equilibrium is said to occur at any price P that equates $Q_S(P)$ and $Q_D(P)$. It is possible that infinitely many competitive equilibria exist. An example of this is shown in Fig. 1. In this case, there are three buyers and three sellers, each with a capacity of 20 MWh. The competitive equilibrium is located where the supply and demand functions intersect, which happens to be along the vertical line segment labeled CE between \$16/MWh and \$17/MWh. Every point on this vertical line segment is a competitive equilibrium. For concreteness, we always take the competitive price to be the midpoint of all possible competitive prices. Hence, in the current example we would take the competitive price to be \$16.50/MWh, the average of the highest and lowest possible competitive prices \$17/MWh and \$16/MWh at the competitive equilibrium quantity 40 MWh.

It is important to recognize that the buyers and sellers in our computational electricity market do not actually solve the competitive profit maximization problems presented above. Rather, these profit maximization problems are used as zero-market-power benchmarks against which our experimental auction outcomes can be compared.

III. EXPERIMENTAL DESIGN

A. Tested Hypotheses

For simplicity, it is assumed that all buyers in the computational electricity market are energy service providers and all sellers are generators, implying that generators do not sell to other generators. Let N_B denote the number of buyers and let

N_S denote the number of sellers. The relative concentration (RCON) of the market is then defined to be

$$\text{RCON} = \frac{N_S}{N_B}.$$

Let CB denote the maximum amount of electricity that each buyer can resell in a retail market and let CS denote the maximum amount of electricity that each seller can generate, both measured in megawatt-hours. Then, the relative capacity (RCAP) of the electricity market is defined as

$$\text{RCAP} = \frac{N_B \cdot \text{CB}}{N_S \cdot \text{CS}}.$$

Let PBCE denote the profits⁷ that buyers would obtain in competitive equilibrium and let PBA denote the profits that buyers instead obtain when prices and quantities are determined in the discriminatory auction. Then, the market power of buyers (MPB) is defined as

$$\text{MPB} = \frac{\text{PBA} - \text{PBCE}}{\text{PBCE}}.$$

If the buyers can exert control over the price of electricity in the auction, i.e., if the buyers can exercise market power, then they should be able to raise their profits above their competitive profit level and MPB should be positive.

Similarly, let PSCE denote the profits that sellers would obtain in competitive equilibrium and let PSA denote the profits that sellers instead obtain in the discriminatory auction. Then, the market power of sellers (MPS) is defined as

$$\text{MPS} = \frac{\text{PSA} - \text{PSCE}}{\text{PSCE}}.$$

If the sellers can exert control over the price of electricity in the auction, i.e., if the sellers can exercise market power, then they should be able to raise their profits above their competitive profit level and MPS should be positive.

The level of total profits achieved by all buyers and sellers in the computational electricity market attains its maximum value in competitive equilibrium by construction. The market power indexes MPB and MPS measure the extent to which the profit levels achieved separately by buyers and sellers in the discriminatory auction differ from the profit levels they would achieve in competitive equilibrium. This redistribution of profits between buyers and sellers may come at the expense of smaller total profits for all buyers and sellers combined.

For example, suppose some *inframarginal* trader (i.e., some trader who would match in competitive equilibrium) engages in opportunistic auction behavior, either underbidding its true willingness to pay or asking more than its true marginal cost. In this case, it could happen that an *extramarginal* trader (i.e., a trader who would fail to match in competitive equilibrium) would be able to match in the discriminatory auction by submitting a higher bid or lower ask than the opportunistic trader, leaving the latter unmatched.

To check for possible inefficiencies arising under the discriminatory auction, we introduce the following efficiency measure.

⁷For expositional simplicity, we refer to the net earnings of the buyers as profits, ignoring all further downstream retail activities by these agents.

The *efficiency* of the market operating under the auction protocol is defined to be the ratio EA of total auction profits to total profits in competitive equilibrium, measured in percentage terms. That is, using previously introduced notation, we define

$$EA = \frac{PBA + PSA}{PBCE + PSCE} \times 100.$$

We test the market power and efficiency implications of discriminatory auction pricing for our computational wholesale electricity market in the form of three hypotheses.

H1—As RCAP Increases, MPB Decreases While MPS Increases, All Else Equal: Intuitively, buyers should have a harder time exercising market power when there is excess demand capacity, i.e., when the maximum amount of electricity that the buyers want to purchase exceeds the amount that the sellers are able to generate. Thus, one might expect the MPB market power index for buyers to decrease and the MPS market power index for sellers to increase with increases in RCAP, all else equal.

H2—As RCON Decreases, MPB Decreases While MPS Increases, All Else Equal: Intuitively, sellers should have an easier time exercising market power as electricity generation becomes concentrated in the hands of fewer sellers per buyer. Thus, the MPB market power index for buyers should decrease and the MPS market power index for sellers should increase as RCON decreases, all else equal. For example, hypothesis H2 is consistent with the claim in [20, p. 952] that the market power exercised by generators in the British electricity spot market would have been reduced substantially if the industry had been subdivided into five generators rather than two.

H3—Most Potential Gains to Trade are Exhausted, i.e., EA is Close to 100%. Any Unrealized Profitable Trades are Those Offering the Smallest Gains: Hypothesis H3 conjectures that the level of total profits achieved by buyers and sellers in the discriminatory auction will be close to the level of total profits achieved under competitive equilibrium, which is the maximum possible level. Consequently, the conjecture is that the size of the total profit pie under the discriminatory auction is essentially independent of the division of this pie among the buyers and sellers and hence essentially independent of any auction-induced market power effects. Hypothesis H3 is consistent with the high efficiency found in a wide variety of human-subject experiments with double auctions [11].

B. Tested Parameter Values

The experimentally tested values for the number N_B of buyers, the number N_S of sellers, the capacity CB of each buyer, and the capacity CS of each seller are given in Table II. The capacities of the buyers and sellers are representative of typical generation and demanded loads.

The capacities for the buyers and sellers are selected to provide the following three test ratios for the relative capacity measure RCAP: 1:2, 1:1, and 2:1. All buyers are assumed to have identical capacities and similarly for all sellers. For simplicity, in this first experimental study, the ATC between each buyer and seller is set at 100 MWh to ensure that the ATC is not a binding

TABLE II
TESTED PARAMETER VALUES

	RCAP		
	1/2	1	2
2	$N_S = 6$ $N_B = 3$ CS=10 CB=10	$N_S = 6$ $N_B = 3$ CS = 10 CB = 20	$N_S = 6$ $N_B = 3$ CS = 10 CB = 40
RCON 1	$N_S = 3$ $N_B = 3$ CS = 20 CB=10	$N_S = 3$ $N_B = 3$ CS = 10 CB = 10	$N_S = 3$ $N_B = 3$ CS = 10 CB = 20
1/2	$N_S = 3$ $N_B = 6$ CS = 40 CB=10	$N_S = 3$ $N_B = 6$ CS = 20 CB = 10	$N_S = 3$ $N_B = 6$ CS = 10 CB = 10

TABLE III
LINEAR REVENUE AND COST CURVES

Buyers	Marginal Revenue
1	\$37/MWh
2	\$17/MWh
3	\$12/MWh
4	\$37/MWh
5	\$17/MWh
6	\$12/MWh
Sellers	Marginal Cost
1	\$35/MWh
2	\$16/MWh
3	\$11/MWh
4	\$35/MWh
5	\$16/MWh
6	\$11/MWh

constraint on any buyer-seller match under these capacity specifications.⁸

Buyers and sellers are assumed to have linear revenue and cost functions subject to capacity constraints, so that their marginal revenues and marginal costs are constant over their quantity choices up to capacity. The cost functions specified for the sellers are scaled linear approximations of the cost functions of actual generating units.

Table III shows the specification for marginal revenue (marginal cost) for each buyer (seller) in the experiments reported below. The fixed costs of the buyers and sellers are set to zero for a simpler model. For a seller, this could be representative of a generator already up and running (i.e., synchronized to the transmission grid) and waiting for a match in the auction to connect to the system and deliver electricity.

The marginal costs of the sellers are chosen to cover three types of operating costs: expensive, medium, and cheap. These three types might be representative of an older generation unit,

⁸As stressed in [21], determining the effects of ATC constraints on market power is a subtle issue, since ATC constraints make it more likely that congestion can be induced strategically for market power advantages. This topic will be taken up in future studies.

an older unit that has been updated, and a new unit, or of different types of fuel usage. Note from Table III that, when all six sellers are simulated, two of each type are included to model the competition between similar companies. The buyers' marginal revenues are similar to the marginal costs of the sellers, but with enough of a difference to yield a competitive equilibrium profit. This assures the existence of a competitive equilibrium price, which is then used to calculate the benchmark profit levels for market power and efficiency.

Buyers and sellers are not permitted to submit bid or ask prices to the auction that would definitely result in negative profits if accepted. To implement this rationality postulate, the set of feasible bid price offers for each buyer is specified to be the interval $[\text{MR} - \$40/\text{MWh}, \text{MR}]$, where MR denotes the buyer's true (constant) marginal revenue. Also, the set of feasible ask price offers for each seller is specified to be interval $[\text{MC}, \text{MC} + \$40/\text{MWh}]$, where MC denotes the seller's true (constant) marginal cost. The lower bound $\text{MR} - \$40/\text{MWh}$ is low enough to encompass all possible ask prices by sellers and the upper bound $\text{MC} + \$40/\text{MWh}$ is high enough to encompass all possible bid prices by buyers.

To check the sensitivity of the market power and efficiency outcomes to the specific values set for the parameters characterizing the MRE reinforcement learning algorithm, the nine RCAP/RCON configurations in Table II are tested three times using three different settings for these parameter values.

Recall from Section II-E that the MRE algorithm is characterized by three parameters: a scaling parameter $s(1)$, a recency parameter r , and an experimentation parameter e . Erev and Roth [18, p. 864] note that a good fit to their experimental data covering twelve distinct types of human-subject games was obtained for all values of these parameters lying in the following ranges: $0 < s(1) < 1000$, $0 < r < 0.20$, and $0.02 < e < 0.30$.

In the first two tests for Table II reported below, the parameter values for the MRE algorithm are calibrated 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. In the first test, each run consists of 1000 auction rounds and in the second test, each run consists of 10 000 auction rounds. The parameter calibration was accomplished in two stages, as follows.

First, given the number of auction rounds per run, the density of the price offers within each trader's feasible price offer range was specified to help ensure an adequate sampling frequency for each possible offer. For the 1000 auction rounds per run case, $K = 30$ possible price offers were randomly selected within each feasible price offer range, implying that each trader could in principle sample each price 33 times during each run. The average profit X achievable in any auction round for this case was set at $X = 15,000$. For the 10 000 auction rounds per run case, $K = 100$ possible price offers were randomly selected within each feasible price offer range, implying that each trader could in principle sample each price 100 times during the course of each run. A value $X = 50,000$ was set for this case.

Second, by direct search, the values for the three MRE algorithm parameters $s(1)$, r , and e were calibrated until the bid or ask price histogram for each of the traders became single peaked by the final auction round in each run. The calibrated parameter

values found for the 1000 auction round case were $s(1) = 1.00$, $r = 0.04$, and $e = 0.97$. The calibrated parameter values found for the 10 000 auction round case were $s(1) = 1.00$, $r = 0.02$, and $e = 0.99$.⁹ For both cases, the calibrated values for $s(1)$ and r fall within the Erev–Roth good fit ranges presented above, whereas the calibrated values for e do not.

In the third test for Table II, the scaling parameter $s(1)$, the recency parameter r , and the experimentation parameter e for the MRE algorithm are instead set equal to the values obtained by Erev and Roth [18, p. 863] by a best overall fit of the RE algorithm to experimental data from twelve distinct types of games run with human subjects. These values are $s(1) = 9.00$, $r = 0.10$, and $e = 0.20$. The MRE algorithm with the latter parameter values is referred to below as the *best fit MRE algorithm*.

IV. EXPERIMENTAL RESULTS

Tables IV–VI report aggregate and individual market power outcomes and efficiency outcomes for the three distinct learning specifications outlined in Section III-B. Specifically, in the first learning specification, each trader is assumed to use the calibrated MRE algorithm with each run consisting of 1000 auction rounds. In the second learning specification, each trader is assumed to use the calibrated MRE algorithm with each run consisting of 10 000 auction rounds. In the third learning specification, each trader is assumed to use the best fit MRE algorithm with each run consisting of 1000 auction rounds.

Each cell in each table corresponds to a unique RCAP/RCON configuration in parallel to Table II. For each table cell, the auction was run 100 times using 100 different seeds for the pseudorandom number generator. For each run, the profit levels attained in the final auction round by buyers as a whole and by sellers as a whole, as well as by individual buyers and sellers, were calculated and compared against competitive profit levels to obtain aggregate and individual MPB and MPS market power indexes. In addition, for each run, the value for the market efficiency measure EA in the last auction round was calculated and recorded.

The means and standard deviations of the aggregate and individual MPB and MPS market power indexes were then calculated across all 100 runs for each table cell. The aggregate results are given at the top of each table cell and the results for individual buyers and sellers are listed underneath. A mean market power outcome with a positive or negative sign is marked with an asterisk if it is substantially different from zero in the sense that the indicated sign does not change when the outcome is either increased or decreased by one standard deviation. Finally, the mean and standard deviation for the market efficiency measure EA calculated across all 100 runs are given at the bottom of each table cell.

As explained in Section III-A, hypothesis H1 predicts that the MPB market power index for buyers should decrease while the MPS market power index for sellers should increase, all

⁹It is interesting to note that under the original RE algorithm, no updating of choice probabilities would occur for this 10 000 auction round case since $e = [K - 1]/K = 0.99$. In contrast, as will be seen below, the MRE algorithm results in good learning and very high market efficiency.

TABLE IV
EXPERIMENTAL MARKET POWER AND EFFICIENCY OUTCOMES FOR THE CALIBRATED MRE ALGORITHM WITH 1000 AUCTION ROUNDS AND PARAMETER VALUES
 $s(1) = 1.00$, $r = 0.04$, AND $e = 0.97$

	1/2		Relative Capacity 1		2	
	MP	StdDev	MP	StdDev	MP	StdDev
2	All Buyers:	-0.27 (0.18)	All Buyers:	-0.23 * (0.17)	All Buyers:	-0.06 (0.28)
	All Sellers:	0.46 (0.88)	All Sellers:	0.29 (0.48)	All Sellers:	-0.24 (0.33)
	Buyer[1]:	-0.24 (0.26)	Buyer[1]:	-0.21* (0.19)	Buyer[1]:	-0.06 (0.28)
	Buyer[2]:	-0.68* (0.50)	Buyer[2]:	-0.87 (0.96)	Buyer[2]:	ZP (0.00)
	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)
	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)
	Seller[2]:	ZP (0.00)	Seller[2]:	1.75 (5.46)	Seller[2]:	-0.24 (0.46)
	Seller[3]:	0.37 (1.97)	Seller[3]:	0.17 (0.85)	Seller[3]:	-0.24 (0.37)
	Seller[4]:	ZP (0.00)	Seller[4]:	ZP (0.00)	Seller[4]:	ZP (0.00)
	Seller[5]:	ZP (0.00)	Seller[5]:	1.39 (4.71)	Seller[5]:	-0.22 (0.44)
Seller[6]:	0.54 (1.32)	Seller[6]:	0.19 (0.80)	Seller[6]:	-0.25 (0.37)	
Efficiency:	96.01 (0.08)	Efficiency:	96.30 (0.11)	Efficiency:	77.60 (0.15)	
Relative Concentration 1	All Buyers:	-0.37* (0.20)	All Buyers:	-0.26* (0.17)	All Buyers:	-0.13 (0.37)
	All Sellers:	0.55 (0.67)	All Sellers:	0.44 (0.56)	All Sellers:	-0.27 (0.37)
	Buyer[1]:	-0.33* (0.20)	Buyer[1]:	-0.24* (0.18)	Buyer[1]:	-0.13 (0.37)
	Buyer[2]:	-0.73* (0.50)	Buyer[2]:	-1.00 (0.00)	Buyer[2]:	ZP (0.00)
	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)
	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)
	Seller[2]:	ZP (0.00)	Seller[2]:	2.37 (6.49)	Seller[2]:	-0.29 (0.51)
	Seller[3]:	0.40 (0.86)	Seller[3]:	0.27 (0.90)	Seller[3]:	-0.25 (0.37)
	Efficiency:	86.88 (0.18)	Efficiency:	96.48 (0.05)	Efficiency:	90.98 (0.24)
	1/2	All Buyers:	-0.33* (0.16)	All Buyers:	-0.25* (0.16)	All Buyers:
All Sellers:		0.55* (0.50)	All Sellers:	0.44 (0.44)	All Sellers:	-0.21 (0.25)
Buyer[1]:		-0.29* (0.19)	Buyer[1]:	-0.21* (0.17)	Buyer[1]:	0.01 (0.43)
Buyer[2]:		-0.68* (0.52)	Buyer[2]:	-1.00 (0.00)	Buyer[2]:	ZP (0.00)
Buyer[3]:		ZP (0.00)	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)
Buyer[4]:		-0.29* (0.17)	Buyer[4]:	-0.25* (0.24)	Buyer[4]:	-0.03 (0.44)
Buyer[5]:		-0.68* (0.49)	Buyer[5]:	-0.98* (0.20)	Buyer[5]:	ZP (0.00)
Buyer[6]:		ZP (0.00)	Buyer[6]:	ZP (0.00)	Buyer[6]:	ZP (0.00)
Seller[1]:		ZP (0.00)	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)
Seller[2]:		ZP (0.00)	Seller[2]:	0.77 (4.60)	Seller[2]:	-0.25 (0.37)
Seller[3]:	0.37 (0.76)	Seller[3]:	0.41 (0.72)	Seller[3]:	-0.18 (0.31)	
Efficiency:	85.53 (0.18)	Efficiency:	96.39 (0.04)	Efficiency:	96.55 (0.13)	

ZP indicates that zero profits were earned both in the auction and in competitive equilibrium.

else equal, in response to increases in RCAP. Looking at the market power outcomes reported in Tables IV–VI, however, it is seen that Hypothesis H1 is not supported under any of the three learning specifications. MPB actually tends to increase and MPS to decrease with increases in RCAP for each given RCON level, a direct contradiction of H1.

Also, hypothesis H2 predicts that the MPB market power index for buyers should decrease while the MPS market power index for sellers should increase, all else equal, in response to decreases in RCON. The latter measure is simply the ratio of the number of sellers to buyers. As seen in Tables IV–VI, however, the changes in buyer and seller market power levels in response to changes in RCON are small and unsystematic for each given level of RCAP, in contradiction to hypothesis H2. This is

particularly true for the experiments comprising 10 000 auction rounds per run reported in Table V.¹⁰

Finally, hypothesis H3 predicts that high market efficiency will be obtained. The mean and standard deviation outcomes for market efficiency EA reported at the bottom of each cell in Tables IV–VI reveal that high market efficiency is indeed generally obtained for each tested RCAP/RCON configuration. Consequently, hypothesis H3 is strongly supported. Note that the calibrated MRE algorithm with 10 000 auction rounds per run

¹⁰In human-subject experiments with actual electricity industry participants, Weiss [4, p. 1] finds that “increasing the number of sellers in a given (electricity) market does not necessarily reduce market power (of sellers), as suggested by most standard theory.” However, in Weiss’s study, this failure of H2 is due to the existence of ATC constraints that effectively create local monopolies at some power grid nodes, constraints that are not present in the current study.

TABLE V
EXPERIMENTAL MARKET POWER AND EFFICIENCY OUTCOMES FOR THE CALIBRATED MRE ALGORITHM WITH 10 000 AUCTION
ROUNDS AND PARAMETER VALUES $s(1) = 1.00$, $r = 0.02$, AND $e = 0.99$

	1/2		Relative Capacity 1		2	
	MP	StdDev	MP	StdDev	MP	StdDev
2	All Buyers:	-0.04 (0.07)	All Buyers:	-0.07 (0.26)	All Buyers:	-0.07 (0.24)
	All Sellers:	0.19 (0.32)	All Sellers:	0.21* (0.19)	All Sellers:	-0.06 (0.19)
	Buyer[1]:	-0.04 (0.06)	Buyer[1]:	-0.07* (0.05)	Buyer[1]:	-0.07 (0.24)
	Buyer[2]:	-0.04 (0.33)	Buyer[2]:	-0.30 (0.47)	Buyer[2]:	ZP (0.00)
	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)
	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)
	Seller[2]:	ZP (0.00)	Seller[2]:	-0.15 (0.79)	Seller[2]:	-0.06 (0.24)
	Seller[3]:	0.23 (0.44)	Seller[3]:	0.26* (0.22)	Seller[3]:	-0.06 (0.17)
	Seller[4]:	ZP (0.00)	Seller[4]:	ZP (0.00)	Seller[4]:	ZP (0.00)
	Seller[5]:	ZP (0.00)	Seller[5]:	-0.30 (0.63)	Seller[5]:	-0.06 (0.25)
Seller[6]:	0.14 (0.36)	Seller[6]:	0.24* (0.21)	Seller[6]:	-0.06 (0.17)	
Efficiency:	100.00 (0.00)	Efficiency:	99.49 (0.01)	Efficiency:	100.00 (0.00)	
Relative Concentration 1	All Buyers:	-0.16* (0.09)	All Buyers:	-0.08* (0.07)	All Buyers:	0.06 (0.24)
	All Sellers:	0.60* (0.38)	All Sellers:	0.22 (0.28)	All Sellers:	-0.05 (0.19)
	Buyer[1]:	-0.14* (0.07)	Buyer[1]:	-0.08* (0.07)	Buyer[1]:	0.06 (0.24)
	Buyer[2]:	-0.30 (0.38)	Buyer[2]:	-0.30 (0.58)	Buyer[2]:	ZP (0.00)
	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)
	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)
	Seller[2]:	ZP (0.00)	Seller[2]:	-0.05 (1.15)	Seller[2]:	-0.05 (0.24)
	Seller[3]:	0.60* (0.38)	Seller[3]:	0.25 (0.32)	Seller[3]:	-0.04 (0.16)
	Efficiency:	94.13 (0.09)	Efficiency:	99.66 (0.01)	Efficiency:	100.00 (0.00)
	1/2	All Buyers:	-0.14* (0.07)	All Buyers:	-0.06* (0.05)	All Buyers:
All Sellers:		0.59* (0.36)	All Sellers:	0.20* (0.19)	All Sellers:	-0.08 (0.16)
Buyer[1]:		-0.14* (0.06)	Buyer[1]:	-0.06 (0.06)	Buyer[1]:	0.10 (0.20)
Buyer[2]:		-0.24 (0.36)	Buyer[2]:	-0.31 (0.60)	Buyer[2]:	ZP (0.00)
Buyer[3]:		ZP (0.00)	Buyer[3]:	ZP (0.00)	Buyer[3]:	ZP (0.00)
Buyer[4]:		-0.12* (0.06)	Buyer[4]:	-0.06 (0.06)	Buyer[4]:	0.10 (0.20)
Buyer[5]:		-0.23 (0.34)	Buyer[5]:	-0.27 (0.64)	Buyer[5]:	ZP (0.00)
Buyer[6]:		ZP (0.00)	Buyer[6]:	ZP (0.00)	Buyer[6]:	ZP (0.00)
Seller[1]:		ZP (0.00)	Seller[1]:	ZP (0.00)	Seller[1]:	ZP (0.00)
Seller[2]:		ZP (0.00)	Seller[2]:	ZP (0.00)	Seller[2]:	-0.10 (0.20)
Seller[3]:	0.59* (0.36)	Seller[3]:	0.20* (0.19)	Seller[3]:	-0.07 (0.14)	
Efficiency:	95.22 (0.09)	Efficiency:	99.56 (0.01)	Efficiency:	100.00 (0.00)	

ZP indicates that zero profits were earned both in the auction and in competitive equilibrium.

strictly dominates the other two learning specifications with regard to market efficiency. That is, the efficiency outcomes reported in each table cell for Table V are as high or higher than the efficiency outcomes reported in the corresponding table cells for Tables IV and VI.

V. DISCUSSION

What explains the failure of the market power hypotheses H1 and H2 in the current electricity market context? On the surface, this failure seems to contradict basic economic intuition. Also, why is the market efficiency hypothesis H3 so strongly supported?

A careful case-by-case examination of the microevents underlying the outcomes reported in Tables IV–VI goes a long way toward dispelling the mystery. Briefly, one sees that the aggregate measures RCAP and RCON are simply too crude to reflect well the opportunities for exercising market power that individual buyers and sellers actually face. To understand the latter, the market microstructure must be carefully examined. In particular, as stressed in [18], it is important to distinguish between market outcomes that are due to market microstructure and market outcomes that are due to learned behavior.

Two different definitions for market power will next be given that permit the separate identification of market power due to structural causes and market power due to learning.

TABLE VI
EXPERIMENTAL MARKET POWER AND EFFICIENCY OUTCOMES FOR THE BEST FIT MRE ALGORITHM WITH 1000 AUCTION ROUNDS AND PARAMETER VALUES
 $s(1) = 9.00, r = 0.10, \text{ AND } e = 0.20$

	Relative Capacity								
	1/2		1		2				
	MP	StdDev	MP	StdDev	MP	StdDev			
2	All Buyers:	-0.13*	(0.09)	All Buyers:	-0.15*	(0.09)	All Buyers:	0.10	(0.30)
	All Sellers:	0.55*	(0.38)	All Sellers:	0.38*	(0.33)	All Sellers:	-0.10	(0.25)
	Buyer[1]:	-0.12*	(0.08)	Buyer[1]:	-0.13*	(0.10)	Buyer[1]:	0.10	(0.30)
	Buyer[2]:	-0.20	(0.40)	Buyer[2]:	-0.75*	(0.33)	Buyer[2]:	ZP	(0.00)
	Buyer[3]:	ZP	(0.00)	Buyer[3]:	ZP	(0.00)	Buyer[3]:	ZP	(0.00)
	Seller[1]:	ZP	(0.00)	Seller[1]:	ZP	(0.00)	Seller[1]:	ZP	(0.00)
	Seller[2]:	ZP	(0.00)	Seller[2]:	-0.50	(1.34)	Seller[2]:	-0.12	(0.34)
	Seller[3]:	0.54	(0.63)	Seller[3]:	0.45*	(0.40)	Seller[3]:	-0.10	(0.22)
	Seller[4]:	ZP	(0.00)	Seller[4]:	ZP	(0.00)	Seller[4]:	ZP	(0.00)
	Seller[5]:	ZP	(0.00)	Seller[5]:	-0.42	(1.67)	Seller[5]:	-0.08	(0.36)
Seller[6]:	0.55	(0.60)	Seller[6]:	0.46*	(0.41)	Seller[6]:	-0.09	(0.24)	
Efficiency:	99.81	(0.02)	Efficiency:	96.30	(0.05)	Efficiency:	99.88	(0.06)	
Relative Concentration 1	All Buyers:	-0.22*	(0.12)	All Buyers:	-0.13*	(0.10)	All Buyers:	0.13	(0.33)
	All Sellers:	0.80*	(0.53)	All Sellers:	0.28	(0.35)	All Sellers:	-0.10	(0.26)
	Buyer[1]:	-0.21*	(0.11)	Buyer[1]:	-0.11*	(0.10)	Buyer[1]:	0.13	(0.33)
	Buyer[2]:	-0.31	(0.44)	Buyer[2]:	-0.80*	(0.40)	Buyer[2]:	ZP	(0.00)
	Buyer[3]:	ZP	(0.00)	Buyer[3]:	ZP	(0.00)	Buyer[3]:	ZP	(0.00)
	Seller[1]:	ZP	(0.00)	Seller[1]:	ZP	(0.00)	Seller[1]:	ZP	(0.00)
	Seller[2]:	ZP	(0.00)	Seller[2]:	-0.37	(1.89)	Seller[2]:	-0.10	(0.34)
	Seller[3]:	0.76*	(0.63)	Seller[3]:	0.34	(0.45)	Seller[3]:	-0.11	(0.24)
	Efficiency:	92.13	(0.09)	Efficiency:	94.59	(0.07)	Efficiency:	100.00	(0.00)
	1/2	All Buyers:	-0.21*	(0.12)	All Buyers:	-0.14*	(0.08)	All Buyers:	0.09
All Sellers:		0.67*	(0.46)	All Sellers:	0.30	(0.31)	All Sellers:	-0.07	(0.19)
Buyer[1]:		-0.18*	(0.12)	Buyer[1]:	-0.14*	(0.10)	Buyer[1]:	0.09	(0.27)
Buyer[2]:		-0.37	(0.47)	Buyer[2]:	-0.77*	(0.44)	Buyer[2]:	ZP	(0.00)
Buyer[3]:		ZP	(0.00)	Buyer[3]:	ZP	(0.00)	Buyer[3]:	ZP	(0.00)
Buyer[4]:		-0.20*	(0.11)	Buyer[4]:	-0.11	(0.11)	Buyer[4]:	0.10	(0.25)
Buyer[5]:		-0.38	(0.47)	Buyer[5]:	-0.73*	(0.46)	Buyer[5]:	ZP	(0.00)
Buyer[6]:		ZP	(0.00)	Buyer[6]:	ZP	(0.00)	Buyer[6]:	ZP	(0.00)
Seller[1]:		ZP	(0.00)	Seller[1]:	ZP	(0.00)	Seller[1]:	ZP	(0.00)
Seller[2]:		ZP	(0.00)	Seller[2]:	0.14	(2.69)	Seller[2]:	-0.08	(0.27)
Seller[3]:	0.63*	(0.55)	Seller[3]:	0.32	(0.48)	Seller[3]:	-0.07	(0.17)	
Efficiency:	91.84	(0.09)	Efficiency:	94.24	(0.07)	Efficiency:	100.00	(0.00)	

ZP indicates that zero profits were earned both in the auction and in competitive equilibrium.

A. Structural Versus Strategic Market Power

First suppose that no trader misrepresents its true reservation price, i.e., suppose each buyer bids its true marginal revenue and each seller asks its true marginal cost. Under the discriminatory midpoint pricing rule, the exact relative positioning of the resulting "true" market demand and supply curves can still confer market power on some buyers and not on others and on some sellers and not on others. Call this *structural market power*.

In the current movement to restructure the electricity industry, all participants are well aware that the choice of auction protocol can substantially affect their relative profitability. Thus, a consideration of the structural market power allocated to different

market participants under alternative auction protocols would presumably be of major interest.

Second, by engaging in unilateral misrepresentation of their true reservation prices, some buyers and/or some sellers may have potentially exercisable market power in addition to (or even in the absence of) structural market power. Call this (*unilateral*) *strategic market power*.

B. Structural Market Power Outcomes

Analytically derived structural market power outcomes are presented in Table VII for the nine tested RCAP/RCON config-

TABLE VII
ANALYTICALLY DERIVED STRUCTURAL MARKET POWER OUTCOMES

	1/2	Relative Capacity 1	2
2	All Buyers: -0.14 All Sellers: 0.56	All Buyers: -0.12 All Sellers: 0.45	All Buyers: 0.04 All Sellers: -0.05
	Buyer[1]: -0.43 Seller[3]: 1.67 Seller[6]: 1.67	Buyer[1]: -0.37 Seller[3]: 1.36 Seller[6]: 1.36	Buyer[1]: 0.12 Seller[3]: -0.16 Seller[6]: -0.16
	All Others: 0.00	All Others: 0.00	All Others: 0.00
Relative Concentration 1	All Buyers: -0.14 All Sellers: 0.56	All Buyers: -0.12 All Sellers: 0.45	All Buyers: 0.04 All Sellers: -0.05
	Buyer[1]: -0.43 Seller[3]: 1.67	Buyer[1]: -0.37 Seller[3]: 1.36	Buyer[1]: 0.12 Seller[3]: -0.16
	All Others: 0.00	All Others: 0.00	All Others: 0.00
1/2	All Buyers: -0.14 All Sellers: 0.56	All Buyers: -0.12 All Sellers: 0.45	All Buyers: 0.04 All Sellers: -0.05
	Buyer[1]: -0.43 Buyer[4]: -0.43 Seller[3]: 1.67	Buyer[1]: -0.37 Buyer[4]: -0.37 Seller[3]: 1.36	Buyer[1]: 0.12 Buyer[4]: 0.12 Seller[3]: -0.16
	All Others: 0.00	All Others: 0.00	All Others: 0.00

urations in Table II.¹¹ The outcomes in Table VII reveal three interesting regularities.

First, buyers have negative structural market power and sellers have positive structural market power for six of the nine tested RCAP/RCON configurations, namely, for the six cells in the first two columns in Table VII. Consequently, in the current electricity market context, it is generally the buyers who are structurally disadvantaged in the auction with regard to market power.

Second, for each level of RCON, the mean structural MPB across all buyers increases as RCAP increases and the mean structural MPS across all sellers decreases as RCAP increases, in direct contradiction to the market power hypothesis H1.

¹¹In cell (1,1) of Table VII, the inframarginal Sellers 3 and 6 have the same marginal cost, but are matched with different buyers at different prices depending on their order of selection. In cell (3, 3), the inframarginal Buyers 1 and 4 have the same marginal revenue, but are matched with different sellers at different prices depending on their order of selection. The structural market power levels reported for these traders are their *expected* market power levels under the assumption that they are randomly ordered for matching purposes.

Third, for each level of RCAP, the mean structural MPB across all buyers and the mean structural MPS across all sellers are invariant to changes in RCON. This directly contradicts the market power hypothesis H2.

Comparing the experimental market power outcomes reported in Tables IV–VI against the analytically derived structural market power outcomes reported in Table VII, one sees that the experimentally determined market power outcomes closely track the structural market power outcomes. Consequently, in the current electricity market context, market microstructure is strongly predictive for observed market power outcomes.

C. Strategic Market Power Outcomes

What about strategic market power? Call a trader *inframarginal* if it would engage in a positive amount of trade in competitive equilibrium and *extramarginal* otherwise. As will be clarified in Section V-E below, every inframarginal buyer

and seller in each of the nine tested RCAP/RCON configurations in Table II has positive strategic market power under the discriminatory auction protocol. That is, assuming all other inframarginal traders bid or ask their true reservation prices, the remaining inframarginal trader can always increase its profits above what it would obtain if it bid or asked its true reservation price. If the trader is a buyer, it can accomplish this by suitably lowering its bid price below its true marginal revenue. If the trader is a seller, it can accomplish this by suitably raising its ask price above its true marginal cost. Under the discriminatory midpoint-pricing rule, this opportunistic behavior would then move the auction price in a direction favorable to the trader.

Nevertheless, structurally disadvantaged traders never learn to effectively exercise strategic market power in our auction experiments. By construction, the exercise of strategic market power is measured by the discrepancies between the experimentally determined market power outcomes in Tables IV–VI and the analytically derived levels for structural market power given in Table VII. While there are some discrepancies in magnitudes, there are no instances in which a trader with negative structural market power attains a positive market power level in the auction. Moreover, instances in which a trader with a positive structural market power level attains a negative market power level in the auction are rare.

Specifically, the only sign discrepancies in mean market power for all buyers or all sellers are as follows: Table IV has two sign discrepancies [mean MPB for all buyers in cells (1, 3) and (2, 3)], Table V has one sign discrepancy [mean MPB for all buyers in cell (1, 3)], and Table VI has no sign discrepancies. In each case, the sign discrepancy in the mean MPB level for all buyers is due to a sign discrepancy occurring for the mean MPB level of a single buyer: namely, Buyer 1. Note, also, that these sign discrepancies are all to the disadvantage of Buyer 1, i.e., its realized auction market power is negative whereas its structural market power is positive. Finally, note that none of the mean MPB values showing a sign discrepancy is marked with an asterisk. This implies that the sign discrepancy disappears within one standard deviation of the reported mean observation.

These findings show that apart from the small number of sign discrepancy cases noted above, learning has no effect on the *relative* exercise of market power by buyers and sellers. When the discriminatory auction protocol gives greater structural market power to buyers, the buyers retain this relative market power advantage in the auction experiments and similarly for sellers. Indeed, when buyers attain a positive mean market power level in the auction, the mean market power level attained by sellers in the auction is negative and vice versa. As will be clarified in Section V-E, this “zero-sum game” finding reflects the high market efficiency levels attained in the auction experiments.

D. Efficiency Outcomes

Apart from roundoff and truncation error, the market efficiency measure EA defined in Section III-A equals 100% in any given auction round if and only if the set of active traders in the auction round coincides with the set of active traders in

competitive equilibrium.¹² This implies, in particular, that no extramarginal trader manages to trade in the auction and that no inframarginal trader fails to trade in the auction.

Comparing the mean and standard deviation outcomes reported in Tables IV–VI for EA, one sees that the highest mean market efficiency outcomes are uniformly attained in Table V. The latter table reports outcomes for the case in which all buyers and sellers use the calibrated MRE algorithm with 10 000 auction rounds per run. In each cell of Table V, the market efficiency measure EA attains a mean value of 94% or better. Nevertheless, generally high mean market efficiency outcomes are also reported in Tables IV and VI for distinctly different settings of the MRE algorithm parameters and with only 1000 auction rounds per run.

The overall implication of these generally high market efficiency levels is that the discriminatory auction essentially reduces to a zero-sum game. That is, total buyer and seller profits are approximately given by total competitive profits in each experiment and the key remaining issue is how these profits are redistributed among buyers and sellers as one switches from competitive equilibrium pricing to discriminatory auction pricing.

Profit distribution under the discriminatory auction is measured by market power. If total profit remains constant as one switches from competitive to auction pricing, then market power simply measures the manner in which the auction redistributes this constant total profit between buyers and sellers. In this case, apart from roundoff error, a positive attained market power level for one type of trader necessarily implies a negative attained market power level for the other. As noted in Section V-C, it is indeed seen in Tables IV–VI that table cells reporting a high mean market efficiency level also generally report a mean market power level for all buyers that is opposite in sign to the mean market power level for all sellers.

E. Microanalysis for Illustrative Cases

To better understand the underlying reasons for the findings reported in Section IV, a more detailed microanalysis will now be given for two of the tested RCAP/RCON combinations in Table II, namely, cell (3, 1) and cell (3, 2).

First, consider cell (3, 1) in Table II with RCAP = RCON = 1/2. The market for this case comprises six buyers, each with capacity 10 MWh, and three sellers, each with capacity 40 MWh. The marginal revenues and marginal costs for these buyers and sellers are listed in Table III. Making use of these capacity, revenue, and cost specifications, the true demand and supply curves can be constructed; these are depicted in Fig. 2. The competitive outcome based on these demand and supply curves is $Q = 40$ MWh and $P = \$14/\text{MWh}$.

The low RCAP value 1/2 for this case implies the existence of excess potential supply, which suggests that buyers should be

¹²This simple characterization for market efficiency has to be slightly qualified in the presence of marginal traders with identical marginal revenues or marginal costs who do not all end up trading in competitive equilibrium. In this case, market efficiency holds regardless of which of these marginal traders actually carries out the marginal competitive equilibrium trades. This indeterminacy is absent for the nine tested configurations in Table II.

avored. Yet the experimental market power outcomes reported for this case in cell (3, 1) in Tables IV–VI show that Seller 3 is the only trader in the discriminatory auction that is successfully able to attain a positive market power level on average. Why is this the case?

It will now be shown that Seller 3 is the only trader that has positive structural market power. On the other hand, all inframarginal buyers and sellers have countervailing strategic market power. This prevents any trader with negative structural market power under the auction protocol from being able to successfully exercise its strategic market power to the point that it attains a positive market power level.

Under the discriminatory auction protocol, it is the 80 MWh unused capacity of Sellers 1 and 2 that constitutes the excess potential supply reflected in $RCAP = 1/2$. These two sellers have relatively high marginal costs of \$35/MWh and \$16/MWh, respectively. If all traders bid and ask their true reservation prices in the auction, Sellers 1 and 2 will not obtain an auction match. Nevertheless, since they also fail to trade in competitive equilibrium, their profits under the auction are the same as under competitive equilibrium. This implies that their structural market power is zero even though their MPS indexes are not well defined.

Similarly, the relatively low marginal revenue \$12/MWh of Buyers 3 and 6 prevents any auction match for these buyers if all traders bid and ask their true reservation prices. Thus, since they also fail to trade in competitive equilibrium, their structural market power is zero even though their MPB indexes are not well defined. Seller 3 has a relatively low marginal cost of \$11/MWh and its total electricity supply capacity is 40 MWh. Buyers 1 and 4 have a relatively high marginal revenue (willingness to pay) of \$37/MWh and their total electricity demand capacity is 20 MWh. Under the discriminatory midpoint pricing rule, assuming all traders bid or ask their true reservation prices, Buyers 1 and 4 would purchase 20 MWh from Seller 3 at a price of \$24/MWh, a price that exceeds the competitive price of \$14/MWh. The profit of Seller 3 on this 20 MWh contract would, thus, be \$260, greater than its competitive profit \$60, and the profit of Buyers 1 and 4 would each be \$130, less than their competitive profit \$230.

Buyers 2 and 5 with marginal revenue \$17/MWh would then end up purchasing Seller 3’s remaining 20 MWh at the competitive price \$14/MWh, giving Seller 3 a (competitive) profit of \$60 and Buyers 2 and 5 each a (competitive) profit of \$30. The MPS structural market power index for Seller 3 is then calculated to be 1.67, the MPB structural market power index for Buyers 1 and 4 is -0.43 , and the MPB structural market power indexes for Buyers 2 and 5 are zero.

These observations imply that, contrary to the implications of hypothesis H1, excess potential supply (a low RCAP value) can fail to ensure that any buyer has positive structural market power in the discriminatory auction. The reason for this is that excess potential supply can come from the excess capacity of high-cost sellers who are extramarginal under competitive equilibrium. Assuming all traders bid or ask their true reservation prices, these sellers will not be able to obtain matches in the discriminatory auction. Consequently, the presence of these extra-marginal sellers confers neither advantage nor disadvantage on any other




Fig. 2. Cell (3, 1) true demand and supply curves for six buyers and three sellers with 10 MWh and 40 MWh capacities, respectively.

trader with regard to structural market power, yet their presence can change the value of RCAP. Similar observations apply for extramarginal buyers in the reverse case of excess potential demand.

What about strategic market power? Consider, once again, cell (3, 1) in Table II with true demand and supply curves depicted in Fig. 2. It turns out that Seller 3 and Buyers 1, 4, 2, and 5 all have some degree of strategic market power under the discriminatory auction protocol.

Specifically, by unilaterally misrepresenting their true willingness to pay by bidding below their true marginal revenue \$37/MWh, Buyers 1 and 4 can each increase their auction profits. Indeed, if either buyer were to bid \$12/MWh plus some small amount epsilon, hence, above the price \$12/MWh at which the extramarginal Buyers 3 and 6 would be able to match, they would obtain a profit close to \$250 on their 10 MWh purchase from Seller 3. The latter profit is greater than their \$230 competitive profits and much greater than the \$130 auction profits that they would earn by bidding their true marginal revenue. The MPB strategic market power indexes for Buyers 1 and 4 are approximately 0.09. Similarly, the MPB strategic market power indexes for Buyers 2 and 5 (who should also strategically bid \$12/MWh plus epsilon) are approximately 0.67.

On the other hand, Seller 3 also has strategic market power. Suppose Seller 3 unilaterally raises its ask price from its true marginal cost \$11/MWh up to \$16/MWh minus epsilon, but no higher (to prevent Seller 2 from matching). Then, the contract price of Seller 3 with Buyer 1 and Buyer 4 for 20 MWh under the discriminatory midpoint pricing rule would be close to \$26.50/MWh. This would give Seller 3 a profit of \$310, higher than the \$260 profits it would earn in the auction by asking its true marginal cost, and much higher than its \$60 competitive profits.

In addition, by asking \$16/MWh minus epsilon, Seller 3’s contract price with Buyers 2 and 5 for 20 MWh would be approximately \$16.50/MWh. This would give Seller 3 a profit of \$110, which is higher than the \$60 profits that Seller 3 would

earn on this 20 MWh contract either in the discriminatory auction with bidding and asking of true reservation prices or in competitive equilibrium. The total profit of Seller 3 would therefore be \$420. Hence, the MPS strategic market power index for Seller 3 is 2.50.

Thus, Seller 3 and Buyers 1, 2, 4, and 5 in cell (3, 1) all have positive strategic market power. As previously determined, however, only Seller 3 has positive structural market power. Examining the results reported for cell (3, 1) in Tables IV–VI, it is seen that none of the buyers succeeds in exercising its strategic market power to the point that it ends up with a positive MPB market power level in the auction. The problem for the buyers is that strategic market power is being exercised on both sides of the market. When ask prices are being raised by sellers at the same time that bid prices are being lowered by buyers, the tendency is for the midpoint price of each bid-ask spread to remain essentially the same.

An examination of the actual price offers submitted to the auction by Seller 3 and Buyers 1, 2, 4, and 5 in cell (3, 1) reveals this effect. Seller 3 learns to submit ask prices higher than its marginal cost and this exercise of strategic market power by Seller 3 foils the attempts by the buyers to lower their auction price by underbidding their true marginal revenues.

Specifically, as seen in Fig. 3, the average ask price of Seller 3 in the final auction round is roughly \$15/MWh regardless of the precise parameter values set for the MRE algorithm. Note from Table III that \$15/MWh exceeds Seller 3's true marginal cost of \$11/MWh and is just below the price \$16/MWh at which the extramarginal Seller 2 could feasibly enter the market.

On the other hand, on average, the bid prices of Buyer 1 and Buyer 4 in the final auction round are roughly \$24/MWh for each tested MRE algorithm specification. This level is lower than their true marginal revenue \$37/MWh, but also higher than the level \$12/MWh at which the extramarginal Buyers 3 and 6 could feasibly enter the market. Also, on average, the bid prices of Buyers 2 and 5 in each final auction round are roughly \$15/MWh, below their true marginal revenue of \$17/MWh, but high enough to prevent entry by Buyers 3 and 6.

Thus, all active traders in cell (3, 1) exercise strategic market power to some degree by asking higher than true marginal costs or bidding below true marginal revenues. Nevertheless, the net result of these countervailing forces is that Seller 3's structural market power advantage prevails.

Regarding market efficiency, bid and ask price data for cell (3, 1) reveal that all inframarginal traders learn to place their bids and asks within ranges that prevent the entry of extramarginal traders. Moreover, the buyers end up bidding above asks, on average, so that coordination failures are relatively infrequent. The result is high market efficiency, despite the fact that Seller 3 achieves a relatively high market power level.

Higher market efficiency is achieved for cell (3, 1) as the overall volatility of the bid and ask prices in the final auction round decreases. As indicated in Fig. 3 for Seller 3, the greatest overall volatility is observed for the calibrated MRE algorithm with 1000 auction rounds per run and the least overall volatility is observed for the calibrated MRE algorithm with 10 000 auction rounds per run. This decline in volatility is particularly marked for Buyers 2 and 5.

Next, consider cell (3, 2) in Table II with $RCAP = 1$ and $RCON = 1/2$. The market for this case comprises six buyers, each with capacity 10 MWh, and three sellers, each with capacity 20 MWh. The marginal revenues and marginal costs for these buyers and sellers are listed in Table III. Making use of these specifications for capacities, revenues, and costs, the true demand and supply curves can be constructed; these are depicted in Fig. 4. The competitive outcome based on these demand and supply curves is $Q = 40$ MWh and $P = \$16.50/\text{MWh}$.

An analysis of the results reported for cell (3, 2) in Tables IV–VI proceeds along lines similar to the foregoing analysis for cell (3, 1), with one interesting exception. As seen by comparing Fig. 4 with Fig. 2, Buyers 2 and 5 and Seller 2 in cell (3, 2) face a much greater challenge with regard to learning how to select their bid and ask prices than any inframarginal trader in cell (3, 1). This is because the price range in which they can successfully match is much narrower.

Indeed, under the calibrated MRE algorithm with 1000 auction rounds per run, it is seen in cell (3, 2) of Table IV that Buyer 2 fails to match at all and Buyer 5 matches only rarely. An examination of bid and ask data for cell (3, 2) in Table IV reveals that the bid prices of these buyers are extremely volatile and they almost never exceed the marginal cost of Seller 2. Consequently, this learning specification does not permit price discovery for these buyers.

Seller 2 in cell (3, 2) of Table IV manages to match frequently enough with Buyer 1 and Buyer 4 to sustain a positive average market power level, although with extremely high standard deviation. The average ask price of Seller 2 is roughly \$20/MWh, which is higher than the marginal revenue \$17/MWh of Buyers 2 and 5. The average ask price of Buyers 1 and 4 in cell (3, 2) of Table IV is roughly \$25/MWh. Interestingly, the average ask price of Seller 3 in cell (3, 2) of Table IV is only about \$16/MWh, which is below the marginal revenue of Buyers 2 and 5. Nevertheless, given the extreme volatility of these latter buyers' bids, matches with these buyers essentially never occur.

Buyers 2 and 5 in cell (3, 2) perform better under the MRE algorithm specifications reported in Tables V and VI. For example, bid data for Buyers 2 and 5 in cell (3, 2) of Table VI reveal that they manage to bid close to their true marginal revenue \$17/MWh in the final auction round in about one third of the 100 runs and this percentage improves for Table V.

In contrast, Seller 2 in cell (3, 2) performs worse in Tables V and VI, failing to match at all in Table V. In Table VI, the average ask price of Seller 2 in cell (3, 2) is only about \$16/MWh, close to its true marginal cost. This permits matches with Buyers 2 and 5, but with very little gains to trade. In contrast, the average ask price of Seller 3 in cell (3, 2) of Table VI is around \$15/MWh. Thus, as the seller submitting the lowest ask on average, Seller 3 is now matched most frequently with Buyers 1 and 4, thus crowding out Seller 2 and forcing Seller 2 to trade with Buyers 2 and 5 for much lower gain. The average ask price of Buyers 1 and 4 in cell (3, 2) of Table VI is about \$24/MWh.

Despite the increased challenge that some traders face to achieve matches in cell (3, 2), the coordination failures mostly involve marginal traders with small gains to trade. Consequently, high market efficiency is still achieved.

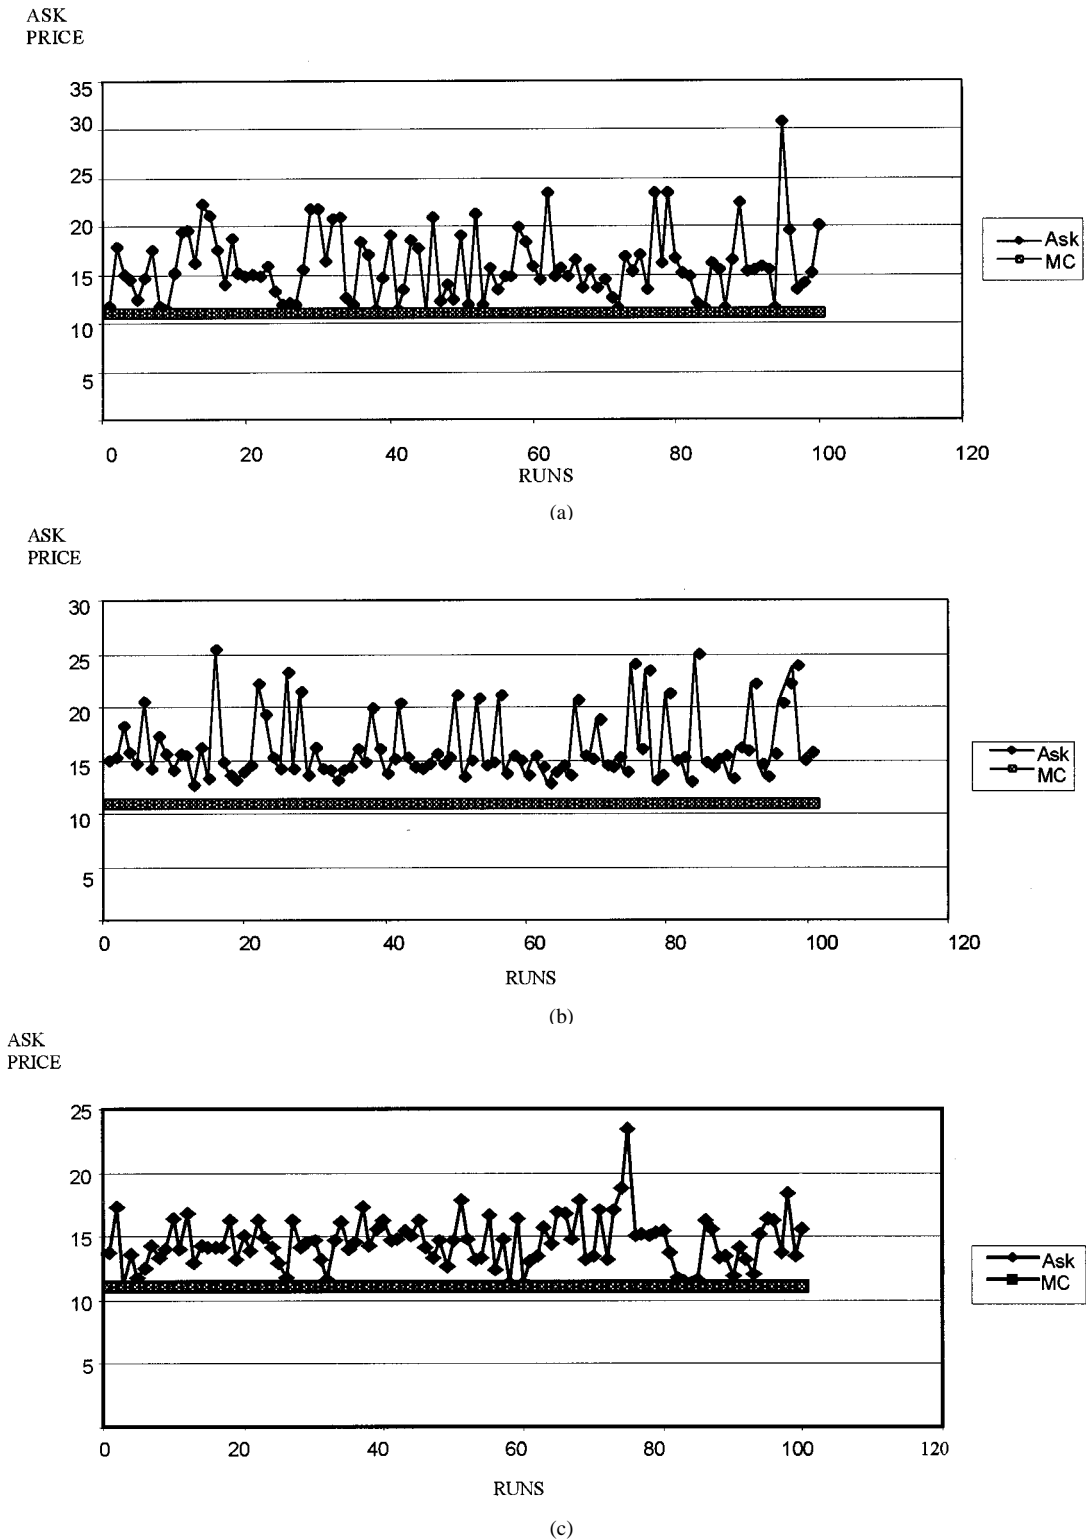


Fig. 3. Plot of Seller 3’s ask price in cell (3, 1) in the final generation across all 100 runs under three different MRE algorithm specifications. (a) Table IV, calibrated MRE algorithm with 1000 auction rounds per run. (b) Table V, calibrated MRE algorithm with 10 000 auction rounds per run. (c) Table VI, best fit MRE algorithm with 1000 auction rounds per run.

VI. CONCLUDING REMARKS

A. Summary of Key Findings

The careful testing of auction protocols for a restructured electricity market by means of an agent-based computational model imposes two requirements on this model. First, the

model should adequately reflect the actual microstructure of the market. Second, the artificial traders in the model should behave in ways that approximate the behavior of real traders.

This study attempts to capture in a computational model the basic features of a restructured wholesale electricity market operating in the short run, abstracting from longer run contracting

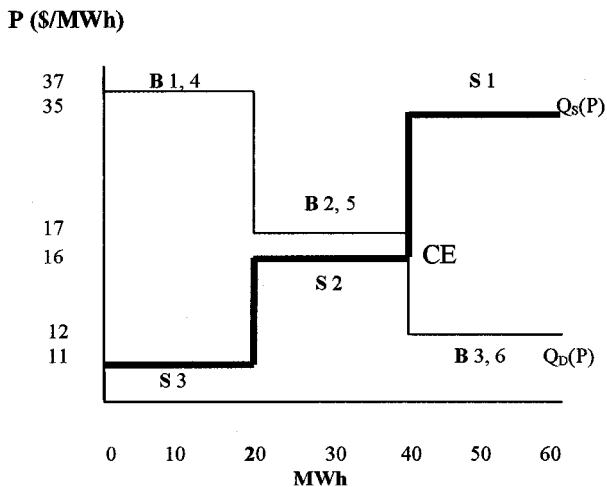


Fig. 4. Cell (3, 2) true demand and supply curves for six buyers and three sellers with 10 MWh and 20 MWh capacities, respectively.

considerations. We have assumed that, under restructuring, a small number of heterogeneous buyers (energy service providers) actively participate in this market along with a small number of heterogeneous sellers (generators). The buyers and sellers submit price and quantity offers repeatedly to a clearinghouse double auction that employs discriminatory midpoint pricing. The capacities, marginal revenues, and marginal costs of the buyers and sellers are private information. These features imply that the buyers and sellers face an inherently complex strategic situation.¹³

Moreover, we have attempted to implement learning in a manner supported by empirical data. Each buyer and seller is assumed to update its price offers over time in accordance with the MRE learning algorithm, a modified version of an individual reinforcement learning algorithm developed by Roth and Erev [16], [17]. The latter researchers developed their algorithm on the basis of extensive data obtained from experiments with human subjects in multiagent decision environments.

As detailed in Section III-A, we test two hypotheses regarding relative market power: H1 and H2. As seen in Section IV, one key finding of this study is that hypothesis H1 is not supported. When RCON is held fixed, the effects of increasing RCAP are often in the opposite direction of the effects predicted by H1.

A second key finding of this study is that hypothesis H2 also receives no support. Holding RCAP fixed, changes in RCON have only small unsystematic effects on market power in contradiction to H2.¹⁴ This latter finding strongly cautions against the common practice of confounding capacity and concentration effects in market power studies by letting firm size and numbers of firms vary together in an uncontrolled way.

¹³Indeed, in game-theoretic terms, it can be shown that the computational electricity market has numerous "pure Nash equilibria," i.e., numerous offer configurations that satisfy the following condition: given the price and quantity offers of all other traders, no individual trader can increase its profits by means of a unilateral deviation from its own current price and quantity offers.

¹⁴Interestingly, a similar finding of unsystematic effects for changes in RCON, taking RCAP as given, is reported in [22] for a computational labor market.

The problem with hypotheses H1 and H2 in the current electricity market context is that they attempt to predict relative market power effects purely on the basis of aggregate aspects of market structure as measured by RCAP and RCON. However, as seen in Section IV, neither aggregate turns out to be well correlated with the actual opportunities open to individual buyers and sellers to exercise market power under the discriminatory auction protocol.

In contrast, as shown in Section V, the microstructure of the electricity market is strongly predictive for the relative ability of buyers and sellers to exercise market power in the discriminatory auction. More precisely, the relative market power levels attained by buyers and sellers when they are permitted to learn to make opportunistic price offers closely track the relative "structural" market power levels they attain when the buyers are instead forced to bid their true willingness to pay and the sellers are instead forced to ask their true marginal costs.

Interestingly, examining the market power results obtained in an earlier electricity study [23, Table III] in which the buyer and seller populations instead each engage in social mimicry learning via a genetic algorithm (GA), it is seen that a similar conclusion holds. Structural market power is strongly predictive for the relative exercise of market power by the buyers and sellers; the effects of GA social mimicry learning on relative market power are small and unsystematic. For ease of comparison, these GA results are reproduced here as Table VIII.¹⁵

Taken together, these relative market power outcomes suggest that the microstructure of our electricity market under the discriminatory auction protocol so strongly channels the behavior of buyers and sellers that the precise form of their learning behavior is largely irrelevant. As noted next, however, this robustness to variations in learning behavior does not extend fully to efficiency outcomes.

As detailed in Section III-A, we also test a basic market efficiency hypothesis (H3). A third key finding of our study is that this hypothesis H3 is strongly supported. The market efficiency measure EA is 90% or better for almost all of the tested RCAP/RCON configurations.

The particular parameter values specified for the MRE learning algorithm used by traders in our current electricity market study do affect the ability of some traders to avoid coordination failure. However, in accordance with hypothesis H3, the trades in question are marginal trades offering the smallest gains and the resulting effects on market efficiency generally tend to be small and unsystematic. Our experiments suggest that the number of auction rounds per run may be a more important determinant of market efficiency than these learning parameter settings per se.

On the other hand, the market efficiency levels of the auction outcomes obtained in [23] under the assumption that the electricity traders instead use GA social mimicry learning are reported in Table VIII. These results show that market efficiency obtained with GA social mimicry learning is substantially degraded relative to market efficiency obtained with individual

¹⁵Table VIII corrects a labeling problem in [23, Table III], namely, the row labeled RCON = 1/2 should instead have been labeled RCON = 2 and vice versa.

TABLE VIII
EXPERIMENTALLY DETERMINED MARKET POWER AND EFFICIENCY OUTCOMES WITH GA SOCIAL LEARNING FROM [23]

	1/2		Relative Capacity 1		2	
2	MP	StdDev	MP	StdDev	MP	StdDev
	All Buyers: -0.53*	(0.35)	All Buyers: -0.50*	(0.39)	All Buyers: -0.37	(0.57)
	All Sellers: 0.19	(1.04)	All Sellers: 0.19	(1.02)	All Sellers: -0.43	(0.54)
	Buyer[1]: -0.52*	(0.36)	Buyer[1]: -0.49*	(0.40)	Buyer[1]: -0.37	(0.57)
	Buyer[2]: -0.63*	(0.62)	Buyer[2]: -0.99*	(0.09)	Buyer[2]: ZP	(0.00)
	Buyer[3]: ZP	(0.00)	Buyer[3]: ZP	(0.00)	Buyer[3]: ZP	(0.00)
	Seller[1]: ZP	(0.00)	Seller[1]: ZP	(0.00)	Seller[1]: ZP	(0.00)
	Seller[2]: ZP	(0.00)	Seller[2]: -0.72	(1.89)	Seller[2]: -0.44	(0.65)
	Seller[3]: -0.12	(1.47)	Seller[3]: 0.30	(1.12)	Seller[3]: -0.42	(0.50)
	Seller[4]: ZP	(0.00)	Seller[4]: ZP	(0.00)	Seller[4]: ZP	(0.00)
Seller[5]: ZP	(0.00)	Seller[5]: ZP	(0.00)	Seller[5]: -0.44	(0.66)	
Seller[6]: 0.50	(1.84)	Seller[6]: 0.27	(1.12)	Seller[6]: -0.41	(0.51)	
Efficiency: 60.50		Efficiency: 65.33		Efficiency: 59.68		
Relative Concentration 1	MP	StdDev	MP	StdDev	MP	StdDev
	All Buyers: -0.66*	(0.29)	All Buyers: -0.60*	(0.34)	All Buyers: -0.44	(0.57)
	All Sellers: 0.25*	(1.13)	All Sellers: 0.24	(1.12)	All Sellers: -0.50*	(0.49)
	Buyer[1]: -0.63*	(0.32)	Buyer[1]: -0.59*	(0.35)	Buyer[1]: -0.44	(0.57)
	Buyer[2]: -0.90*	(0.35)	Buyer[2]: -1.00	(0.00)	Buyer[2]: ZP	(0.00)
	Buyer[3]: ZP	(0.00)	Buyer[3]: ZP	(0.00)	Buyer[3]: ZP	(0.00)
	Seller[1]: ZP	(0.00)	Seller[1]: ZP	(0.00)	Seller[1]: ZP	(0.00)
	Seller[2]: ZP	(0.00)	Seller[2]: -0.20	(3.30)	Seller[2]: -0.52	(0.58)
	Seller[3]: 0.19*	(1.17)	Seller[3]: 0.28	(1.26)	Seller[3]: -0.48	(0.49)
	Efficiency: 51.06		Efficiency: 58.67		Efficiency: 52.68	
1/2	MP	StdDev	MP	StdDev	MP	StdDev
	All Buyers: -0.59*	(0.28)	All Buyers: -0.53*	(0.32)	All Buyers: -0.30	(0.54)
	All Sellers: 0.71	(1.15)	All Sellers: 0.67	(1.12)	All Sellers: -0.30	(0.44)
	Buyer[1]: -0.55*	(0.30)	Buyer[1]: -0.51*	(0.32)	Buyer[1]: -0.31	(0.58)
	Buyer[2]: -0.83*	(0.46)	Buyer[2]: -1.00	(0.00)	Buyer[2]: ZP	(0.00)
	Buyer[3]: ZP	(0.00)	Buyer[3]: ZP	(0.00)	Buyer[3]: ZP	(0.00)
	Buyer[4]: -0.56*	(0.29)	Buyer[4]: -0.52*	(0.32)	Buyer[4]: -0.30	(0.56)
	Buyer[5]: -0.79*	(0.50)	Buyer[5]: -1.00	(0.00)	Buyer[5]: ZP	(0.00)
	Buyer[6]: ZP	(0.00)	Buyer[6]: ZP	(0.00)	Buyer[6]: ZP	(0.00)
	Seller[1]: ZP	(0.00)	Seller[1]: ZP	(0.00)	Seller[1]: ZP	(0.00)
Seller[2]: ZP	(0.00)	Seller[2]: 0.05	(4.53)	Seller[2]: -0.38	(0.57)	
Seller[3]: 0.70	(1.17)	Seller[3]: 0.73	(1.29)	Seller[3]: -0.24	(0.48)	
Efficiency: 65.38		Efficiency: 73.67		Efficiency: 70.00		

ZP indicates that zero profits were earned both in the auction and in competitive equilibrium.

MRE learning. Consequently, market efficiency is not robust with respect to switches from individual to social learning.¹⁶

As detailed in Section V-E, a careful examination of the bid and ask behavior of individual buyers and sellers in our electricity market provides an explanation for these market power and efficiency findings. Since both trader types can submit opportunistic price offers, each type has countervailing strategic market power that holds in check the ability of the other type to offset the structural market power biases inherent in the discrim-

inatory auction protocol per se. This is true whether the traders use individual or social learning.

Nevertheless, under individual MRE learning, inframarginal traders are better able to home in on bids and asks that are sufficiently close to their true reservation prices to ensure positive bid-ask spreads and prevent entry by extramarginal traders. Coordination failure is, thus, largely avoided, meaning that the set of trades conducted under the discriminatory auction agrees closely with the set of trades that would occur in competitive equilibrium. The result is high market efficiency, despite the fact that some buyers and sellers have substantially positive or negative structural market power levels. The root cause of the inefficiency under GA social mimicry learning is a relatively high

¹⁶For an example of an oligopoly market in which a switch from individual to social learning results in substantially higher average output, see [23, Fig. 5]. In Vriend's example, all firms have identical costs.

coordination failure rate due to inappropriate mimicry by structurally distinct traders.

B. Comparison With Other Studies

Our findings regarding relative market power support the conclusion reached by Weiss [4] for human-subject electricity market experiments that active bidding by buyers may limit the ability of sellers to exercise market power. However, since Weiss only considers nodal and uniform pricing, he is not led to consider the distinction between structural and strategic market power.

In our context, the discriminatory auction protocol inherently allocates market power to some traders in preference to others, even in the absence of opportunistic bids and asks. On the other hand, all buyers and sellers can attempt to secure strategic market power in the sense that buyers can attempt to increase their profits by bidding lower than their true marginal revenues and sellers can attempt to increase their profits by asking higher than their true marginal costs. Therefore, our market power conclusion must be nuanced by saying that the presence of active traders on each side of the market reduces the ability of structurally disadvantaged traders to overcome the structural market power biases inherent in the auction protocol through the exercise of strategic market power. In addition, the ability to exercise strategic market power is further limited in our context by the threat of entry by extramarginal traders.

Our findings regarding market efficiency are reminiscent of the conclusions reached by Gode and Sunder [18] and other previous researchers regarding the efficiency of continuous double auctions. A *continuous double auction* is a double auction in which bids and asks are continuously received, trades can occur at any time, and bids and asks are accepted by the traders themselves rather than matched by a clearinghouse. As noted in [5, pp. 5–6], continuous double auctions have been observed in human-subject experiments to induce very efficient outcomes under a wide range of treatment conditions, much more so than traditional economic theory would suggest.

A key question raised by Gode and Sunder [18] is the extent to which the efficiency of any given market mechanism is attributable to trader rationality or inherent in the design of the mechanism. Their findings for continuous double auctions with zero-intelligence traders suggest that efficiency is inherent in the continuous double-auction mechanism per se.

Similarly, our market efficiency findings would seem to suggest that efficiency is inherent in the design of the discriminatory clearinghouse double-auction mechanism. Nevertheless, this conclusion is tempered by two additional findings.

First, market efficiency is seriously degraded when the buyer and seller populations each use GA social mimicry learning instead of individual MRE learning. This form of social mimicry is not particularly appropriate in the current electricity context since buyers have different marginal revenues and sellers have different marginal costs. Second, as indicated in footnote 6 of Section II-E, market efficiency can also be seriously degraded when buyers and sellers learn in accordance with the original RE learning algorithm, implying that they do not respond to the unfavorable stimulus of zero profits.

These two additional findings suggest that the following caution is in order. While the discriminatory clearinghouse double auction may reliably deliver high market efficiency when buyers and sellers refrain from inappropriate learning behavior, it may not be robust against the active exercise of bad judgment.

The extent to which our market power and efficiency findings generalize to wholesale electricity markets operating under different auction protocols is an interesting open question. For example, would our findings generalize to clearinghouse double-auctions with uniform pricing? Or to continuous double auctions exhibiting the various special types of rules (opening price rules, priority rules, etc.) listed in [25] on the basis of a survey of systems in actual operation?

Under alternative auction protocols, the learning behavior of traders might have more substantial effects on market power or market efficiency because the traders have a greater leeway for the exercise of strategic market power. In any case, it might be that the MRE learning algorithm applied in the current study is too simplistic to capture fully the strategic opportunities open to the traders.

For example, Camerer and Ho [26] have developed an individual learning algorithm that permits traders to use past observations to form beliefs about what other traders will do in the future. Would the observed strategic behavior of traders under Camerer–Ho learning differ significantly from what we have observed using individual MRE learning and GA social mimicry learning? In particular, would the use of Camerer–Ho learning permit traders to overcome structural market power biases through strategic pricing? The inability of the traders in our current experiments to overcome structural market power biases through strategic pricing appears to be due more to the symmetric design of the double auction (simultaneous bids and asks) than to any lack of learning power per se. However, a more systematic examination of this issue is clearly needed.

In addition, the representation of the traders' actions in the current study is very simplistic: each trader submits a single price offer and a single quantity offer to the auction in each auction round. In contrast, as described in [10], each generator in the England and Wales day-ahead electricity market submits price and quantity offers for up to three incremental levels of output for each of its generating units, i.e., in each auction round, each generator submits a supply *function* rather than a single price-quantity pair. The domain of possible offers by each generator is thus enormously enlarged from a set of points to a set of functions.

To handle these real-world features, it seems essential to permit the traders to engage in more comprehensive forms of learning that include inductive reasoning (experimentation with new ideas) as well as aspects of reinforcement learning, social mimicry, and forecasting of future events. As discussed in [27], social scientists are just beginning to appreciate the care and attention needed to model computationally the learning behavior of multiple social agents interacting in complex real-world contexts.

These issues, critically important for the computational modeling of restructured electricity markets, will be addressed in future studies.

ACKNOWLEDGMENT

The authors would like to thank D. Fogel and two anonymous referees for thoughtful constructive comments. The authors would also like to thank D. Lane, H. Quirnbach, G. Sheblé, and especially D. Koesrindartoto and A. Kroujiline for many helpful discussions regarding this work. Finally, special thanks are extended to D. Ashlock and G. Sheblé for the use of their basic EPRI Simulator code that was used in modified form for this study. The executable files used to run all experiments reported in this paper are available from the third author upon request.

REFERENCES

- [1] H. C. Quirnbach and J. P. Acton, "Structural reform in electric power: A framework for analysis," RAND, Santa Monica, CA, Doc. R-3596-NSF/RC, 1987.
- [2] T. J. Brennan, K. L. Palmer, R. J. Kopp, A. J. Krupnick, V. Staglino, and D. Burttraw, *A Shock to the System: Restructuring America's Electricity Industry*. Washington, DC: Resources for the Future, 1996.
- [3] D. M. Newbery, *Privatization, Restructuring, and Regulation of Network Utilities*. Cambridge, MA: MIT Press, 1999.
- [4] J. Weiss, *Market Power Issues in the Restructuring of the Electricity Industry: An Experimental Investigation*. Cambridge, MA: Harvard Bus. School, 1999.
- [5] D. Friedman, "The double auction institution: A survey," in *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, Eds. Reading, MA: Addison-Wesley, 1993, vol. XIV, Santa Fe Institute Studies in the Sciences of Complexity, pp. 3–25.
- [6] G. B. Sheblé, *Computational Auction Mechanisms for Restructured Power Industry Operation*. Norwell, MA: Kluwer, 1999.
- [7] K. Klemperer, "Auction theory: A guide to the literature," in *The Economic Theory of Auctions*, P. Klemperer, Ed. Cheltenham, U.K.: Edward Elgar, 2000, pp. 3–62.
- [8] N.-H. M. von der Fehr and D. Harbord, "Spot market competition in the U.K. electricity industry," *Econom. J.*, vol. 103, pp. 531–546, 1993.
- [9] P. Klemperer, "What really matters in auction design," Nuffield College, Oxford University, Oxford, U.K., Working Paper, 2000, to be published.
- [10] J. Bower and D. Bunn, "Experimental analysis of the efficiency of uniform-price versus discriminatory price auctions in the England and Wales electricity market," *J. Econom. Dyn. Control*, vol. 25, no. 3–4, pp. 561–592, 2001.
- [11] D. Friedman and J. Rust, Eds., *The Double Auction Market: Institutions, Theories, and Evidence*. Reading, MA: Addison-Wesley, 1993, vol. XIV, Santa Fe Institute Studies in the Sciences of Complexity.
- [12] C. A. Holt, "Industrial organization: A survey of laboratory research," in *Handbook of Experimental Economics*, J. A. Kagel and A. E. Roth, Eds. Princeton, NJ: Princeton Univ. Press, 1995, pp. 348–443.
- [13] J. A. Kagel, "Auctions: A survey of experimental research," in *Handbook of Experimental Economics*, J. A. Kagel and A. E. Roth, Eds. Princeton, NJ: Princeton Univ. Press, 1995, pp. 501–535.
- [14] J. Rust, J. H. Miller, and R. Palmer, "Characterizing effective trading strategies: Insights from a computerized double auction tournament," *J. Econom. Dyn. Control*, vol. 18, pp. 61–96, 1994.
- [15] D. Lane, A. Kroujiline, V. Petrov, and G. Sheblé, "Electricity market power: Marginal cost and relative capacity effects," in *Proceedings of the 2000 Congress on Evolutionary Computation*, vol. II, A. Alzala, Ed.. Piscataway, NJ, 2000, pp. 1048–1054.
- [16] A. E. Roth and I. Erev, "Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term," *Games Econ. Beh.*, vol. 8, pp. 164–212, 1995.
- [17] I. Erev and A. E. Roth, "Predicting how people play games with unique, mixed-strategy equilibria," *Amer. Econom. Rev.*, vol. 88, pp. 848–881, 1998.
- [18] D. K. Gode and S. Sunder, "Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality," *J. Pol. Econom.*, vol. 101, pp. 119–137, 1993.
- [19] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: Cambridge Univ. Press, 1998.
- [20] R. J. Green and D. M. Newbery, "Competition in the British electricity spot market," *J. Pol. Econom.*, vol. 100, pp. 929–953, 1992.

- [21] F. L. Alvarado, "Market power: A definition," in *Proceedings of the 1998 Conference on Bulk Power Systems Dynamics and Control: IV Restructuring*, Aug. 1998.
- [22] L. Tesfatsion, "Structure, behavior, and market power in an evolutionary labor market with adaptive search," *J. Econom. Dyn. Control*, vol. 25, no. 3–4, pp. 419–457, 2001.
- [23] J. Nicolaisen, M. Smith, V. Petrov, and L. Tesfatsion, "Relative capacity and concentration effects on electricity market power," in *Proceedings of the 2000 Congress on Evolutionary Computation*, A. Alzala, Ed. Piscataway, NJ: IEEE Press, 2000, vol. II, pp. 1041–1047.
- [24] N. Vriend, "An illustration of the essential difference between individual and social learning and its consequences for computational analyzes," *J. Econom. Dyn. Control*, vol. 24, pp. 1–19, 2000.
- [25] I. Domowitz, "Automating the continuous double auction in practice: Automated trade execution systems in financial markets," in *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, Eds. Reading, MA: Addison-Wesley, 1993, vol. XIV, Santa Fe Institute Studies in the Sciences of Complexity, pp. 27–60.
- [26] C. Camerer and T.-H. Ho, "Experience-weighted attraction learning in normal form games," *Econometrica*, vol. 67, pp. 827–874, 1999.
- [27] E. Chattoe, (1998) Just how (Un)realistic are evolutionary algorithms as representations of social processes?. *J. Artificial Societ. Soc. Sim.* [Online] <http://www.soc.surrey.ac.uk/JASSS/1/3/2.html>



James Nicolaisen received the B.S. degree in physics from Morningside College, Sioux City, IA, in 1998. He is currently working towards the M.S. degree in electrical engineering with a minor in economics at Iowa State University, Ames.



Valentin Petrov received the M.S. degree in physics from Sofia University, Sofia, Bulgaria, in 1993 and the M.S. degree in nonlinear optics from Bowling Green State University, Bowling Green, OH, in 1998. He is currently working towards the M.S. degree in computer engineering with a minor in economics at Iowa State University, Ames.

He was previously with Varian Associates, Palo Alto, CA.



Leigh Tesfatsion received the Ph.D. degree in economics with a minor in mathematics from the University of Minnesota, Minneapolis, in 1975.

She is currently a Professor of Economics with Iowa State University, Ames, with a courtesy appointment as a Professor of Mathematics. Her current research interests focus on agent-based computational economics (ACE), the computational study of economies modeled as evolving systems of autonomous interacting agents, which she is applying to the study of market power in labor

markets and to the study of restructuring in electricity markets.

Dr. Tesfatsion is a Co-Editor in charge of the Complexity-at-Large section of *Complexity* as well as an Associate Editor of *Journal of Economic Dynamics and Control*, *Journal of Public Economic Theory*, *Applied Mathematics and Computation*, and IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION.