Capacity Withholding in Restructured Wholesale Power Markets: An Agent-Based Test Bed Study

Hongyan Li, *Member, IEEE*, and Leigh Tesfatsion, *Member, IEEE*

***Abstract*—This study uses a dynamic 5-bus test case implemented via the AMES Wholesale Power Market Test Bed to investigate strategic capacity withholding by generation companies (GenCos) in restructured wholesale power markets under systematically varied demand conditions. The strategic behaviors of the GenCos are simulated by means of a stochastic reinforcement learning algorithm motivated by human-subject laboratory experiments. The learning GenCos attempt to improve their earnings over time by strategic selection of their reported supply offers. This strategic selection can involve both *physical* capacity withholding (reporting of lower-than-true maximum operating capacity) and *economic* capacity withholding (reporting of higher-than-true marginal costs). We explore the ability of demand conditions to mitigate incentives for capacity withholding by letting demand bids vary from 100% fixed demand to 100% price-sensitive demand.**

***Index Terms*—Capacity withholding, demand-bid price sensitivity, restructured wholesale power markets, locational marginal pricing, multi-agent stochastic reinforcement learning, dynamic 5-bus test case, AMES Wholesale Power Market Test Bed**

I. Introduction

The U.S. Federal Energy Regulatory Commission (FERC) in an April 2003 white paper [1] proposed a market design for common adoption by U.S. wholesale power markets. Core features of this market design include: central management and oversight by an independent market operator; a two-settlement system consisting of a bid/offer-based day-ahead market supported by a parallel real-time market to ensure continual balancing of supply and demand for power; and management of transmission grid congestion by means of locational marginal pricing.

Joskow [2] estimates that over 50% of generating capacity in the U.S. is now operating under some variant of FERC's market design. Energy regions that have adopted (or plan to adopt) this design include the midwest (MISO), New England (ISO-NE), New York (NYISO), the mid-atlantic states (PJM), California (CAISO), the southwest (SPP), and Texas (ERCOT).

The complexity of FERC's market design – together with the relative recency of its adoption (implying short data series) – makes it extremely difficult to study its dynamic performance properties using standard analytical and statistical modeling tools. A key unresolved issue is the extent to which the complicated rules and regulations governing market operations under the design might encourage strategic bid/offer behaviors on the part of market participants that reduce overall market performance over time. A related issue is the extent to which grid congestion and load-pocket formation can be strategically manipulated to benefit certain market participants at the expense of others.

Fortunately, powerful new agent-based modeling tools are now available that can handle this degree of complexity. As detailed at [3], these tools are already fruitfully being applied to the study of restructured wholesale power markets.

For example, in a series of studies ([4], [5], [6]) we study *economic capacity withholding* (reporting of higher-than-true marginal costs) by profit-seeking generation companies (GenCos) participating in a 5-bus wholesale power market operating under FERC's market design. The GenCos strategically determine their supply offers over time using VRE reinforcement learning, a variant of a stochastic reinforcement learning algorithm developed by Alvin Roth and Ido Erev ([7], [8]) on the basis of human-subject laboratory studies. These economic capacity-withholding experiments were conducted using the AMES Wholesale Power Market Test Bed, an open-source computational laboratory specifically designed for the systematic experimental study of FERC's market design.[1]

Also, Tellidou and Bakirtzis [10] study *physical capacity withholding* (reporting of lower-than-true maximum operating capacity) as well as economic capacity withholding within an agent-based computational modeling of an energy auction market operating over a two-bus transmission grid with a fixed daily demand (load) profile. Their simulated GenCos have constant marginal costs and decide on hourly point quantity-price supply offers via SA-Q learning, a modified version of Q-learning. Each GenCo has the same learning parameters. The authors find that the GenCos are able to learn over time to exercise capacity withholding even if the only information available to them is public price data.

In this study we extend this earlier work. We use the AMES test bed to conduct systematic physical and economic capacity-withholding experiments for a dynamic 5-bus test case under alternative demand-bid price sensitivity conditions ranging

Latest revision: 19 October 2008. This work is scheduled to appear in the *Proceedings* of the 2009 IEEE Power Systems Conference and Exposition (PSCE), Seattle, WA. It has been supported in part by the National Science Foundation under Grant NSF-0527460 and by a grant from the ISU Electric Power Research Center.

Hongyan Li (corresponding author: lihy@iastate.edu), Electrical and Computer Engineering Department, Iowa State University, Ames, IA 50011 USA; and Leigh Tesfatsion (tesfatsi@iastate.edu), Economics Department, Iowa State University, Ames, IA 50011 USA.

[1] AMES is an acronym for *A*gent-based *M*odeling of *E*lectricity *S*ystems. The first version of AMES was formally released by the developers (H. Li, J. Sun, and L. Tesfatsion) as open-source software at the 2007 IEEE Power and Energy Society General Meeting. Downloads, manuals, and tutorial information for all AMES version releases to date can be accessed at the AMES homepage [9].

from 100% fixed demand (no price sensitivity) to 100% price sensitivity. GenCos and Load-Serving Entities (LSEs) participate in a day-ahead energy market operating over a 5-bus transmission grid with congestion managed by locational marginal pricing. As in actual ISO-managed energy markets such as the MISO [11], the supply offers of the GenCos consist of reported marginal cost functions over reported operating capacity intervals and the demand bids of the LSEs are combinations of fixed demand bids and price-sensitive demand bid functions. The GenCos rely on VRE reinforcement learning to determine their reported supply offers over time.

Real-world restructured wholesale power markets are sequential open-ended games in that multiple participant traders must decide on bids/offers for electric power on a daily basis, with no fixed horizon. Presumably, then, the traders will attempt to optimize their learning methods over time as they gain market experience.

In recognition of this learning-to-learn issue, we take preliminary steps in this study to help ensure that the GenCos' learning methods are calibrated to their decision environment. We conduct initial experiments for the dynamic 5-bus test case with 100% fixed demand involving extensive parameter sweeps for key VRE learning parameters. We use these initial learning experiments to determine GenCo-individuated *sweet spot* VRE learning parameter values resulting in the highest daily net earnings for the GenCos. We then set each GenCo's VRE learning parameters to its sweet-spot values for all subsequent experiments.

Section II outlines the main features of the AMES test bed. Section III explains the experimental design used to explore GenCo capacity withholding under systematically varied settings for demand-bid price-sensitivity when GenCos have sweet-spot VRE learning capabilities. Section IV explains more carefully how we determined these sweet-spot VRE learning capabilities. Experimental findings for GenCo capacity withholding are reported in Section V. Concluding remarks are given in Section VI.

II. The AMES Test Bed (Version 2.01)

A. Overview

This study uses Version 2.01 of the AMES Wholesale Power Market Test Bed to conduct all reported experiments. AMES(V2.01) incorporates core features of the wholesale power market design proposed by the U.S. FERC [1]; see Fig. 1. A detailed description of these features can be found in materials provided at the AMES homepage [9].

Below is a summary description of the logical flow of events in the AMES(V2.01) wholesale power market:

- The AMES wholesale power market operates over an *AC transmission grid* starting on day 1 and continuing through a user-specified maximum day (unless terminated earlier in accordance with a user-specified stopping rule). Each day D consists of 24 successive hours H = 00,01, ...,23.
- The AMES wholesale power market includes an *Independent System Operator (ISO)* and a collection of energy traders consisting of *Load-Serving Entities (LSEs)* and *Generation Companies (GenCos)* distributed across the busses of the transmission grid. Each of these entities is implemented as a software program encapsulating both methods and data; see, e.g., the schematic depiction of a GenCo in Fig. 2
- The objective of the ISO is the reliable attainment of appropriately constrained *operational efficiency* for the wholesale power market, i.e., the maximization of total net benefits subject to generation and transmission constraints.
- In an attempt to attain this objective, the ISO undertakes the daily operation of a *day-ahead market* settled by means of *locational marginal pricing (LMP)*. Roughly stated, a *locational marginal price* at any particular transmission grid bus is the least cost of servicing demand for one additional megawatt (MW) of power at that bus.[2]
- The objective of each LSE is to secure power for its downstream (retail) customers. During the morning of

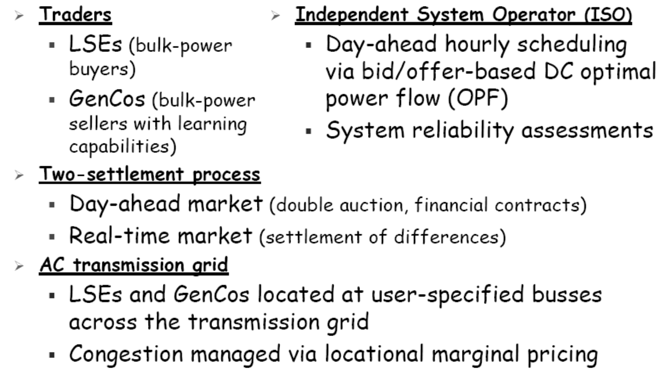

Fig. 1. AMES test bed architecture

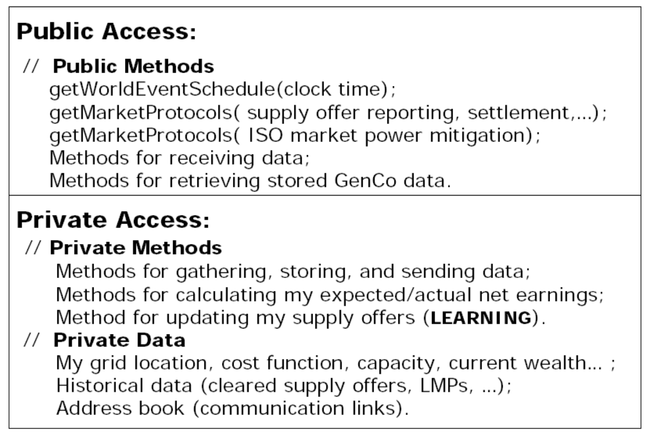

Fig. 2. AMES GenCo: A cognitive agent with learning capabilities

[2]In reality, LMPs are shadow prices for "nodal balance constraints" constituting part of the constraint set of optimal power flow problems and are derived as derivatives of the optimized power flow objective function with respect to particular types of perturbations of these constraints. Moreover, these nodal balance constraints are imposed at "pricing nodes" that might not correspond to actual physical bus locations on the grid. For expositional simplicity, throughout this study we use the standard engineering short-hand description for LMPs as valuations for single-unit increases in demand and we treat pricing nodes as coincident with transmission grid busses. For a more rigorous explanation and derivation of LMPs, see [6].

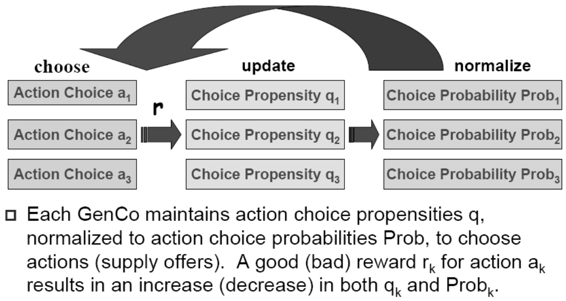

Fig. 3. AMES GenCos use stochastic reinforcement learning to determine the supply offers they report to the ISO for the day-ahead market.

each day D, each LSE reports a demand bid to the ISO for the day-ahead market for day D+1. Each demand bid consists of two parts: a *fixed demand bid* (i.e., a 24-hour load profile); and 24 *price-sensitive demand bids* (one for each hour), each consisting of a linear demand function defined over a purchase capacity interval. LSEs have no learning capabilities; LSE demand bids are user-specified at the beginning of each simulation run.

- The objective of each GenCo is to secure for itself the highest possible net earnings each day. During the morning of each day D, each GenCo i uses its current action choice probabilities to choose a *supply offer* from its action domain AD_i to report to the ISO for use in all 24 hours of the day-ahead market for day D+1.
- Each supply offer in AD_i consists of a linear marginal cost function defined over an operating capacity interval. GenCo i's ability to vary its choice of a supply offer from its action domain AD_i permits it to adjust the ordinate/slope of its reported marginal cost function and/or the upper limit of its reported operating capacity interval in an attempt to increase its daily net earnings.
- After receiving demand bids from LSEs and supply offers from GenCos during the morning of day D, the ISO determines and publicly reports hourly power supply commitments and LMPs for the day-ahead market for day D+1 as the solution to hourly bid/offer-based *DC optimal power flow (DC-OPF)* problems. *Transmission grid congestion* is managed by the inclusion of congestion cost components in LMPs.
- At the end of each day D, the ISO settles all commitments for the day-ahead market for day D+1 on the basis of the LMPs for the day-ahead market for day D+1.
- At the end of each day D, each GenCo i uses *stochastic reinforcement learning* to update the action choice probabilities currently assigned to the supply offers in its action domain AD_i, taking into account its day-D settlement payment ("reward"). In particular, as depicted in Fig. 3, if the supply offer reported by GenCo i on day D results in a relatively good reward, GenCo i increases the probability of choosing this supply offer on day D+1, and conversely.
- There are no system disturbances (e.g., weather changes) or shocks (e.g., forced generation outages or line outages). Consequently, the binding financial contracts determined in the day-ahead market are carried out as planned and traders have no need to engage in real-time (spot) market trading.
- Each LSE and GenCo has an initial holding of money that changes over time as it accumulates earnings and losses.
- There is no entry of traders into, or exit of traders from, the wholesale power market. LSEs and GenCos are currently allowed to go into debt (negative money holdings) without penalty or forced exit.

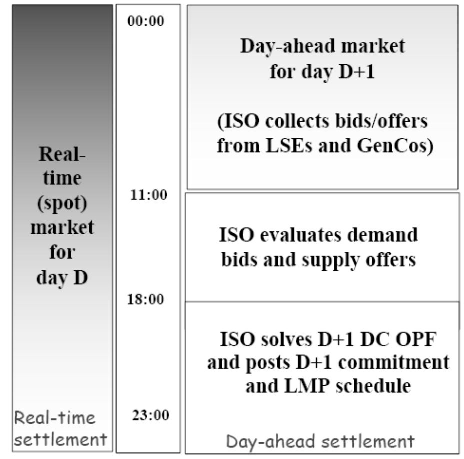

Fig. 4. AMES ISO activities during a typical day D

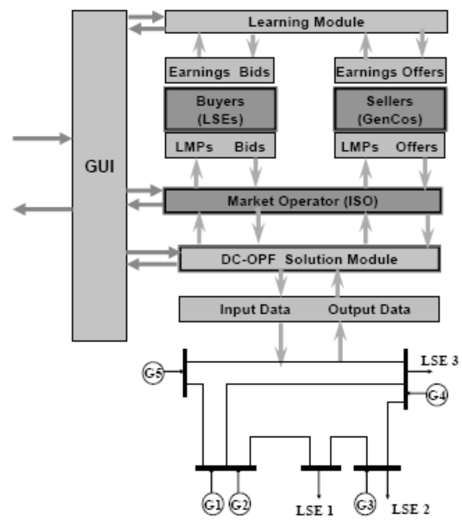

Fig. 5. Illustration of AMES dynamics on a typical day D in the absence of system disturbances or shocks for the special case of a 5-bus grid

The activities of the ISO on a typical day D are depicted in Fig. 4. The overall dynamical flow of activities in the wholesale power market on a typical day D in the absence of system disturbances or shocks is depicted in Fig. 5.

B. Demand Bids and Supply Offers

On each day D, each LSE j reports 24 demand bids for use in the 24 hours of the day-ahead market for day D+1. The demand bid for hour H consists of a *fixed demand bid* p^F_{Lj}(H) (in MWs) and a *price-sensitive demand bid function*

$$D_{jH}(p^S_{Lj}(H)) \;=\; c_j(H) \;-\; 2d_j(H)\cdot p^S_{Lj}(H) \qquad (1)$$

defined over a *true purchase capacity interval*

$$0 \leq p^S_{Lj}(H) \leq SLMax_j(H) , \qquad (2)$$

where p^S_{Lj}(H) is real electric power (in MWs). The expression $D_{jH}(p^S_{Lj}(H))$ denotes LSE j's *true purchase reservation value* for p^S_{Lj}(H), i.e., the maximum dollar payment it is truly willing to make (per MWh) for p^S_{Lj}(H).

On each day D, each GenCo i's *true marginal cost function* for each hour H of the day-ahead market for day D+1 takes the form of a linear function

$$MC_i(p_{Gi}(H)) = a_i + 2b_i \cdot p_{Gi}(H) \qquad (3)$$

defined over a *true operating capacity interval*

$$Cap^L_i \leq p_{Gi}(H) \leq Cap^U_i , \qquad (4)$$

where p_{Gi}(H) is real electric power (in MWs). The expression $MC_i(p_{Gi}(H))$ denotes GenCo i's *true sale reservation value* for p_{Gi}(H), i.e., the minimum dollar payment it is truly willing to accept (per MWh) for p_{Gi}(H).

On each day D, each GenCo i submits one *reported supply offer* to the ISO for use in each hour H of the day-ahead market for day D+1. This reported supply offer consists of a *reported marginal cost function*

$$MC^R_i(p_{Gi}) = a^R_i + 2b^R_i \cdot p_{Gi} \qquad (5)$$

defined over a *reported operating capacity interval*

$$Cap^L_i \leq p_{Gi} \leq Cap^{RU}_i , \qquad (6)$$

where p_{Gi} is real electric power (in MWs). The expression $MC^R_i(p_{Gi})$ denotes GenCo i's *reported sale reservation value* for p_{Gi}, i.e., the minimum dollar payment it *reports* it is willing to accept (per MWh) for p_{Gi}.

To avoid operating at a point where the true marginal cost of its last supplied MW of power exceeds the marginal benefit (received payment), GenCo i's reported marginal cost functions (5) lie on or above its true marginal cost function (3). In addition, to avoid infeasible commitments, GenCo i's *reported maximum operating capacity* Cap^{RU}_i in (6) never exceeds its true maximum operating capacity Cap^U_i in (4).

Note from the above discussion that each reported supply offer for GenCo i can be summarized in the form of a vector (a^R_i,b^R_i,Cap^{RU}_i).

C. GenCo Costs and Net Earnings

Total variable cost refers to the costs sustained by a supplier that vary with the level of its operations, whereas *fixed cost* refers to the costs sustained by a supplier independently of its level of operations. *Total cost* refers to the sum of the two.

For the specific context at hand, the *true total variable cost function* for GenCo i for each hour H takes the form

$$TVC_i(p_{Gi}) = \int_0^{p_{Gi}} MC_i(p)dp = a_i \cdot p_{Gi} + b_i \cdot [p_{Gi}]^2 , \qquad (7)$$

and the *true total cost function* for GenCo i for each hour H takes the form

$$TC_i(p_{Gi}) = [TVC_i(p_{Gi}) + FCost_i] , \qquad (8)$$

where p_{Gi} denotes any real-power generation level in (4). By definition, then, the *fixed cost* for GenCo i in each hour H takes the form $TC_i(0) = FCost_i$.

Net earnings are defined as revenues minus true total *variable* cost. Suppose, in particular, that GenCo i is located at bus $k(i)$ and is committed at a generation level p_{Gi} at price $LMP_{k(i)}$ for hour H of the day-ahead market for day D+1. Then the the net earnings of GenCo i for hour H of day D+1 are given by

$$NE_i(H,D) = LMP_{k(i)} * p_{Gi} - TVC_i(p_{Gi}) . \qquad (9)$$

The net earnings of GenCo i over all 24 hours of day D+1, received in settlement from the ISO at the end of day D, are then given by

$$NE_i(D) = \sum_{H=00}^{H=23} NE_i(H,D) . \qquad (10)$$

D. Determination of LMPs and Power Commitments

The AMES ISO computes hourly LMPs and power commitments for the day-ahead market by solving bid/offer-based DC Optimal Power Flow (OPF) problems that approximate underlying AC-OPF problems. To handle these computations the AMES ISO makes repeated calls to *DCOPFJ*, an accurate and efficient Java DC-OPF solver developed as open-source software by Sun and Tesfatsion ([5], [6]).[3] DCOPFJ consists of a strictly convex quadratic programming solver wrapped in an outer SI-pu data conversion shell.

III. Experimental Design

All market performance experiments reported in this study are based on a *dynamic 5-bus test case* characterized by the following structural, institutional, and behavioral conditions:

- The 5-bus transmission grid configuration is as depicted in Fig. 6, with transmission grid, LSE, and GenCo structural attributes as presented in Li et al. [12].[4]
- The five GenCos in Fig. 6 are individual plant owners with distinct maximum operating capacities as follows: 110MW for GenCo 1 (G1); 100MW for GenCo 2 (G2); 520MW for GenCo 3 (G3); 200MW for GenCo 4 (G4); and 600 MW for GenCo 5 (G5). Note that the next-to-largest GenCo 3 is favorably situated in a potential "load pocket" with respect to the three LSEs.
- GenCo 4 (a "peaking unit") has the most costly generation. Next in line is GenCo 3. The three remaining GenCos 1, 2, and 5 have moderate costs.
- The daily fixed demand (load) profiles for the three LSEs are the same from one day to the next. As depicted in Fig. 7, each daily fixed demand profile peaks at hour 17.[5]

[3] A stand-alone version of DCOPFJ can be obtained at the software site for the IEEE Taskforce on Open-Source Software for Power Systems [13].

[4] The 5-bus transmission grid depicted in Fig. 6 is due to Lally [14]. This grid configuration is now used extensively in ISO-NE/PJM training manuals to derive quantity and price solutions at a given point in time assuming ISOs have complete and correct information about grid, LSE, and GenCo structural attributes.

[5] These profile shapes are adopted from a case study presented in Shahidehpour et al. [15, p. 296-297].

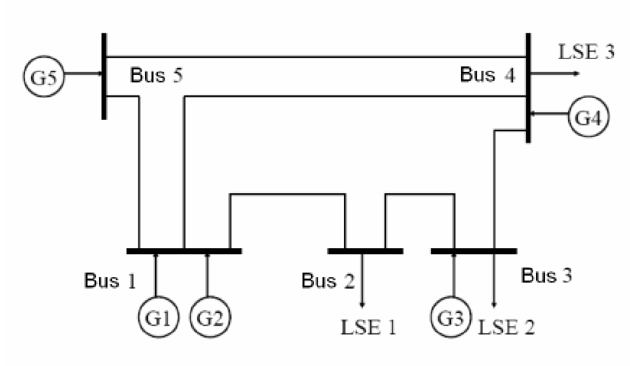

Fig. 6. 5-bus transmission grid for the dynamic 5-bus test case

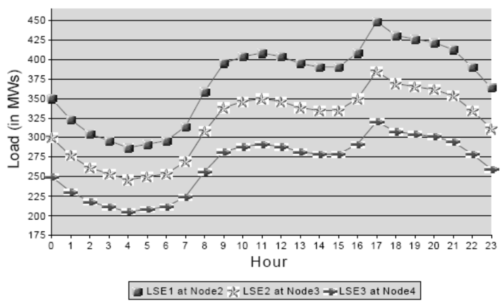

Fig. 7. Daily LSE fixed demand (load) profiles for the dynamic 5-bus test case

- The VRE learning parameters for each of the five GenCos are set at “sweet spot” values for which the GenCos as a whole earn the highest average daily net earnings.[6]

To control for purely random effects, we conducted thirty runs for each treatment using thirty distinct random seed values; see Li et al. [12] for the precise numerical values used. Also, unless otherwise indicated, experiments were conducted with all five AMES stopping rules flagged “on.” The stopping day for each run is referred to as the *final day* for that run.

Our primary treatment factor is the extent to which each GenCo can exercise physical capacity withholding by reporting lower-than-true maximum operating capacities. As clarified more carefully in Section V, we investigate two shrinkage rates for reported maximum operating capacities: 5% shrinkage (relative to true maximum operating capacity); and 10% shrinkage (relative to true maximum operating capacity).

Another treatment factor we consider is relative demand-bid price sensitivity. As our measure for this factor, we construct a ratio R of maximum potential price-sensitive demand to maximum potential total demand. More precisely, for each LSE j and each hour H, let

$$\mathrm{R}_j(H) \;=\; \frac{\mathrm{SLMax}_j(H)}{\mathrm{MPTD}_j(H)} \quad . \tag{11}$$

In (11) the expression SLMax$_j$(H) denotes LSE j’s *maximum potential price-sensitive demand* in hour H as measured by

[6]In particular, as explained in the following Section IV, we use Case(1,1) in Table I corresponding to the basic learning parameter settings $\alpha = 1$ and $\beta = 100$.

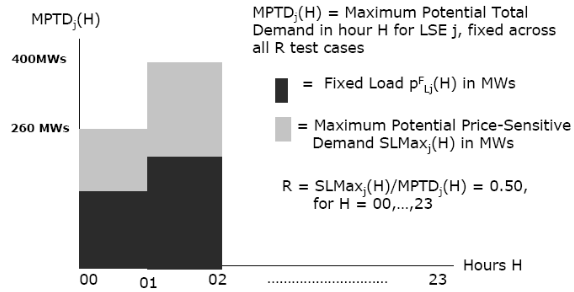

Fig. 8. Illustration of the construction of the R ratio for measuring relative demand-bid price sensitivity for the special case R=0.5

the upper bound of its purchase capacity interval (2), and

$$\mathrm{MPTD}_j(H) \;=\; [\mathrm{p}^F_{Lj}(H) + \mathrm{SLMax}_j(H)] \tag{12}$$

denotes LSE j’s *maximum potential total demand* in hour H as the sum of its fixed demand p^F_{Lj}(H) and its maximum potential price-sensitive demand SLMax$_j$(H) in hour H. The construction of the R ratio is illustrated in Fig. 8.

For our price-sensitive demand experiments we start by setting all of the R values (11) for each LSE j and each hour H equal to R=0.0 (the pure fixed-demand case). We then systematically increase R by tenths, ending with the value R=1.0 (the pure price-sensitive demand case). A positive R value indicates that the LSEs are able to exercise at least some degree of price resistance.

The maximum potential price-sensitive hourly demands SLMax$_j$(H) for each LSE j are thus systematically increased across experiments. However, we control for confounding effects arising from changes in overall demand capacity as follows: For each LSE j and each hour H, the denominator value MPTD$_j$(H) in (12) is held constant across experiments by appropriate reductions in the fixed demand p^F_{Lj}(H) as SLMax$_j$(H) is increased. Specifically, MPTD$_j$(H) is set equal across all experiments to BP^F_{Lj}(H), the hour-H fixed-demand level BP^F(H) for LSE j specified in Li et al. [12] for their benchmark dynamic 5-bus test case. Consequently, for each tested R value,

$$\mathrm{p}^F_{Lj}(H) \;=\; [\text{1-R}] * \mathrm{BP}^F_{Lj}(H) \; ; \tag{13}$$

$$\mathrm{SLMax}_j(H) \;=\; \mathrm{R} * \mathrm{BP}^F_{Lj}(H). \tag{14}$$

Moreover, as R is incrementally increased from R=0.0 to R=1.0, we control for confounding effects arising from changes in the LSEs’ price-sensitive demand bids by setting the ordinate and slope parameters $\{$(c$_j$(H),d$_j$(H)): H=00,...,23$\}$ to fixed values for each LSE j. A listing of the specific numerical values used can be found in Li et al. [12].

IV. Procedure for Determination of Sweet-Spot VRE Learning Parameter Values

This section reports on initial learning experiments conducted with the dynamic 5-bus test case outlined in Section III with 100% fixed demand and no physical capacity withholding. The purpose of these initial learning experiments is to determine “sweet-spot” VRE learning parameter values for the GenCos that perform reasonably well for their particular decision environment.

Reasonability is judged in terms of the average daily net earnings (Avg DNE) ultimately attained by the GenCos as a result of the supply offers (actions) they learn to report to the ISO over time. Avg DNE is calculated as the daily net earnings (10) earned on the final day D averaged across all five GenCos and across all thirty runs.

As detailed in Appendix A, the VRE reinforcement learning algorithm for each GenCo i is characterized by the following four parameters:

- GenCo i's *initial action choice propensity level* $q_i(1)$, which determines GenCo i's initial aspiration level for daily net earnings at the beginning of day 1;
- GenCo i's *temperature cooling rate* T_i, which controls the extent to which differences in GenCo i's action choice propensities translate into differences in GenCo i's action choice probabilities;
- GenCo i's *recency parameter* r_i, which controls the relative weight GenCo i places on current versus past "rewards" (daily net earnings outcomes) when it updates its action choice propensity values;
- GenCo i's *experimentation parameter* e_i, which dampens the growth of GenCo i's chosen-action propensity levels and controls the extent to which a reward resulting from a currently chosen action affects GenCo i's updating of its action choice propensities for non-chosen actions.

In extensive VRE learning experiments conducted for the benchmark dynamic 5-bus test case under alternative settings for the recency and experimentation parameters r and e over their full feasible ranges from 0 to 1, Pentapalli [16] determined that high Avg DNE outcomes were generally obtained with r=0.04 and e=0.96 for each GenCo. Consequently, throughout the present study we set r=0.04 and e=0.96 for each GenCo i.

Clearly the values set for the initial action choice propensity level $q_i(1)$ and temperature cooling rate T_i for each GenCo i should reasonably be calibrated to the particular earnings opportunities it faces. The following normalization is used to achieve this individual calibration while minimizing the total number of parameter values to be experimentally determined.

We first define a derived parameter

$$\alpha \;=\; \frac{q_i(1)}{\mathrm{MaxDNE}_i} \;,\; i = 1, \ldots, I \;, \tag{15}$$

where MaxDNE_i is an estimate for GenCo i's maximum possible daily net earnings derived from its action domain AD_i assuming "competitive" marginal-cost pricing (sales price = reported marginal cost). Specifically, letting s_i^R = (a_i^R,b_i^R,Cap_i^{RU}) denote a generic supply offer in AD_i,[7]

$$MaxDNE_i \;=\; 24 * \left(\max_{s_i^R \in AD_i} \left[HNE(s_i^R) \right] \right) \;, \tag{16}$$

[7] Compare (16) with definition (10) for the actual net earnings of GenCo i over all 24 hours of the day-ahead market for day D+1 under LMP pricing. The LMP received by GenCo i at a positive generation commitment level p_{Gi} in any hour H can exceed GenCo i's reported marginal cost at p_{Gi} for hour H if GenCo i has a binding upper operating capacity limit at p_{Gi}. This is why MaxDNE_i is characterized as an estimate rather than a true upper bound for GenCo i's maximum possible daily net earnings.

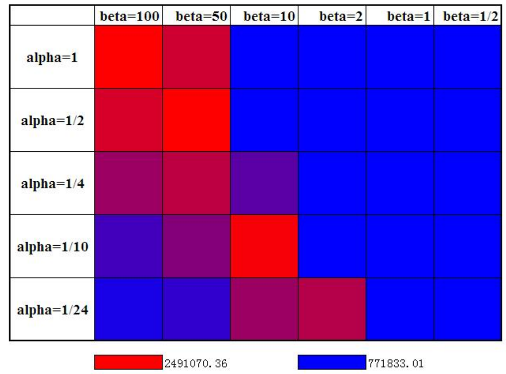

Fig. 9. A heat-map depiction of average daily net earnings (Avg DNE) outcomes under alternative (α, β) VRE learning parameter combinations. Lighter shades indicate higher Avg DNE.

where the hourly net earnings function HNE(s_i^R) satisfies

$$HNE(s_i^R) = MC_i^R(Cap_i^{RU}) * Cap_i^{RU} - TVC_i(Cap_i^{RU}) \;. \tag{17}$$

Given a (typically) distinct positive value MaxDNE_i for each GenCo i, a non-negative setting for α determines a distinct initial earnings aspiration level $q_i(1)$ for each GenCo i. Low α values correspond to pessimistic aspiration levels (relative to MaxDNE_i), and conversely.

We then define a second derived parameter

$$\beta \;=\; \frac{q_i(1)}{T_i} \;\;,\; i = 1, \ldots, I \;. \tag{18}$$

Given a non-negative value for α in (15) – and hence a value for $q_i(1)$ for each GenCo i – a non-negative setting for β in (18) determines a temperature cooling rate T_i for each GenCo i. Low β values correspond to high temperature cooling rates, and conversely.

Table I reports experimental findings for Avg DNE under alternative values for α and β. Fig. 9 provides a heat-map depiction of these Avg DNE findings.

An interesting "sweet spot" pattern is immediately evident in Fig. 9: namely, the (α,β) combinations associated with the highest Avg DNE outcomes are along a nonlinear ridge line spanning combinations from (high,high)=(1,100) in the northwest corner to (low,moderate)=(1/24,2) in the south-central region. What causes this nonlinear coupled dependence of Avg DNE on α and β?

A high α value reflecting an optimistically high initial earnings aspiration tends to induce experimentation with alternative action choices due to "disappointment" with the actual net earnings outcomes resulting from early action choices (reflected in large drops in propensity values for these chosen actions). Experimentation can facilitate the eventual discovery of good actions. Conversely, a low α value reflecting a pessimistically low initial earnings aspiration tends to induce premature fixation on an early action choice due to the unexpectedly high earnings outcome resulting from this choice (refelected in a large increase in the propensity value for this chosen action).

Nevertheless, these α effects can be amplified or offset by β effects. A high β value (low T value) amplifies the tendency to premature fixation by amplifying differences in propensity levels across action choices. A moderately low β value can prevent premature fixation by dampening the effects of propensity changes on action choice probabilities. However, a sufficiently low β value results in action choice probability distributions that are essentially uniform across the GenCo's action domain, negating all of the GenCo's efforts to learn which actions result in highest daily net earnings. This deleterious effect is seen in the uniformly low Avg DNE outcomes achieved in Table I and Fig. 9 for the lowest tested β levels 1 and 1/2.

Based on the Avg DNE findings presented in Table I and depicted in Fig. 9, we use the *sweet-spot* VRE learning parameter settings (α, β) = (1,100) in all of the capacity-withholding experiments reported in Section V.[8]

V. Report of Key Findings for Capacity Withholding

A. Benchmark Case: No Physical Capacity Withholding

This subsection presents findings for dynamic 5-bus test case experiments in which the GenCos can exercise economic capacity withholding but not physical capacity withholding. That is, the GenCos can learn over time to report higher-than-true marginal cost functions, but each GenCo i always reports a maximum operating capacity equal to its true maximum operating capacity Cap_i^U.[9]

Table II reports Avg DNE outcomes for each GenCo i, as well as overall Avg DNE, calculated for day D=100. Results are reported for R values ranging from R=0.0 (100% fixed demand) to R=1.0 (100% price-sensitive demand).

Consider, first, the results for R=0.0. Not surprisingly, GenCo 3 attains the highest Avg DNE. As seen from Fig. 6, GenCo 3 is located within a potential "load pocket". Indeed, the branch connecting Bus 1 to Bus 2 is typically congested around the peak demand hour 17, and GenCo 3 exploits the resulting load-pocket opportunity by engaging in substantial economic capacity withholding.

GenCo 5 has the largest operating capacity and it is committed at a higher power output each day, on average, than GenCo 3. However, it does not end up with as high a mark-up over true marginal cost as GenCo 3 and hence attains a lower Avg DNE. GenCo 4 is committed at about half of its operating capacity at a moderate mark-up over true marginal cost, on average, and the Avg DNE of GenCo 4 is approximately the same as for GenCo 5.

Finally, the two smallest-capacity GenCos 1 and 2 are both located at Bus 1, hence they are in direct rivalry with each other. Moreover, the branch connecting Bus 1 to Bus 2 exhibits persistent congestion around the peak demand hour 17, hence GenCos 1 and 2 are partially blocked from servicing the demand at Busses 2, 3, and 4 during this peak demand time. Consequently, GenCos 1 and 2 are committed at relatively low power outputs at relatively low mark-ups over true marginal cost, on average, and both attain relatively low Avg DNEs.

Table II also shows that the individual Avg DNE_i for each GenCo i dramatically declines as R increases from 0.0 to 1.0. However, these declines are at different rates for different GenCos, resulting in changes in their shares in overall Avg DNE.

For example, given R=0.0 (100% fixed demand), GenCo 5's share of Avg DNE is smaller than that of GenCo 3 despite having the largest operating capacity of all GenCos. However, given R=1.0 (100% price-sensitive demand), GenCo 5 has the highest share of Avg DNE of all GenCos; it is now being committed at the highest power level, and this outweighs the fact that GenCo 5 is exercising less economic capacity withholding than GenCo 3.

The underlying reason for these relative changes in fortune is that total demand substantially declines in moving from R=0.0 to R=1.0. All GenCos are forced to compete with each other for the reduced demand. Eventually, all GenCos lose their pivotal supplier status, and any GenCo aggressively engaging in economic capacity withholding risks being undercut by rival supply offers. In particular, GenCos with relatively low true marginal costs are more favored in this environment since higher-cost GenCos could fail to be committed at all.

B. Common Maximum Physical Capacity Withholding Rates

This subsection presents findings for dynamic 5-bus test case experiments in which the GenCos can exercise physical capacity withholding as well as economic capacity withholding. More precisely, the GenCos can report supply offers for which their reported maximum operating capacities are strictly less than their true maximum operating capacities by a percentage no greater than a *maximum shrinkage rate*, either 5% or 10%.[10]

Table III reports GenCo average % capacity shrinkages and % Avg DNE changes (relative to the benchmark case of no physical capacity withholding) calculated for day D=100 when the maximum shrinkage rate is 5%. Results are reported under systematically varied demand conditions ranging from R=0.0 (100% fixed demand) to R=1.0 (100% price-sensitive demand). The results in Table III display several regularities, as follows:

- For each R value, each GenCo's average % capacity shrinkage is well below the maximum shrinkage rate.

[8]For completeness and replicability purposes, we also note here that the following parameter settings for action domain construction were used in the Section V experiments for each GenCo i: M1_i=10; M2_i=10; M3_i=1 (experiments with no physical capacity withholding); M3_i=10 (experiments with physical capacity withholding); RIMin_i^C = 1 (experiments with no physical capacity witholding); RIMin_i^C=0.95 (experiments with a maximum of 5% physical capacity withholding permitted); and RIMin_i^C=0.90 (experiments with a maximum of 10% physical capacity withholding permitted). The cardinality of the action domain for each GenCo i is determined as the product M_i of M1_i, M2_i, and M3_i. See Sun and Tesfatsion [5] for a detailed description of action domain construction for the AMES GenCos.

[9]More precisely, each GenCo i's action domain AD_i consists of 100 possible supply offers. Each possible supply offer is a "marginal cost" function that lies on or above GenCo i's true marginal cost function and that spans GenCo i's true operating capacity interval [Cap^L,Cap^U].

[10]More precisely, each GenCo i's action domain AD_i consists of 1000 possible supply offers. GenCo i can choose from among 10 equally-spaced shrinkage rates s from 0% to the maximum shrinkage rate. Each shrinkage rate s has the form s = 100%·[Cap^U - Cap^{RU}(s)]/[Cap^U - Cap^L], which determines a reported maximum operating capacity Cap^{RU}(s). For each shrinkage rate s there are 100 possible supply offers, each consisting of a "marginal cost" function that lies on or above GenCo i's true marginal cost function and that spans the operating capacity interval [Cap^L,Cap^{RU}(s)].

- For each R value, the five GenCos have similar average % capacity shrinkages.
- For each R value, the % Avg DNE change is positive for some GenCos and negative for others (relative to benchmark).

The latter net earnings finding reflects how extraordinarily difficult it is for individual GenCos operating in dynamic wholesale power markets with multiple rivals to ensure that strategic changes in their reported supply offers indeed result in higher average daily net earnings for themselves.

For example, for R=0.0, GenCo 1, GenCo 2, and GenCo 5 attain higher average daily net earnings (relative to benchmark) while GenCo 3 and GenCo 4 substantially lose ground. Examining the micro data, it is seen that GenCo 1, GenCo 2, and GenCo 5 are being committed on average at somewhat higher power levels (relative to benchmark) while, at the same time, the LMPs at their busses are higher as well (relative to benchmark). In contrast, GenCo 3 and GenCo 4 are being committed on average at somewhat smaller power levels (relative to benchmark) and the LMPs at their busses are much lower (relative to benchmark).

LMPs and power commitments are, of course, *system* outcomes determined by the totality of demand bids and supply offers reported into the day-ahead market in interaction with nonlinear power flow on the grid. They are not under the control of individual GenCos, yet they determine the individual daily net earnings of these GenCos.

The pattern of findings seen in Table III for a 5% maximum shrinkage rate is also seen in Table IV, which repeats the experiments of Table III for a 10% maximum shrinkage rate. In addition, the following regularities are also evident:

- For each R value, each GenCo i's average % capacity shrinkage under a 10% maximum shrinkage rate is almost twice its average % capacity shrinkage under a 5% maximum shrinkage rate.
- Nevertheless, for each R value, each GenCo i's % Avg DNE change is very similar in sign and magnitude no matter which maximum shrinkage rate is in effect.

Given the latter finding, in the next subsection focusing on capacity withholding by a single GenCo we only report results for the case of a 5% maximum shrinkage rate.

C. Physical Capacity Withholding by a Single GenCo

This subsection presents average % capacity shrinkages and % Avg DNE changes (relative to the benchmark no-shrinkage case) calculated for day D=100 in dynamic 5-bus test case experiments in which only a single GenCo engages in capacity shrinkage. The maximum shrinkage rate is fixed at 5%.

Specifically, in Table V the single Genco is the relatively cheap and small GenCo 1. In Table VI the single GenCo is the relatively more expensive and large GenCo 3. And in Table VII the single Genco is the relative cheap yet largest GenCo 5.

Comparing the eighteen R/GenCo cases presented in these three tables, the following intriguing regularities are seen:

- For the set of cases corresponding to any one R value, all of the GenCos engaging in capacity shrinkage have similar average % capacity shrinkages well below the maximum permitted shrinkage rate of 5%; yet their % Avg DNE changes (relative to benchmark) are highly dissimilar.
- For many cases, the single GenCo engaging in capacity shrinkage has a *negative* % Avg DNE change (relative to benchmark). This is particularly true for cases involving GenCo 1 and GenCo 5.
- For almost all cases, at least one GenCo *not* engaging in capacity shrinkage has a *positive* % Avg DNE change (relative to benchmark). This is particularly true when GenCo 1 engages in capacity shrinkage.
- For no case is it true either that *all* GenCos end up having a *positive* % Avg DNE change (relative to benchmark) or that *all* GenCos end up having a *negative* % Avg DNE change (relative to benchmark).

A clear understanding of these results must await a more extensive examination of outcomes at a micro level. Yet one implication seems clear. Capacity withholding has interesting cross-effects that could potentially be exploited by GenCos who own multiple generation plants located at multiple busses.

VI. Concluding Remarks

Restructured wholesale power markets are sequential open-ended games. A careful explanation of the findings presented in Section V will thus require a detailed micro examination of learning behaviors and market interactions over time.

For example, maximum potential total demand (12) in our dynamic 5-bus test case experiments is always less than 90% of the true total operating capacity of the five GenCos. Even during peak demand times when congestion partially blocks relatively cheap generation at Bus 1 from servicing demand at Busses 2, 3, and 4 there is always enough potential operating capacity to satisfy demand. Consequently, it would seem that strategic capacity withholding to induce higher net earnings should not be a serious problem.

What is missing from this high-level analysis, however, is a determination of the pivotal supplier status of different GenCos with regard to fixed demand, meaning that fixed demand *cannot* be met *without* their operating capacity. Pivotal supplier status relative to fixed demand implies substantial opportunities for the exercise of market power through capacity withholding. An additional complicating aspect is that capacity withholding on the part of some GenCos can induce pivotal supplier status (and higher net earnings) for others.

This issue will be addressed in future studies.

Appendix A
VRE Reinforcement Learning

This section describes the implementation of the VRE reinforcement learning algorithm for an arbitrary AMES(V2.01) GenCo i.

Suppose it is the beginning of the initial day D=1. Each GenCo i must choose an action (supply offer) from its action domain AD_i to report to the ISO for the day-ahead market in day D+1, where AD_i consists of M_i possible actions.

The *initial propensity* of GenCo i to choose action $m \in AD_i$ is given by $\mathrm{q}_{im}(1)$ for m = 1,...,M_i. AMES(V2.01) permits the user to set these initial propensity levels to any real numbers. However, the assumption used in this study is that GenCo i's initial propensity levels are all set equal to some common value $q_i(1)$, as follows:

$$q_{im}(1) \quad = \quad q_i(1) \;\; \text{for all supply offers } m \in AD_i \tag{19}$$

Now consider the beginning of any day $D \geq 1$, and suppose the current propensity of GenCo i to choose action m in AD_i is given by $q_{im}(D)$. The *choice probabilities* that GenCo i uses to select an action for day D are then constructed from these propensities as follows:

$$p_{im}(D) = \frac{\exp(q_{im}(D)/T_i)}{\sum_{j=1}^{M_i} \exp(q_{ij}(D)/T_i)} \;,\; m \in AD_i \tag{20}$$

In (20), T_i is a *temperature cooling rate* that affects the degree to which GenCo i makes use of propensity values in determining its choice probabilities. As $T_i \rightarrow \infty$, then $p_{im}(D) \rightarrow 1/M_i$, so that in the limit GenCo i pays no attention to propensity values in forming its choice probabilities. On the other hand, as $T_i \rightarrow 0$, the choice probabilities (20) become increasingly peaked over the particular actions m having the highest propensity values $q_{im}(D)$, thereby increasing the probability that these actions will be chosen.

At the end of day D, the current propensity $q_{im}(D)$ that GenCo i associates with each action m in AD_i is updated in accordance with the following rule. Let m' denote the action that was *actually* selected and reported into the day-ahead market by GenCo i in day D. Also, let $NE_{im'}(D)$ denote the *actual* daily net earnings (10) attained by GenCo i at the end of day D as its settlement payment for all 24 hours of the day-ahead market for day D+1. Then, for each action m in AD_i,

$$q_{im}(D+1) \;\; = \;\; [1-r_i]q_{im}(D) \, + \, Response_{im}(D) \;\; , \tag{21}$$

where

$$Response_{im}(D) = \begin{cases} [1-e_i] \cdot NE_{im'}(D) & \text{if} \;\; m = m' \\ \\ e_i \cdot q_{im}(D)/[M_i - 1] & \text{if} \;\; m \neq m', \end{cases} \tag{22}$$

and $m \neq m'$ implies $M_i \geq 2$. The *recency parameter* r_i in (21) determines the relative weight placed on past versus current daily net earnings payments in the updating of the action choice propensities over time. The *experimentation parameter* e_i in (22) dampens the growth of the chosen action's propensity level. It also permits reinforcement to spill over to some extent from a chosen action to other actions to encourage continued experimentation with a wide variety of actions in the early stages of the learning process.

ACKNOWLEDGMENT

The authors thank Chen-Ching Liu, Junjie Sun, and members of the ISU Electric Energy Economics (E3) Group for many useful discussions related to the topic of this study.

REFERENCES

[1] FERC, *Notice of White Paper*, U.S. Federal Energy Regulatory Commission, April 2003.

[2] P. Joskow, "Markets for power in the united states: An interim assessment," *The Energy Journal*, vol. 27, no. 1, pp. 1–36, 2006.

[3] L. Tesfatsion, "Agent-Based Computational Economics Research on Restructured Electricity Markets," hosted by the Economics Department, Iowa State University, www.econ.iastate.edu/tesfatsi/aelect.htm

[4] H. Li, J. Sun, and L. Tesfatsion, "Dynamic LMP Reponse Under Alternative Price-Cap and Price-Sensitive Demand Scenarios," *Proceedings*, IEEE Power Engineering Society General Meeting, Pittsburgh, PA, July 2008.

[5] J. Sun and L. Tesfatsion, "Dynamic testing of wholesale power market designs: An open-source agent-based framework," *Computational Economics*, Vol. 30, pp. 291–327, 2007. Working paper available: www.econ.iastate.edu/tesfatsi/DynTestAMES.JSLT.pdf

[6] J. Sun and L. Tesfatsion, "DC optimal power flow formulation and solution using QuadProgJ," *Proceedings*, IEEE Power and Energy Society General Meeting, Tampa, Florida, June 2007. Working paper available: www.econ.iastate.edu/tesfatsi/DC-OPF.JSLT.pdf

[7] A. E. Roth and E. Ido, "Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term," *Games and Econ. Behavior*, Vol. 8, pp. 164–212, 1995.

[8] I. Erev and A. E. Roth, "Predicting how people play games with unique mixed-strategy equilibria," *American Economic Review*, Vol. 88, pp. 848–881, 1998.

[9] L. Tesfatsion, "AMES Wholesale Power Market Test Bed Homepage," hosted by the Economics Department, Iowa State University, http://www.econ.iastate.edu/tesfatsi/AMESMarketHome.htm

[10] A. Tellidou and A. Bakirtzis, "Agent-Based Analysis of Capacity Withholding and Tacit Collusion in Electricity Markets," *IEEE Transactions on Power Systems*, Vol. 22, No. 4, November, pp. 1735–1742, 2007.

[11] MISO, *Homepage*, Midwest ISO, Inc., 2008. [Online]. Available: www.midwestiso.org/

[12] H. Li, J. Sun, and L. Tesfatsion, "Separation and Volatility of Locational Marginal Prices in Restructured Wholesale Power Markets," Working Paper, Economics Department, Iowa State University, 2008, in progress.

[13] IEEE Task Force on Open-Source Software for Power Systems, ewh.ieee.org/cmte/psace/CAMS_taskforce/software.htm

[14] J. Lally, "Financial transmission rights: Auction example," in *Financial Transmission Rights Draft 01-10-02*, m-06 ed, ISO New England, Inc., January 2002, section 6.

[15] M. Shahidehpour, H. Yamin, and Z. Li, *Market Operations in Electric Power Systems*. New York, NY: IEEE/Wiley-Interscience, John Wiley & Sons, Inc., 2002.

[16] M. Pentapalli, "A comparative study of Roth-Erev and modified Roth-Erev reinforcement learning algorithms for uniform-price double auctions," M.S. Thesis, Computer Science Department, Iowa State University, March 2008.

Hongyan Li received his M.S. degree in Electric Power Systems from Xian Jiaotong University (China) in 1997. He served as an engineer for the Sifang & Huaneng Power System Control Corporation (China) for seven years. He is currently pursuing a Ph.D. degree in Electrical and Computer Engineering at Iowa State University. His principal research area is power system economics, with a particular focus on the efficiency and reliability of restructured wholesale power markets.

Leigh Tesfatsion received her Ph.D. degree in Economics from the University of Minnesota in 1975. She is currently Professor of Economics and Mathematics at Iowa State University. Her principal research area is Agent-based Computational Economics (ACE), the computational study of economic processes modeled as dynamic systems of interacting agents, with a particular focus on restructured electricity markets. She is an active participant in IEEE Power Engineering Society working groups and task forces focusing on power economics issues and a co-organizer of the ISU Electric Energy Economics (E3) Group. She serves as associate editor for a number of journals, including the *Journal of Energy Markets*.

TABLE I

AVERAGE GENCO DAILY NET EARNINGS FOR THE FINAL DAY UNDER ALTERNATIVE GENCO VRE LEARNING PARAMETER SPECIFICATIONS (α, β) WITH NO PRICE-SENSITIVE DEMAND AND NO PHYSICAL CAPACITY WITHHOLDING.

		beta=100	beta=50	beta=10	beta=2	beta=1	beta=1/2
alpha=1	GenCo 1	69,219.61	5,578.19	19,786.27	19,825.24	19,825.24	19,863.80
	GenCo 2	54,548.72	3,040.47	12,299.14	11,765.93	11,765.93	11,765.93
	GenCo 3	1,725,216.72	1,765,140.21	529,014.82	548,883.32	548,883.32	547,658.35
	GenCo 4	321,907.08	196,769.51	31,510.57	29,790.04	29,790.04	29,762.60
	GenCo 5	270,754.58	187,954.06	190,968.98	189,378.40	189,378.40	189,396.96
	Avg DNE	2,441,646.71	2,158,482.44	783,579.77	799,642.93	799,642.93	798,447.64
	St. Dev.	(558,896.97)	(730,657.92)	(399,309.98)	(433,253.98)	(433,253.98)	(434,433.11)
alpha=1/2	GenCo 1	79,875.97	74,182.72	20,959.35	20,114.75	19,825.24	19,825.24
	GenCo 2	64,817.10	61,235.66	14,366.78	12,241.57	11,765.93	11,765.93
	GenCo 3	1,462,304.20	1,737,816.84	537,044.33	520,518.36	548,883.32	548,883.32
	GenCo 4	306,198.90	337,814.49	32,397.27	29,790.04	29,790.04	29,790.04
	GenCo 5	276,640.75	280,020.65	192,046.57	189,168.30	189,378.40	189,378.40
	Avg DNE	2,189,836.92	2,491,070.36	796,814.29	771,833.01	799,642.93	799,642.93
	St. Dev.	(534,136.31)	(496,068.85)	(400,651.49)	(400,321.80)	(433,253.98)	(433,253.98)
alpha=1/4	GenCo 1	87,629.74	79,100.46	14,920.20	20,187.91	20,114.75	19,825.24
	GenCo 2	76,471.25	65,279.31	9,170.37	12,323.55	12,241.57	11,765.93
	GenCo 3	1,115,033.21	1,328,446.25	1,074,869.72	525,030.40	520,518.36	548,883.32
	GenCo 4	258,044.34	305,601.87	95,151.47	30,809.31	29,790.04	29,790.04
	GenCo 5	256,589.11	270,324.35	188,384.38	190,457.78	189,168.30	189,378.40
	Avg DNE	1,793,767.65	2,048,752.24	1,382,496.14	778,808.95	771,833.01	799,642.93
	St. Dev.	(529,846.55)	(610,971.13)	(920,990.49)	(398,145.02)	(400,321.80)	(433,253.98)
alpha=1/10	GenCo 1	50,093.01	78,026.20	74,886.81	20,959.35	19,786.27	20,114.75
	GenCo 2	47,977.10	69,290.98	60,364.79	14,366.78	12,299.14	12,241.57
	GenCo 3	767,282.28	1,042,911.66	1,662,257.14	537,044.33	529,014.82	522,830.32
	GenCo 4	153,075.18	225,113.47	318,609.44	32,397.27	31,510.57	29,790.04
	GenCo 5	182,152.53	235,383.32	274,076.15	192,046.57	190,968.98	189,168.30
	Avg DNE	1,200,580.10	1,650,725.62	2,390,194.34	796,814.29	783,579.77	774,144.97
	St. Dev.	(510,232.72)	(665,255.14)	(561,884.41)	(400,651.49)	(399,309.98)	(398,842.95)
alpha=1/24	GenCo 1	37,197.65	53,395.14	79,422.74	9,329.03	22,190.32	20,787.98
	GenCo 2	38,089.68	50,074.97	65,366.01	5,317.87	14,854.24	12,330.86
	GenCo 3	635,691.68	682,930.40	1,178,427.61	1,615,272.93	549,196.62	528,895.76
	GenCo 4	83,253.22	130,439.19	249,815.58	184,160.28	32,701.84	31,621.09
	GenCo 5	183,685.63	193,689.29	248,395.06	193,984.82	192,528.20	192,354.51
	Avg DNE	977,917.86	1,110,528.99	1,821,427.00	2,008,064.92	811,471.22	785,990.18
	St. Dev.	(403,198.74)	(454,288.77)	(549,146.76)	(862,134.72)	(399,997.12)	(397,810.42)

TABLE II

BENCHMARK CASE: AVERAGE GENCO DAILY NET EARNINGS FOR DAY 100 WHEN ONLY ECONOMIC CAPACITY WITHHOLDING IS PERMITTED

GenCo	R=0.0	R=0.2	R=0.4	R=0.6	R=0.8	R=1.0
1	20,196.11	25,023.98	18,747.53	15,805.06	10,531.84	4,854.98
2	16,366.03	21,665.28	15,385.46	13,475.27	11,314.75	5,324.56
3	1,237,976.82	176,717.76	51,536.59	17,695.26	7,630.61	3,398.80
4	165,133.57	11,746.23	2,523.19	534.78	135.97	66.88
5	188,646.36	179,072.71	147,584.54	114,626.37	81,038.34	54,220.06
Avg DNE	1,628,318.89	414,225.96	235,777.31	162,136.73	110,651.51	67,865.28
St. Dev.	(878,152.41)	(195,242.12)	(116,151.45)	(67,832.44)	(30,366.66)	(17,494.45)

TABLE III

CAPACITY SHRINKAGE CASE 1 (5% MAXIMUM): AVERAGE GENCO % CAPACITY SHRINKAGE AND % DAILY NET EARNINGS CHANGES FOR DAY 100 RELATIVE TO THE BENCHMARK CASE

	R=0.0		R=0.2		R=0.4		R=0.6		R=0.8		R=1.0	
GenCo	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE
1	-0.023 (0.014)	0.124	-0.023 (0.015)	-0.109	-0.023 (0.014)	0.050	-0.024 (0.014)	-0.097	-0.025 (0.014)	-0.119	-0.026 (0.013)	-0.026
2	-0.017 (0.013)	0.333	-0.017 (0.013)	-0.021	-0.017 (0.014)	0.272	-0.017 (0.013)	0.140	-0.017 (0.013)	-0.059	-0.018 (0.013)	0.001
3	-0.022 (0.013)	-0.577	-0.021 (0.013)	0.033	-0.019 (0.013)	-0.006	-0.019 (0.013)	-0.043	-0.019 (0.013)	0.038	-0.019 (0.013)	0.012
4	-0.021 (0.013)	-0.784	-0.020 (0.013)	-0.073	-0.021 (0.013)	0.323	-0.020 (0.013)	0.092	-0.020 (0.013)	0.021	-0.020 (0.013)	0.022
5	-0.022 (0.013)	0.107	-0.021 (0.012)	0.035	-0.020 (0.012)	0.060	-0.020 (0.013)	0.002	-0.020 (0.013)	-0.019	-0.020 (0.013)	-0.028

TABLE IV
CAPACITY SHRINKAGE CASE 2 (10% MAXIMUM): AVERAGE GENCO % CAPACITY SHRINKAGE AND % DAILY NET EARNINGS CHANGES FOR DAY 100 RELATIVE TO THE BENCHMARK CASE

	R=0.0		R=0.2		R=0.4		R=0.6		R=0.8		R=1.0	
GenCo	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE
1	-0.046 (0.028)	0.124	-0.046 (0.029)	-0.101	-0.046 (0.028)	0.050	-0.049 (0.028)	-0.101	-0.050 (0.027)	-0.126	-0.052 (0.026)	-0.032
2	-0.034 (0.027)	0.327	-0.033 (0.027)	-0.015	-0.034 (0.027)	0.275	-0.034 (0.026)	0.137	-0.035 (0.027)	-0.066	-0.037 (0.025)	-0.002
3	-0.042 (0.028)	-0.566	-0.043 (0.026)	0.043	-0.038 (0.026)	0.007	-0.037 (0.025)	-0.019	-0.037 (0.025)	0.051	-0.037 (0.025)	0.021
4	-0.041 (0.026)	-0.778	-0.041 (0.025)	-0.038	-0.043 (0.026)	0.366	-0.040 (0.026)	0.142	-0.040 (0.026)	0.040	-0.040 (0.026)	0.043
5	-0.044 (0.026)	0.120	-0.042 (0.025)	0.081	-0.040 (0.024)	0.071	-0.041 (0.025)	0.007	-0.040 (0.027)	-0.017	-0.040 (0.026)	-0.026

TABLE V
CAPACITY SHRINKAGE CASE 3 (5% MAXIMUM, GENCO 1 ONLY): GENCO 1'S AVERAGE % CAPACITY SHRINKAGE AND ALL GENCOS' AVERAGE % DAILY NET EARNINGS CHANGES FOR DAY 100 RELATIVE TO THE BENCHMARK CASE

	R=0.0		R=0.2		R=0.4		R=0.6		R=0.8		R=1.0	
GenCo	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE
1	-0.025 (0.013)	-0.061	-0.022 (0.014)	-0.093	-0.023 (0.014)	-0.100	-0.024 (0.014)	-0.072	-0.025 (0.013)	-0.097	-0.026 (0.013)	0.001
2	0.000 (0.000)	-0.012	0.000 (0.000)	0.017	0.000 (0.000)	0.016	0.000 (0.000)	0.014	0.000 (0.000)	-0.004	0.000 (0.000)	-0.007
3	0.000 (0.000)	0.012	0.000 (0.000)	0.011	0.000 (0.000)	0.010	0.000 (0.000)	0.009	0.000 (0.000)	0.011	0.000 (0.000)	0.003
4	0.000 (0.000)	0.016	0.000 (0.000)	0.007	0.000 (0.000)	0.015	0.000 (0.000)	0.022	0.000 (0.000)	0.018	0.000 (0.000)	0.006
5	0.000 (0.000)	0.019	0.000 (0.000)	0.000	0.000 (0.000)	0.010	0.000 (0.000)	0.011	0.000 (0.000)	-0.002	0.000 (0.000)	-0.003

TABLE VI
CAPACITY SHRINKAGE CASE 4 (5% MAXIMUM, GENCO 3 ONLY): GENCO 3'S AVERAGE % CAPACITY SHRINKAGE AND ALL GENCOS' AVERAGE % DAILY NET EARNINGS CHANGES FOR DAY 100 RELATIVE TO THE BENCHMARK CASE

	R=0.0		R=0.2		R=0.4		R=0.6		R=0.8		R=1.0	
GenCo	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE
1	0.000 (0.000)	0.181	0.000 (0.000)	-0.054	0.000 (0.000)	0.006	0.000 (0.000)	0.020	0.000 (0.000)	0.000	0.000 (0.000)	0.000
2	0.000 (0.000)	0.181	0.000 (0.000)	-0.094	0.000 (0.000)	0.001	0.000 (0.000)	-0.005	0.000 (0.000)	0.000	0.000 (0.000)	0.000
3	-0.023 (0.013)	-0.573	-0.021 (0.013)	0.052	-0.019 (0.013)	0.004	-0.019 (0.013)	0.011	-0.019 (0.013)	0.013	-0.019 (0.013)	-0.003
4	0.000 (0.000)	-0.778	0.000 (0.000)	-0.063	0.000 (0.000)	0.010	0.000 (0.000)	0.011	0.000 (0.000)	0.010	0.000 (0.000)	0.008
5	0.000 (0.000)	-0.005	0.000 (0.000)	-0.014	0.000 (0.000)	0.005	0.000 (0.000)	-0.002	0.000 (0.000)	0.000	0.000 (0.000)	0.000

TABLE VII
CAPACITY SHRINKAGE CASE 5 (5% MAXIMUM, GENCO 5 ONLY): GENCO 5'S AVERAGE % CAPACITY SHRINKAGE AND ALL GENCOS' AVERAGE % DAILY NET EARNINGS CHANGES FOR DAY 100 RELATIVE TO THE BENCHMARK CASE

	R=0.0		R=0.2		R=0.4		R=0.6		R=0.8		R=1.0	
GenCo	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE	% ΔCap	% ΔNE
1	0.000 (0.000)	-0.103	0.000 (0.000)	-0.153	0.000 (0.000)	0.172	0.000 (0.000)	0.003	0.000 (0.000)	-0.079	0.000 (0.000)	-0.004
2	0.000 (0.000)	0.024	0.000 (0.000)	-0.169	0.000 (0.000)	0.106	0.000 (0.000)	-0.039	0.000 (0.000)	-0.060	0.000 (0.000)	0.062
3	0.000 (0.000)	0.008	0.000 (0.000)	-0.015	0.000 (0.000)	-0.027	0.000 (0.000)	-0.022	0.000 (0.000)	-0.004	0.000 (0.000)	0.009
4	0.000 (0.000)	-0.005	0.000 (0.000)	0.073	0.000 (0.000)	0.457	0.000 (0.000)	0.023	0.000 (0.000)	0.019	0.000 (0.000)	0.006
5	-0.022 (0.013)	0.086	-0.021 (0.012)	-0.011	-0.021 (0.012)	0.047	-0.020 (0.012)	0.000	-0.020 (0.013)	-0.018	-0.020 (0.013)	-0.030