Separation and Volatility of Locational Marginal Prices in Restructured Wholesale Power Markets

Hongyan Li, *Student Member, IEEE*, Junjie Sun, and Leigh Tesfatsion, *Member, IEEE*

***Abstract*—This study uses an agent-based test bed ("AMES") to investigate separation and volatility of locational marginal prices (LMPs) in an ISO-managed restructured wholesale power market operating over an AC transmission grid. Particular attention is focused on the dynamic and cross-sectional response of LMPs to systematic changes in demand-bid price sensitivities and supply-offer price cap levels under varied learning specifications for the generation companies. Also explored is the extent to which the supply offers of the marginal (price-determining) generation companies induce correlations among neighboring LMPs.**

***Index Terms*—Restructured wholesale power markets, agent-based modeling,, locational marginal prices (LMPs), LMP separation, LMP volatility, multi-agent learning, demand-bid price sensitivity, supply-offer price caps, AMES Wholesale Power Market Test Bed**

I. Introduction

THE wholesale power market design proposed by the U.S. Federal Energy Regulatory Commission (FERC) in an April 2003 white paper [1] recommends that *locational marginal prices (LMPs)* be used for the management of transmission grid congestion. Under locational marginal pricing, the solution of suitably formulated optimal power flow problems at successive time intervals is used to price electric power in accordance with both the location and timing of its injection into, or withdrawal from, the transmission grid.

Locational marginal pricing has now been implemented – or scheduled for implementation – in U.S. energy regions in the midwest (MISO), New England (ISO-NE), New York (NYISO), the mid-atlantic states (PJM), California (CAISO), the southwest (SPP), and Texas (ERCOT) [2]. Consequently, it is critical for market operators in these regions to understand how LMPs respond under alternative structural conditions, institutional arrangements, and behavioral dispositions of participant traders.

Previous studies have derived analytical expressions for LMPs at a point in time, conditional on given grid, demand, and supply conditions; see, for example, [3] and [4]. However, only recently have researchers begun to pay attention to the dynamic response of LMP solution paths to changed circumstances, particularly when traders have learning capabilities permitting them to strategically adjust their trade behaviors over time.

Latest revision: Revised 2 March 2010. This work has been supported in part by the National Science Foundation under Grant NSF-0527460 and by grants from the ISU Electric Power Research Center.

Hongyan Li (lihy@iastate.edu), Electrical and Computer Engineering Department, Iowa State University, Ames, IA 50011 USA; Junjie Sun (Junjie.Sun@occ.treas.gov), OCC, U.S. Treasury, Washington D.C., and Leigh Tesfatsion (corresponding author: tesfatsi@iastate.edu), Economics Department, Iowa State University, Ames, IA 50011 USA.

For example, Sueyoshi and Tadiparthi [5] use an agent-based test bed to examine price response under alternative transmission line limit conditions for a wholesale power market separated into multiple zones (collections of wholesale power sellers and buyers). Prices separate across any two zones functionally disconnected by a binding constraint on their intertie line. Sellers and buyers use reinforcement learning to determine their supply offers and demand bids (price-quantity pairs) for day-ahead and real-time markets. One of the key experimental findings of the authors is that the average level and volatility of day-ahead prices both increase as the number of capacity-limited intertie lines is systematically increased.

Sun and Tesfatsion [6] report findings for several dynamic 5-bus test case experiments to illustrate the capabilities of AMES(V1.31), an agent-based test bed for restructured wholesale power markets.[1] AMES(V1.31) models a wholesale power market with *load-serving entities (LSEs)* and *generation companies (GenCos)* for which the GenCos have learning capabilities. The market operates over an AC transmission grid in accordance with core features of the market design proposed by the U.S. FERC [1]. The market operator uses LSE fixed demands (loads) and GenCo supply offers as input data for DC optimal power flow problems to determine hourly dispatch and LMP levels for the day-ahead market. A key finding of the authors is that the GenCos quickly learn to tacitly collude on supply offers resulting in LMPs substantially higher than competitive price levels.

Li et al. [7] report preliminary experimental findings to demonstrate the capabilities of AMES(V2.01), an extension of AMES(V1.31). This extended version permits LSEs to submit price-sensitive as well as fixed demand bids for the day-ahead market.[2] In addition, the market operator is permitted to impose a price cap on the supply offers submitted by GenCos for the day-ahead market in an attempt to mitigate their ability to exercise market power.[3]

Building on Li et al. [7], this study uses AMES (V2.02) to

[1] AMES is an acronym for *A*gent-based *M*odeling of *E*lectricity *S*ystems. AMES(V1.31) was formally released as open-source software at the IEEE Power and Energy Society General Meeting in June 2007. Downloads, manuals, and tutorial information for all AMES version releases to date can be accessed at the AMES homepage [8].

[2] Although price-sensitive demand bids are permitted in U.S. restructured wholesale power markets operating under the FERC market design, most demand is still in the form of price-insensitive loads. For example, the actual ratio of cleared price-sensitive demand to cleared fixed demand in the MISO [9] is currently only about 1%.

[3] Price-cap policies differ widely across U.S. restructured wholesale power markets. For example, MISO [9] currently imposes a price cap on supply offers only under extreme conditions. Consequently, this price cap is more of a "damage control" device than a device for controlling market power.

undertake a comprehensive and systematic investigation of the effects of changes in learning parameters, demand-bid price sensitivities, and supply-offer price caps on LMP separation and volatility over time. Each GenCo uses stochastic reinforcement learning to adaptively choose its supply offers on the basis of its past net earnings outcomes. Careful attention is paid both to dynamic market performance effects and to spatial cross-correlation effects. The primary objective is to gain a more fundamental understanding of how learning, network externalities, and GenCo pivotal and marginal supplier status interact to determine the distribution of LMPs both across the grid (separation) and over time (volatility).

As usual for agent-based models, AMES is most easily explained in general descriptive terms (verbal summaries, flow diagrams, and pseudo-code) as an accompaniment to the actual source code available at the AMES homepage [8].[4] This general description is provided in Section II.

Section III explains the AMES experimental design used to explore dynamic market performance under systematically varied settings for the following three treatment factors: (i) GenCo learning (absent, or present with different learning parameter settings); (ii) the degree to which LSE demand bids are price sensitive (0 to 100%); and (iii) the level of the supply-offer price cap (infinite, high, moderate, or low relative to average peak-hour LMP). Experimental findings for dynamic market performance are reported in Section IV. Section V presents experimental findings for spatial cross-correlations, including GenCo-LMP cross-correlations, LMP-LMP cross-correlations, and comparisons with empirical LMP-LMP cross-correlations determined from MISO price data.

Concluding remarks are given in Section VI. Appendices A through D provide additional discussion of technical details: namely, the market operator's DC-OPF problem formulation; the construction of supply-offer action domains for the GenCos; the learning method used by GenCos to choose their supply offers from their action domains; and the precise construction of the measures used to characterize dynamic market performance.

As an aid to the reader, annotated listings of the admissible exogenous variables and endogenous variables used in subsequent sections of this study are provided in Table I and Table II, respectively.

II. The AMES Test Bed (Version 2.02)

A. Overview

AMES(V2.02) incorporates, in simplified form, core features of the wholesale power market design proposed by the U.S. FERC [1]; see Figure 1. A detailed description of many of these features can be found in Sun and Tesfatsion ([6], [15]). Below is a summary description of the logical flow of daily market events:[5]

[4]Additional materials provided at the AMES homepage [8] include instructions for downloading the source code, instructions for setting up AMES as a standard Java package in an integrated development environment (e.g., NetBeans), a user manual, tutorials, and research publications.

[5]Interested readers might also wish to view ref. [12], which recasts AMES in more standard state-space equation form. The result is a highly nonlinear and highly coupled system of first-order stochastic difference equations describing an open-ended dynamic game.

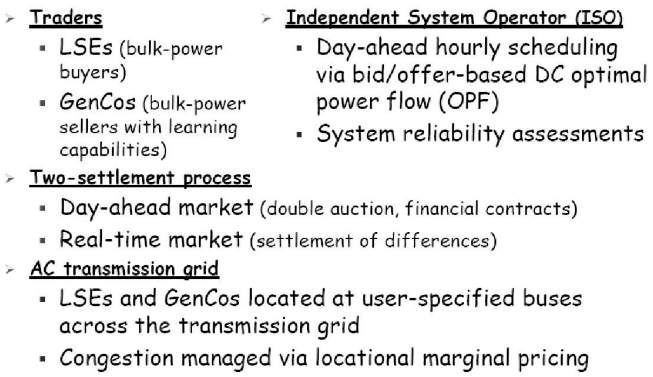

Fig. 1. AMES test bed architecture.

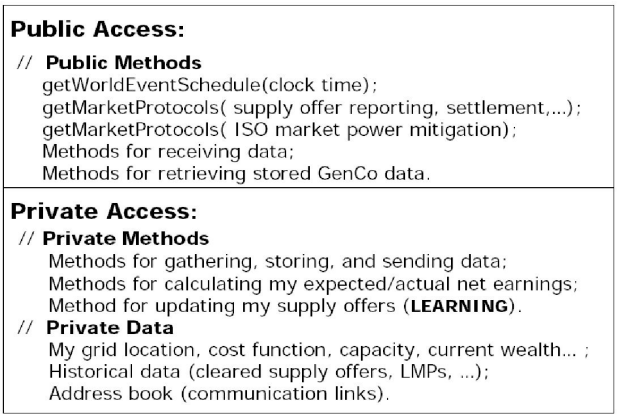

Fig. 2. AMES GenCo: A cognitive agent with learning capabilities.

- 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 buses 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 surplus 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)*, i.e., the determination of prices for electric power in accordance with both the location and timing of its injection into, or

withdrawal from, the transmission grid.[6]

- The objective of each LSE is to secure for itself the highest possible net earnings each day through the purchase of power in the day-ahead market and the resale of this power to its downstream (retail) customers.
- During the morning of 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 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 through the sale of power in the day-ahead market.
- 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.[7] 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 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 dispatch and LMP levels 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 of the LSE and GenCo payment obligations 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 on each day D for the day-ahead market for day D+1 are carried out as planned; traders have no need to engage in real-time 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.

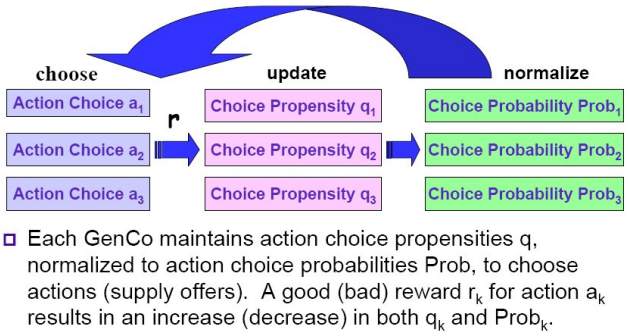

Fig. 3. AMES GenCos use stochastic reinforcement learning to determine the supply offers they report to the ISO for the day-ahead market.

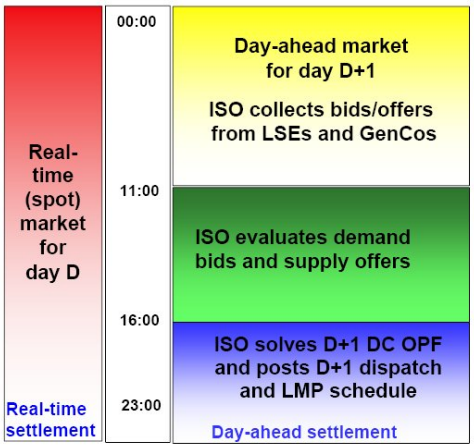

Fig. 4. AMES ISO activities during a typical day D.

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

For each day D, the demand bid reported by LSE j for each hour H of the day-ahead market in day D+1 consists of a *fixed demand bid* $\text{p}^F_{Lj}(\text{H})$ (MW) and a *price-sensitive demand bid*

[6]Roughly stated, a *locational marginal price (LMP)* at any particular transmission grid bus k during any particular time period T is the least cost to the system of servicing demand for one additional megawatt (MW) of power at bus k during period T. See Liu et al. [10] for a careful discussion of LMP derivation from optimal power flow solutions.

[7]In the MISO [9], GenCos each day are actually permitted to report a *separate* supply offer for each hour of the day-ahead market. In order to simplify the learning problem for GenCos, the current version of AMES restricts GenCos to the daily reporting of only one supply offer for the day-ahead market. Interestingly, the latter restriction is imposed on GenCos by the ISO-NE [11] in its particular implementation of FERC's market design. Baldick and Hogan [13, pp. 18-20] conjecture that imposing such limits on the ability of GenCos to report distinct hourly supply offers could reduce their ability to exercise market power.

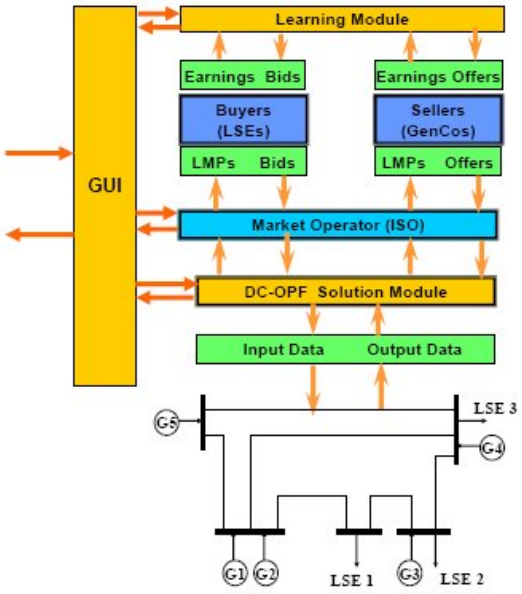

Fig. 5. Illustration of AMES dynamics on a typical day D in the absence of system disturbances or shocks for the special case of a 5-bus grid.

function

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

defined over a *true purchase capacity interval*

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

for real power p^S_{Lj}(H). The expression $D_{jH}(p^S_{Lj}(H))$ denotes LSE j's *true purchase reservation value* for energy evaluated at p^S_{Lj}(H), i.e., the maximum dollar amount it is truly willing to pay per MWh.

Also, for each day D, the single *supply offer* reported by GenCo i for use in each hour H of the day-ahead market for day D+1 consists of a *reported marginal cost function*

$$MC^R_i(p_{Gi}) \;=\; a^R_i + 2b^R_i p_{Gi} \quad (\$/MWh) \qquad (3)$$

defined over a *reported operating capacity interval*

$$Cap^L_i \;\leq\; p_{Gi} \;\leq\; Cap^{RU}_i \quad (MW) \qquad (4)$$

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

To avoid operating at a point where true incremental cost exceeds payment received for its last supplied MW of power, GenCo i's reported marginal cost functions always lie on or above its *true marginal cost function*

$$MC_i(p_{Gi}) \;=\; a_i + 2b_i p_{Gi} \quad (\$/MWh) \;\; . \qquad (5)$$

Also, to avoid infeasible dispatch levels, GenCo i always reports an upper operating capacity level Cap^{RU}_i that lies within GenCo i's *true operating capacity interval*

$$Cap^L_i \;\leq\; p_{Gi} \;\leq\; Cap^U_i \quad (MW) \;\; . \qquad (6)$$

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. Costs, Profits, and Net Earnings

At the beginning of any planning period, the *avoidable costs* of a GenCo refer to the production costs that the GenCo *can* avoid incurring during the period by shutting down, by resale of purchased assets, or by other actions. Conversely, the *sunk costs* of the GenCo refer to the production costs that the GenCo *cannot* avoid incurring during the period because of irrevocable commitments, lack of asset resale value, or other circumstances. *Total costs* refer to the sum of the two.

For the specific context at hand, it is assumed that GenCos do not have any avoidable fixed costs. Thus, the *true avoidable cost function* for GenCo i for any hour H is simply the integral of its marginal cost function, as follows:

$$VCost_i(p_{Gi}) = \int_0^{p_{Gi}} MC_i(p)dp = a_i p_{Gi} + b_i[p_{Gi}]^2 \;\; (\$/h), \qquad (7)$$

The *true total cost function* for GenCo i for any hour H then takes the form

$$TC_i(p_{Gi}) \;=\; [VCost_i(p_{Gi}) + SCost_i] \quad (\$/h), \qquad (8)$$

where p_{Gi} (in MWs) denotes any feasible real-power generation level for GenCo i in hour H and $SCost_i$ (\$/h) denotes GenCo i's pro-rated sunk costs for hour H.

Profit is defined as revenues minus true total costs. On the other hand, *net earnings* are defined as revenues minus true total *avoidable* costs. Suppose, in particular, that GenCo i is located at bus $k(i)$ and is dispatched 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 profit of GenCo i for hour H of day D+1, incurred at the end of day D, is given by

$$\pi_i(H,D) \;=\; LMP_{k(i)} * p_{Gi} - TC_i(p_{Gi}) \quad (\$/h). \qquad (9)$$

On the other hand, the net earnings of GenCo i for hour H of day D+1, incurred at the end of day D, are given by

$$NE_i(H,D) \;=\; LMP_{k(i)} * p_{Gi} - VCost_i(p_{Gi}) \quad (\$/h). \quad (10)$$

The net earnings of GenCo i over all 24 hours of day D+1, incurred at the end of day D, are then given by

$$NE_i(D) \;=\; \sum_{H=00}^{H=23} NE_i(H,D) \quad (\$/h). \qquad (11)$$

As will be seen in Section III, we make use of estimates $MaxDNE_i$ for each 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,[8]

$$MaxDNE_i \;=\; 24 * \left(\max_{s^R_i \in AD_i} \left[HNE(s^R_i) \right] \right) \quad (\$/h), \quad (12)$$

[8]Compare (12) with definition (11) 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 dispatch 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.

where $s_i^R = (a_i^R, b_i^R, Cap_i^{RU})$ denotes a generic supply offer in its action domain AD_i and the hourly net earnings function $HNE(s_i^R)$ ($/h) is given by

$$HNE(s_i^R) = MC_i^R(Cap_i^{RU}) * Cap_i^{RU} - VCost_i(Cap_i^{RU}) \ . \quad (13)$$

D. Supply-Offer Price Cap

As explained more carefully in Appendix A, the goal of the ISO is the reliable attainment of appropriately constrained *operational efficiency* for the wholesale power market. That is, the ISO attempts to maximize the total net surplus accruing to LSEs and GenCos from hourly bulk power trades subject to various transmission and generation constraints.

The ISO is concerned about loss of operational efficiency due to the possible exercise of "market power" by GenCos through strategic reporting of supply offers. Specifically, a GenCo has *market power* if the GenCo can use capacity withholding to increase its net earnings. *Capacity withholding* can take two possible forms: *economic withholding*, i.e., reporting a higher-than-true marginal cost function; and *physical withholding*, i.e., reporting a less-than-true upper operating capacity limit. As one possible approach to GenCo market power mitigation, the ISO can impose a *supply-offer price cap (PCap)*. Under such a policy, the maximum sale reservation value $MC_i^R(Cap_i^{RU})$ reported by any GenCo i cannot exceed PCap.

E. Determination of Dispatch and LMP Levels

As detailed in Appendix A, the ISO computes hourly power dispatch levels and LMPs 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 ISO makes repeated calls to *DCOPFJ*, an accurate and efficient DC-OPF solver developed by Sun and Tesfatsion ([14], [15]). DCOPFJ consists of a strictly convex quadratic programming solver wrapped in an outer SI-pu data conversion shell.

F. GenCo Action Domain Construction

The construction of the *action domain* (supply offer choice set) AD_i for each GenCo i is a critical modeling issue. Empirical sensibility suggests these action domains should permit flexible choice from among a wide range of possible supply offers, and that the degree of flexibility should be roughly similar across the GenCos. On the other hand, computational practicality suggests the number of supply offers included in each action domain should not be unduly large.

Appendix B briefly summarizes how action domains have been constructed for the GenCos in accordance with these objectives. A rigorous detailed discussion and illustration of action domain construction for the AMES GenCos can be found in Sun and Tesfatsion [6, Appendix].

G. GenCo Learning

The essential idea of stochastic reinforcement learning is that the probability of choosing an action should be increased (reinforced) if the corresponding reward is relatively good and decreased if the corresponding reward is relatively poor. As detailed in Appendix C, each GenCo determines its supply offers by means of *VRE reinforcement learning (VRE-RL)*, a variant of a stochastic reinforcement learning algorithm developed by Alvin Roth and Ido Erev ([16], [17]) on the basis of human-subject experiments.

Each GenCo's learning is implemented by means of a Java reinforcement learning module, *JReLM*, developed by Gieseler [18]. The user can tailor the settings of each GenCo's learning parameters to its situation, in particular to its cost attributes, its operating capacity, and its anticipated net earnings.

H. Graphical User Interface

AMES has a graphical user interface (GUI) with separate screens for carrying out the following functions: (a) creation, modification, analysis, and storage of case studies; (b) initialization and editing of the structural attributes of the transmission grid; (c) initialization and editing of the structural attributes of LSEs and GenCos; (d) specification of learning parameter settings for GenCos; (e) specification of simulation controls (e.g., the simulation stopping rule); and (f) customization of table and chart output displays.

III. Experimental Design

The experimental design for this study is based on the *benchmark dynamic 5-bus test case* presented in Table III. This benchmark case is characterized by the following structural, institutional, and behavioral conditions:

- The wholesale power market operates over a 5-bus transmission grid as depicted in Fig. 6, with branch reactances, locations of LSEs and GenCos, and initial hour-0 LSE fixed demand levels adopted from Lally [19].[9]
- True GenCo cost and capacity attributes are as depicted in Figure 7. GenCos range from GenCo 5, a relatively large coal-fired baseload unit with low marginal operating costs, to GenCo 4, a relatively small gas-fired peaking unit with relatively high marginal operating costs.
- Demand is 100% fixed (no price sensitivity) with LSE daily fixed demand profiles adopted from a case study presented in Shahidehpour et al. [20, p. 296-297]; see Fig. 8. Hourly load varies from light (hour 4:00) to peak (hour 17:00), which systematically affects the market power potential of the GenCos; see Fig. 9.
- GenCos are non-learners, meaning they report supply offers to the ISO that convey their true marginal cost functions (5) and true operating capacity limits (6).
- There is no supply-offer price cap.

[9]Lally's transmission grid configuration is now used extensively in ISO-NE/PJM training manuals to derive DC-OPF solutions at a given point in time. An implicit assumption in these derivations is that the ISO knows the true structural attributes of the LSEs and GenCos. No mention is made of the possibility that LSEs and GenCos in real-world ISO-managed wholesale power markets might learn to exercise market power over time through strategic reporting of these attributes.

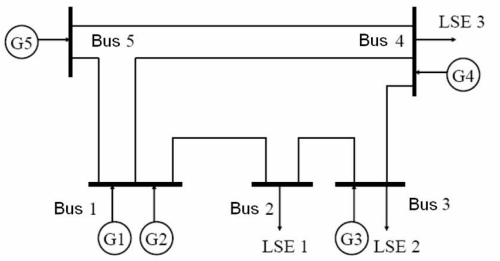

Fig. 6. Transmission grid for the benchmark dynamic 5-bus test case.

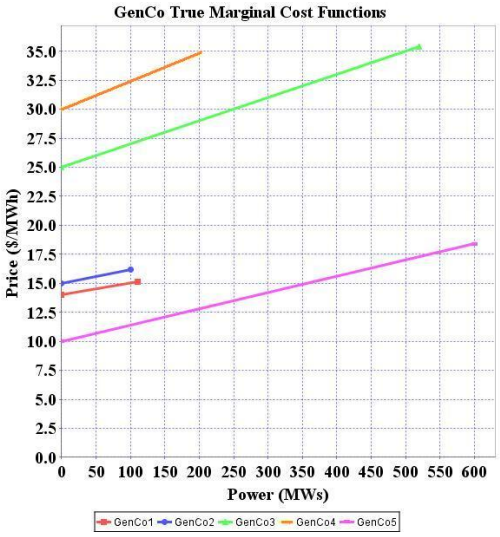

Fig. 7. GenCo true marginal cost functions and true capacity attributes for the benchmark dynamic 5-bus test case.

Each experiment reported in this study extends the benchmark dynamic 5-bus test case by systematically varying one or more treatment factors. Three types of treatment factors are considered: GenCo learning capabilities; LSE demand-bid price sensitivity; and an ISO-imposed supply-offer price cap.

With regard to GenCo learning, in each experiment one of the following two treatments is imposed. Either (i) the GenCos are non-learners, or (ii) each GenCo i is a learner that makes daily use of the VRE-RL algorithm to adjust the ordinate and slope parameters $\{a_i^R, b_i^R\}$ of its reported marginal cost

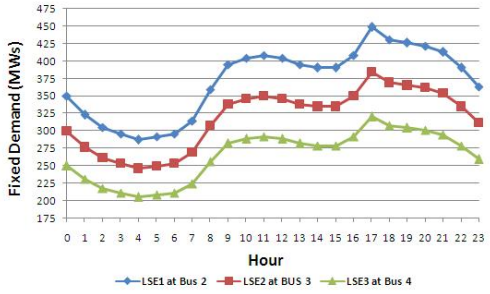

Fig. 8. Daily LSE fixed demand (load) profiles for the benchmark dynamic 5-bus test case.

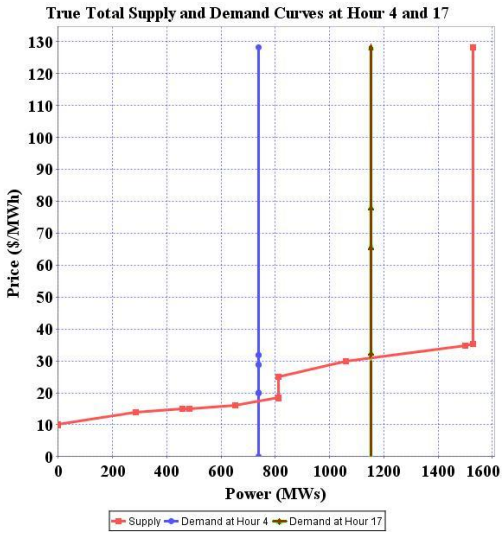

Fig. 9. True total supply and demand curves for hours 4:00 and 17:00 for the benchmark dynamic 5-bus test case. Demand for this benchmark case is 100% fixed (R=0.0).

function (3) in pursuit of increased net earnings.[10]

For the learning treatments, the action domain AD_i for each GenCo i is constructed as in [7] to include 100 candidate supply offer choices, and the VRE-RL recency and experimentation parameters r_i and e_i for each GenCo i are fixed at 0.04 and 0.96, respectively, in keeping with the VRE-RL parameter sensitivity results determined in [22]. A range of settings is then systematically tested for each GenCo i's VRE-RL initial propensity and temperature parameters $q(1)_i$ and T_i; see Table IV for a precise listing of tested values.

When GenCos have learning capabilities, random effects are present in their supply offer selections. To control for these random effects, we generated thirty seed values via the standard Java class "random;" see Table IV for a listing of these seed values. For each learning treatment we then used these thirty seed values to implement thirty distinct runs, each 1000 simulated days in length.[11]

The second treatment factor we consider is the ratio R of maximum potential price-sensitive demand to maximum potential total demand. More precisely, for each LSE j and each hour H, let

$$R_j(H) \;=\; \frac{SLMax_j(H)}{MPTD_j(H)} \quad . \tag{14}$$

In (14) the expression $SLMax_j(H)$ denotes LSE j's *maximum potential price-sensitive demand* in hour H as measured by the upper bound of its purchase capacity interval (2), and

$$MPTD_j(H) \;=\; [p^F_{Lj}(H) + SLMax_j(H)] \;\; (MW) \tag{15}$$

[10]Recall that a detailed description of the VRE-RL algorithm is provided in Appendix C. In this study the GenCos are only allowed to exercise economic withholding of capacity through strategic marginal cost reporting. In a separate study [21] we explore the consequences of permitting GenCos to engage in physical withholding of capacity through strategic reporting of their operating capacity limits.

[11]As a stability check, we sampled each run at multiple intermediate days. In all cases, the outcome variables of interest showed essentially no change from day 500 to day 1000.

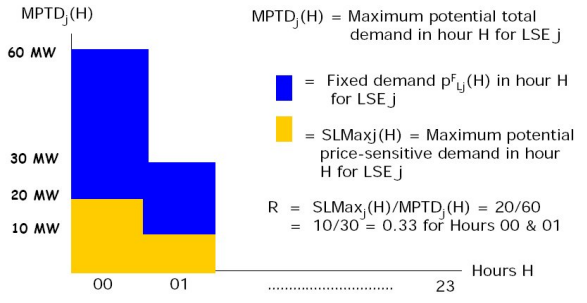

Fig. 10. Illustration of the R ratio construction for measuring relative demand-bid price sensitivity. The depicted special case is for R=0.33

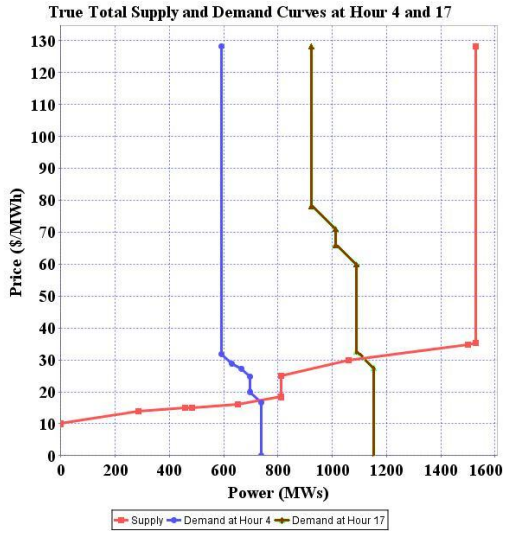

Fig. 11. True total supply and demand curves for hours 4:00 and 17:00 for the benchmark dynamic 5-bus test case extended to include 20% potential price-sensitive demand (R=0.2).

denotes LSE j's *maximum potential total demand* in hour H as the sum of its fixed demand and its maximum potential price-sensitive demand in hour H. The construction of the R ratio is illustrated in Figure 10 for the special case R=0.33.

For our price-sensitive demand experiments we start by setting all of the R values (14) for each LSE j and each hour H equal to R=0.0 (the 100% fixed-demand case). We then systematically increase R by tenths, ending with the value R=1.0 (the 100% price-sensitive demand case). A positive R value indicates that the LSEs are able to exercise at least some degree of price resistance. Compare, for example, the true total demand curves in Fig. 9 with 100% fixed demand (R=0.0) to the true total demand curves in Fig. 11 with 20% potential price-sensitive demand (R=0.2).

The maximum potential price-sensitive hourly demands $\text{SLMax}_j(\text{H})$ 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 $\text{MPTD}_j(\text{H})$ in (15) is held constant across experiments by appropriate reductions in the fixed demand $\text{p}^F_{Lj}(\text{H})$ as $\text{SLMax}_j(\text{H})$ is increased. Specifically, $\text{MPTJ}_j(\text{H})$ is set equal across all experiments to $\text{BP}^F_{Lj}(\text{H})$, the hour-H fixed-demand level $\text{BP}^F(\text{H})$ for LSE j depicted in Table III for the benchark dynamic 5-bus test case. Consequently, for each tested R value,

$$p^F_{Lj}(H) = [1-R] * BP^F_{Lj}(H)\ ; \quad (16)$$

$$SLMax_j(H) = R * BP^F_{Lj}(H). \quad (17)$$

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 holding fixed the ordinate and slope values $\{(\text{c}_j(\text{H}),\text{d}_j(\text{H})): \text{H}=00,...,23\}$ for each LSE j. A listing of the specific numerical values used can be found in Table V.

In particular, for conceptual consistency with the benchmark dynamic 5-bus test case with no price-sensitive demand (R=0.0), LSE j's ordinate value $\text{c}_j(\text{H})$ is set equal to the hour-H LMP solution $\text{LMP}_{k(j)}(\text{H})$ for this benchmark case, where $k(j)$ denotes the particular bus k at which LSE j is located. This guarantees that no price-sensitive demand would be cleared in the benchmark case even if the LSEs were permitted to report price-sensitive demand bids as well as fixed demand bids. Also, the ratio $\text{c}_j(\text{H})/2\text{d}_j(\text{H})$ for each LSE j is set to ensure that it is greater or equal to the $\text{SLMax}_j(\text{H})$ value determined by (17), as required by the admissibility restrictions imposed on LSE j's demand function $\text{D}_{Hj}(\text{p})$ in Table I.

The third treatment factor we consider is PCap ($/MWh), an ISO-imposed supply-offer price cap. In experiments in which PCap is imposed, GenCos are not permitted to report marginal costs (sale reservation values) that rise above PCap. Consequently, each GenCo i selects its daily supply offer so that its maximum reported sale reservation value, $\text{MC}^R_i(\text{Cap}^{RU}_i)$, does not exceed PCap.

As will be seen in the following Section IV – in particular Table XIII – the mean outcome for average hourly LMP with GenCo learning and with 100% fixed demand (R=0.0) is approximately 140 ($/MWh). We therefore tested six PCap settings centered around this "normal" value, as follows: (a) no PCap; (b) a high value 160; (c) a normal value 140; (d) a moderately low value 120; (e) a low value 100; and (f) a very low value 80.

IV. Report of Key Findings for Dynamic Market Performance

This section uses the experimental design outlined in Section III to test the effects on dynamic market performance of changes in GenCo learning capabilities, demand-bid price-sensitivities, and supply-offer price caps. Dynamic market performance is characterized by the following seven measures:

- total GenCo daily net earnings (Total Gen DNE)
- average hourly cleared price-sensitive and fixed demand (Avg Total Demand) for LSEs
- average hourly true total avoidable costs (Avg TrueTVCost) for GenCos
- average hourly reported total avoidable costs (Avg RepTVCost) for GenCos
- average hourly Lerner index values (Avg LI) for GenCos
- LMP spiking across 24 hours of a designated day
- LMP volatility range across 24 hours of a designated day

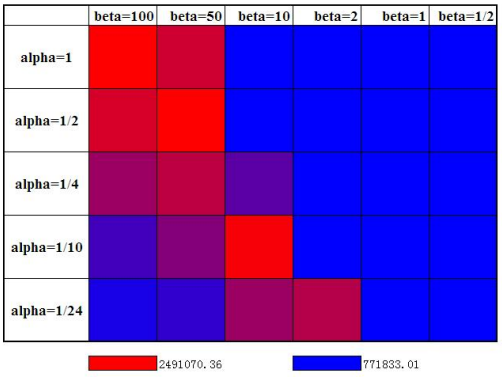

Fig. 12. A heat-map depiction of mean outcomes for total GenCo daily net earnings (Total Gen DNE) under alternative settings for the VRE reinforcement learning parameters (α, β) for the benchmark dynamic 5-bus test case extended to include GenCo learning. Red (lighter) shades indicate higher Total Gen DNE.

For no-learning treatments, we calculate each of these measures for a typical day. For learning treatments, we calculate the mean of each of these measures across thirty runs for the final (1000th) simulated day. The precise definitions and calculations of these measures (with accompanying standard deviations) are provided in Appendix D.

A. Learning Calibration

As a prelude to conducting experiments with GenCo learning, we first used intensive parameter sweeps to determine suitable settings for two potentially critical VRE-RL learning parameters for each GenCo i. In particular, as indicated in Table IV, we systematically tested a range of values for α_i and β_i, defined as follows:

- GenCo i's net earnings aspirations at the beginning of the initial day 1, as captured by the ratio α_i of its initial propensity level $q_i(1)$ to its maximum possible daily net earnings $MaxDNE_i$ defined in (12);
- the ratio β_i = $q_i(1)/T_i$ of GenCo i's initial propensity level $q_i(1)$ to its temperature parameter T_i.

Figure 12 depicts experimental findings for mean Total Gen DNE outcomes under alternative settings for[12]

$$\alpha = \frac{q_i(1)}{MaxDNE_i}, \quad \beta = \frac{q_i(1)}{T_i}, \; i = 1, \ldots, I \qquad (18)$$

assuming common α and β values across GenCos, 100% fixed demand, and no supply-offer price cap. An interesting pattern is immediately evident. The (α,β) combinations associated with the highest mean Total Gen DNE outcomes lie 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 mean Total Gen DNE on α and β?

The settings for α and β have distinct but correlated effects on the degree to which each GenCo experiments with different actions, i.e., different ordinate and slope values a^R and b^R for its reported marginal cost function (3). All else equal, high α values reflecting optimistically high initial net earnings expectations tend to induce experimentation with many different actions due to "disappointment" with the net earnings outcomes that result from each choice. Conversely, low α values reflecting pessimistically low initial net earnings expectations tend to induce premature fixation on an early chosen action due to the "surprisingly high" net earnings that result from this choice.

High β values reflecting high cooling levels (low temperature parameter settings) amplify the tendency to premature fixation in the case of low α values by amplifying differences in propensity levels across action choices. Moderately low β values can prevent premature fixation by dampening the effects of propensity differences on action choice probabilities. However, extremely low β values result in action choice probability distributions that are essentially uniform across each GenCo's action domain, negating all GenCo efforts to learn which actions result in the highest daily net earnings. This deleterious effect is seen in the uniformly low mean Total Gen DNE outcomes achieved in Fig. 12 for the lowest tested β levels 1 and 1/2.

As seen in Fig. 12, the GenCos attain their highest mean Total Gen DNE outcomes for the *sweet-spot* VRE learning parameter settings (α, β) = (1,100). These sweet-spot settings are used in all of the learning treatments reported in the remainder of this study.

B. Pure Learning Experiments

Fig. 13 depicts hourly bus LMP levels, GenCo dispatch levels, and branch power flows during a typical day for the benchmark dynamic 5-bus test case. Recall this benchmark case involves non-learning GenCos, 100% fixed demand, and no supply-offer price cap. Fig. 14 depicts hourly LMP levels, GenCo dispatch levels, and branch power flows for day 1000 of a typical run (ID=03) for this benchmark case after extension to include learning GenCos.

Comparing Figures 13 and 14, the most significant pure learning effect is clearly the substantial increase in LMP outcomes for each bus in each hour, ranging from an approximate 2-fold increase for buses 1 and 5 to an approximate 6-fold increase for buses 2 and 3. In addition, learning also affects the GenCo dispatch levels. For example, the dispatch level for the peaker-unit GenCo 4 located at bus 4 is higher in each hour whereas the dispatch level for the small GenCo 1 located at bus 1 is markedly lower in every hour except the peak-demand hour 17. In contrast, branch power flows appear to be relatively unaffected.

Tables VI through XI provide more detailed numerical data regarding the effects of learning on hourly bus LMPs, GenCo dispatch levels and branch power flows. These numerical data help to explain the learning effects seen in Figures 13 and 14.

[12]As detailed in Section II, the estimate $MaxDNE_i$ for each GenCo i's maximum possible daily net earnings is exogenously derived from its action domain AD_i. Consequently, each specification for α and β in (18) determines a distinct initial propensity level $q_i(1)$ and temperature level T_i for each GenCo i. In our earlier study [7], identical initial propensity and temperature levels were set for all learning GenCos: namely, q(1)=6000 and T=1000. This was unsatisfactory since the "prior anticipated net earnings" q(1) were then set commonly across GenCos with different costs and locations without regard for their distinct earnings opportunities.

Fig. 13. Hourly bus LMP levels, GenCo dispatch levels, and branch power flows during a typical day for the benchmark dynamic 5-bus test case (no GenCo learning).

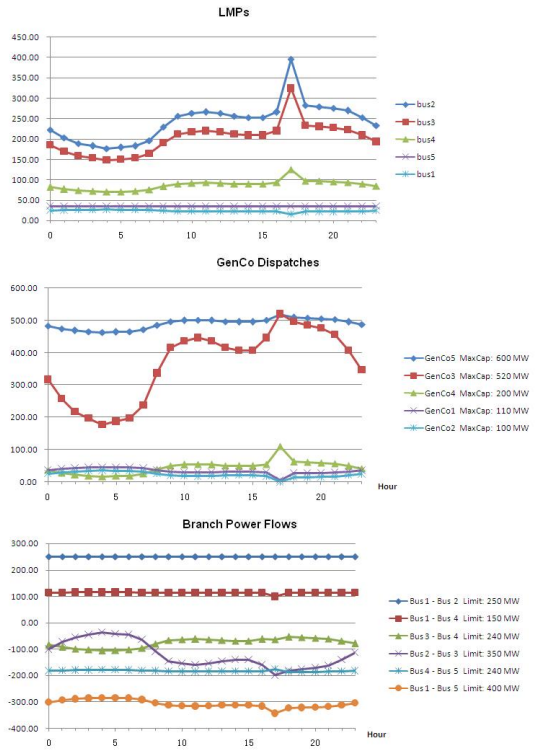

Fig. 14. Hourly bus LMP levels, GenCo dispatch levels, and branch power flows for a typical run (ID=03) of the benchmark dynamic 5-bus test case extended to include GenCo learning.

Consider, first, the benchmark no-learning case. To understand the pattern of LMPs reported in Table VI for this benchmark case, it is important to understand congestion effects. Note from Table X that the branch 1-2 connecting bus 1 and bus 2 is congested in every hour; all other branches are uncongested in every hour.

The congestion on branch 1-2 creates a potential *load pocket* for GenCo 3 in the following sense. As seen from the depiction of the 5-bus transmission grid in Figure 6, the fixed load from LSEs 1, 2, and 3 is located at buses 2, 3, and 4, and GenCo 3 at bus 3 is centrally located relative to this load. The congestion on branch 1-2 results in the semi-islanding of this load from the less-expensive power of GenCos 1, 2, and 5 located at buses 1 and 5. Consequently, the ISO must dispatch GenCo 3 to meet the bulk of this load, particularly during the peak-demand hour 17, no matter what the expense.

To fully understand the pattern of hourly bus LMPs reported in Table VI for the benchmark no-learning case, however, it is also essential to consider limits on generation operating capacity. A GenCo i is said to be *marginal* if it is operating at a point where it is not constrained either by its lower or upper operating capacity limits Cap_i^L and Cap_i^U in (6). As is well known, the LMP at each bus with a marginal GenCo is given by the marginal cost of this GenCo, whereas the LMP at each bus without a marginal GenCo is given by a weighted linear combination of the marginal costs of the marginal GenCos; see, e.g., [4].

From Table VIII, it is seen that GenCo 1 is only marginal during the peak-demand hour 17 and the non-peak hour 18, whereas GenCo 2 and GenCo 3 are marginal in every hour except hour 17. Also, GenCo 4 is only marginal during the peak-demand hour 17, and GenCo 5 is marginal in every hour.

It follows that the LMP at bus 1 (with GenCos 1 and 2) is determined at the peak-demand hour 17 by the marginal cost of the marginal GenCo 1. For all non-peak hours apart from hour 18 the LMP at bus 1 is determined by the marginal cost of the marginal GenCo 2. For the non-peak hour 18 the LMP at bus 1 is determined by the equalized marginal costs of the marginal GenCos 1 and 2. Note from Figure 7 that GenCo 1 is a relatively cheap generation source, and GenCo 2 is only slightly more expensive than GenCo 1. Consequently, as seen in Table VI, the LMP at bus 1 is relatively low in all hours, particularly so in hours 17 and 18 when GenCo 1 is marginal. Similar arguments explain the relatively low LMP level for bus 5 in all hours.

In contrast, apart from hours 17 and 18 the LMP at bus 2 (with no generation) is determined as a weighted linear combination of the marginal costs of the marginal GenCos 2, 3, and 5. For the peak-demand hour 17 the LMP at bus 2 is determined as a weighted linear combination of the marginal costs of the marginal GenCos 1, 4, and 5. For the non-peak hour 18 the LMP at bus 2 is determined as a weighted linear combination of the marginal costs of the marginal GenCos 1, 2, and 5. As seen in Table VI, the need to dispatch the expensive peaker unit, GenCo 4, during the peak-demand hour 17 due to the congestion on branch 1-2 results in an approximate doubling of the LMP at bus 2 during this hour relative to other hours. Similar arguments explain the relatively

large bump in LMP at bus 3 and bus 4 during hour 17.

Now consider, instead, the hourly bus LMP outcomes reported in Table VII for the benchmark case extended to include learning GenCos. Comparing these outcomes to the outcomes reported in Table VI for the no-learning case, it is immediately seen that the LMPs attained with learning GenCos are substantially higher in all hours. What explains this?

As seen in Table XI, the branch 1-2 connecting bus 1 and bus 2 is congested at all hours with learning GenCos, just as it was for non-learning GenCos. On the other hand, comparing the dispatch outcomes reported in Table IX for learning GenCos with the dispatch outcomes reported in Table VIII for non-learning GenCos, it is seen that learning changes these dispatch levels and hence the marginal status of the GenCos. In particular, every GenCo is now marginal in every hour, apart from GenCo 2 and GenCo 3 in the peak-demand hour 17.

The explanation for these dispatch effects is that the learning GenCos, in particular the two largest GenCos 3 and 5, quickly learn to report higher-than-true marginal costs to the ISO. This economic withholding means that the dispatch merit order calculated by the ISO from reported marginal cost functions no longer coincides with the true merit order based on true marginal cost functions, which in turn affects the ISO's dispatch schedule.

Economically, however, the most serious effect of this economic witholding is not the changed dispatch levels per se but rather the resulting increase in LMPs. The price rise relative to the benchmark no-learning case is particularly dramatic for the load-pocket buses 2 through 4 during the peak-demand hour 17.

The opportunity for learning GenCos to profitably undertake substantial economic withholding arises from the fact that LSE demand in the benchmark case is 100% fixed (no price sensitivity). The ISO is forced to meet fixed demand in every hour, no matter how expensive the required generation might be. Consequently, the GenCos rapidly come to understand, through trial-and-error reinforcement learning, that their most profitable strategy is to implicitly collude on high reported marginal cost functions. Since all GenCos end up exercising economic withholding, the overall effect on the dispatch schedule and resulting branch power flows is relatively modest; but the increase in hourly bus LMPs is substantial.

These findings suggest the importance of encouraging a greater sensitivity of LSE demand to price. The following subsection explores what happens when LSE demands are systematically varied from 100% fixed to 100% price sensitive, both with and without GenCo learning.

C. Price-Sensitivity Experiments With and Without Learning

Table XII presents dynamic market performance findings for the benchmark dynamic 5-bus test case extended to include alternative settings for R (relative demand-bid price sensitivity). Table XIII presents parallel mean-outcome findings for a modified version of this experiment in which GenCos have learning capabilities and report strategically chosen marginal cost functions to the ISO.

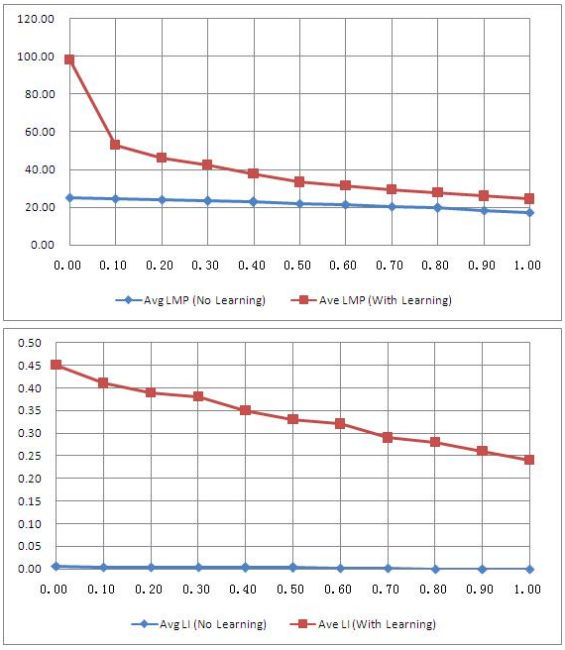

Fig. 15. Mean outcomes for average hourly LMPs and LI levels on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and demand varying from R=0.0 (100% fixed) to R=1.0 (100% price sensitive).

As seen in Table XII, in the absence of GenCo learning an incremental increase in R starting from the benchmark case R=0.0 (100% fixed demand) has the usual intuitively-expected effects. Avg LMP, Avg Total Demand, Avg TrueTVCost, and Avg LI all monotonically decline with increases in R. Indeed, except for the presence of binding operating-capacity constraints for some of the GenCos for the cases in which Avg Total Demand is relatively high, all of the Avg LI outcomes in the absence of GenCo learning would be zero.[13]

Comparing the no-learning results presented in Table XII to the results with GenCo learning presented in Table XIII, it is seen that GenCo learning has strong effects. Mean outcomes for Avg LMP, Avg Total Demand, Avg TrueTVCost, and Avg LI all monotonically decline with increases in R, as before. However, as highlighted in Fig. 15, mean Avg LMP and mean Avg LI are substantially higher for each level of R even though mean Avg Total Demand is lower for each positive level of R.

In addition, mean Avg TrueTVCost under learning is higher than its corresponding no-learning level at R=0.0 and R=0.2 due to out-of-merit order dispatch. However, as R continues to increase, mean Avg TrueTVCost under learning falls below its corresponding no-learning level due to the relatively stronger contraction in mean Avg Total Demand. Moreover, under learning, mean Avg RepTVCost is substantially higher than

[13]In no-learning treatments the GenCos report their true cost and capacity conditions to the ISO each day. Consequently, the GenCos do not deliberately exercise market power, i.e., they do not engage in either economic or physical withholding of capacity. Nevertheless, for reasons explained in [23], a binding operating-capacity constraint on a GenCo G located at a bus k typically causes the LMP at bus k to separate from the marginal cost of G. In standard economic terminology, the cleared units of capacity-constrained GenCos are *strictly inframarginal*, meaning they are not the units at the intersection of demand and supply that determine the market clearing price. This separation results in a non-zero value for this GenCo's LI value.

mean Avg TrueTVCost at each level of R.

The explanation for these effects is that the profit-seeking GenCos quickly learn to implicitly collude on higher-than-true reported marginal costs. This implicit collusion occurs even when demand bids are fully price sensitive (R=1.0) and the GenCos are forced to compete for limited demand.

Real-world day-ahead markets are meant to operate as *double auctions*, i.e., as two-sided auctions with actively managed demand bids as well as actively managed supply offers. As elaborated in ([23],[24]), theoretical, empirical, and human-subject experimental studies all provide strong support for the general efficiency of the double-auction market form. A cautionary implication of the findings in this subsection is that the preponderance of passive fixed demand in real-world day-ahead markets (due largely to a lack of retail market restructuring) prevents the proper operation of these markets as double auctions. Given essentially vertical demand curves unresponsive to price, the only way that ISOs can hope to control the exercise of seller market power is through the imposition of strong mitigation rules that constrain seller supply-offer behaviors.

D. Price Sensitivity and ISO Net Surplus Extraction

As detailed in [25], in a standard ISO-managed day-ahead market the ISO settles the market for each day D+1 at the end of day D. Each power buyer is obliged at the end of day D to purchase from the ISO its cleared power demand for hour H of day D+1 at a price given by the particular LMP determined for its particular bus location. Similarly, each power seller is obliged at the end of day D to sell to the ISO its cleared power supply for hour H of day D+1 at a price given by the particular LMP determined for its particular bus location.

For the 5-bus test case at hand, the power buyers are LSEs and the power sellers are GenCos. Also, hourly LMPs for the day-ahead market are determined by the ISO by means of bid/offer-based DC OPF with losses assumed to be zero. The difference between the total LSE payments collected by the ISO on day D and the total GenCo revenues paid out by the ISO on day D is referred to below as the *ISO daily net surplus (on day D)*.[14] In the absence of transmission grid congestion, a uniform LMP is determined across the grid and ISO daily net surplus is necessarily zero. On the other hand, congestion causes the LMPs to separate across the grid, and in this case it can be shown that the ISO daily net surplus must be non-negative and typically will be positive; see [26, Prop. 2.1].

Table XIV reports total GenCo daily net earnings, total GenCo daily revenues, total LSE daily payments, and ISO daily net surplus during a typical day for the benchmark dynamic 5-bus test case extended to permit six different treatments for R (demand-bid price sensitivity). For each R treatment, congestion occurs on branch 1-2. Consequently, as expected, ISO daily net surplus is strictly positive for each R treatment.

However, Table XIV also reveals two new findings. First, when R=0.0 (100% fixed demand) the ISO daily net surplus is substantial, comprising approximately 28% of total LSE daily payments. This implies that total LSE daily payments are substantially higher than total GenCo daily revenues. As indicated in Table VI, this differential is due to the relatively high LMP for bus 2, which has the largest load (LSE 1) and no generation, and to the relatively low LMPs for bus 1 and bus 5, which have generation but no load.

Second, total GenCo daily net earnings, total GenCo daily revenues, total LSE daily payments, and ISO daily net surplus all undergo marked monotonic declines as R increases to R=1.0 (100% price-sensitive demand). The explanation for these monotonic declines can be deduced from the findings in the previous subsection IV-C as follows.

Given low R values, the LSEs have very little price resistance; they are willing to pay any price to satisfy their fixed demands, and their fixed demands constitute the bulk of their total demands. Around the peak-demand hour 17, due in part to congestion on branch 1-2, the ISO must dispatch the most expensive GenCos 3, 4, and 5 to meet the large LSE fixed demand, i.e., these GenCos are pivotal suppliers for hour 17. This results in relatively high LMPs. As R increases, however, the LSEs are increasingly able to resist high prices through demand withholding. As indicated in Table XII, this results in lower LMPs, lower total demand, and lower avoidable costs of production. Total GenCo daily revenues and total LSE daily payments are thus lower, and total GenCo daily net earnings are also lower because the decline in total GenCo daily avoidable costs is more than offset by the decline in total GenCo daily revenues. Similarly, ISO daily net surplus is lower because the decline in total GenCo daily revenues is more than offset by the decline in total LSE daily payments.

Table XV reports mean outcomes for total GenCo daily net earnings, total GenCo daily revenues, total LSE daily payments, and ISO daily net surplus on day 1000 for a parallel set of price-sensitivity experiments with learning GenCos. As in the previous no-learning case, six different treatments for R are tested. Congestion again occurs on branch 1-2 in all hours under all R treatments, resulting in positive mean outcomes for ISO daily net surplus. Moreover, the mean outcomes for total GenCo daily net earnings, total GenCo daily revenues, total LSE daily payments, and ISO daily net surplus all monotonically decline with increases in R as a result of the corresponding decreases in LMPs, total demand, and true avoidable costs of production; see Table XIII.

Comparing the outcomes in Tables XIV and XV, however, it is seen that learning does have economically significant effects. For each tested R level the mean total GenCo daily net earnings under learning are substantially higher than for the no-learning case, ranging from a 26-fold increase for R=0.0 to a 2.5-fold increase for R=1.0. In particular, even when demand is 100% price sensitive (R=1.0) the GenCos still learn to exercise some degree of market power through economic withholding of capacity.

Also, for R=0.0 the mean outcomes for total GenCo daily revenues, total LSE daily payments, and ISO daily net surplus under learning are all substantially higher than their corresponding values under no learning. Specifically, mean total GenCo daily revenues are over five times higher, mean total

[14]Other terms used in the literature include "congestion rent" (when the surplus is due entirely to congestion) and "merchandising surplus."

LSE payments are over six times higher, and mean ISO daily net surplus is over ten times higher. Since total demand for R=0.0 is the same under learning and no learning, the ten-fold increase in mean ISO daily net surplus under learning implies that the mean LMP paid by the LSEs is substantially higher than the mean LMP received by the GenCos. As indicated in Table VII, this is due to the approximately six-fold increase under learning in the mean LMP for bus 2, which has the largest load (LSE 1) and no generation, and to the much smaller increases under learning in the mean LMPs for bus 1 and bus 5, which have generation but no load.

As R increases, however, the mean outcomes for total GenCo daily revenues, total LSE daily payments, and ISO daily net surplus under learning eventually drop below their corresponding values under no learning. For mean total GenCo daily revenues the switch point is at R=1.0, whereas for mean total LSE daily payments and mean ISO daily net surplus the switch point is at R=0.6.

The explanation for these switch points can be deduced from the LMP and total demand findings presented in Tables XII and XIII for the no-learning and learning cases. When GenCos are learners, low R values (implying large fixed demands) provide pivotal suppliers with a substantial opportunity to engage in profitable economic withholding. This dramatically increases LMPs relative to the no-learning case, particularly at the load-only bus 2. Since total demand for the learning case is only modestly lower than for the no-learning case for low R values, the end result is substantially higher total GenCo daily revenues, total LSE daily payments, and ISO daily net surplus.

On the other hand, as R increases and the LSEs acquire an increasing ability to resist high prices through demand withdrawal, the learning GenCos are increasingly forced to compete with each other for dispatch by lowering their reported marginal costs. This competitive process results in lower LMPs. However, comparing the outcomes in Tables XII and XIII, it is seen that the LMPs resulting under learning remain higher than under no learning for all R values, which in turn induces the LSEs to engage in greater demand withholding under learning.

The result is that mean total GenCo daily revenues, mean total LSE daily payments, and mean ISO daily net surplus under learning all fall below their corresponding no-learning values as R approaches 1.0 due to the relatively strong contraction in total demand under learning. As can be verified from the total GenCo daily net earnings data provided in Table XV, the most expensive GenCo 4 is at the greatest disadvantage in this competitive process while the least expensive GenCo 5 is most advantaged.

A general conclusion drawn from the findings presented in Tables XIV and XV is that ISO net surplus is substantially enhanced both by grid congestion (LMP separation) and by a lack of demand-bid price sensitivity (low R values) leading to GenCo economic withholding of capacity. Overall market efficiency is compromised when congestion and economic withholding result in out-of-merit-order dispatch of generation. System reliability can also be compromised when capacity is effectively withheld from market operations.

The ISOs for U.S. restructured wholesale power markets are typically organized as independent not-for-profit entities with a fiduciary responsibility for ensuring market efficiency and reliable system operations. Maximization of ISO net surplus is certainly not the intended objective of ISOs. Nevertheless, ISO net surplus represents a revenue stream for ISOs whose use is discretionary to the ISO.

In some markets (e.g., MISO) the ISO net surplus is used in part to encourage new transmission investment through the subsidization of financial transmission rights for those who invest in new transmission capacity, which should in principle lead to a reduction of ISO net surplus to the extent that congestion is alleviated. However, transmission investment needs can arise for reasons other than congestion (e.g., the need to reach distributed energy sources), and congestion might better be alleviated by more local generation rather than by more transmission capacity. Consequently, it is unclear whether the benefits flowing from the current uses of ISO net surplus outweigh the potential costs; indeed, these uses could even amplify the problem of misallocated resources. In any case the basic misalignment between intended ISO objectives and ISO revenue incentives remains.

A cautionary implication of the findings in this subsection, then, is that transparent reporting and oversight of ISO operations ("monitoring of the monitors") is desirable because LMP pricing as currently practiced does not achieve proper incentive alignments for ISOs.

E. Price Cap Experiments With and Without Learning

Table XVI reports mean outcomes for average hourly LMP during day 1000 for the benchmark dynamic 5-bus test extended to include GenCo learning and a supply-offer price cap (PCap).[15] None of the tested PCap settings is binding on supply offers in the absence of GenCo learning. To see this, note in Figure 7 that the highest true marginal cost for any GenCo over its true capacity operating interval is only about 35.40 ($/MWh), which is much lower than the lowest tested PCap value 80 ($/MWh). Consequently, the average hourly LMP outcome 25.18 ($/MWh) with non-learning GenCos provides a common benchmark of comparison for all of the PCap treatments with learning GenCos.

Although the tested PCap settings are not binding for non-learning GenCos, they can be binding on the marginal cost functions reported by learning GenCos. In this study it is assumed that GenCos with learning capabilities whose reported marginal cost functions are constrained by PCap are not willing to supply power at reported marginal costs that exceed PCap. Rather, they reduce their reported maximum capacities until their reported marginal costs at their reported maximum capacities are no greater than PCap.

[15]For the subsequent interpretation of these findings, it is important to recall from Section III that PCap is a price cap on GenCo-reported supply offers (marginal cost functions) and not on LMPs per se. LMPs are system marginal costs subject to network effects, not GenCo marginal costs. As discussed more carefully in [23], in the presence of grid congestion the LMPs at buses without marginal GenCos can strictly exceed the marginal cost of each GenCo. Consequently, PCap is not necessarily an upper bound on LMPs.

Consequently, in PCap experiments with learning GenCos, capacity shrinkage can result in a total offered supply that is below total fixed demand. Careful attention must therefore be paid to the possible occurrence of *inadequacy events (IE)*, i.e., hours during which GenCo offered supply is less than LSE fixed demand.

In Table XVI two different methods are used to account for IE effects. The first method ("with IE") sets a *reserve price* of 1000 ($/MWh) during any hour in which an IE occurs, and this reserve price is used as the LMP at each bus for this hour. The second method ("without IE") simply ignores hours during which an IE occurs.

The outcomes reported in Table XVI show that no IE occurs in the absence of a supply-offer price cap; offered supply is adequate to meet fixed demand in each hour. However, as PCap is successively lowered from 160 to 80, the frequency of IE increases from 4% to 31.1%. Ignoring hours in which IE occurs, it appears that the imposition of a successively lower PCap results in a successively lower mean Avg LMP value, although this value is still substantially higher than in the no-learning case. However, when IE hours are taken into account by imposition of the reserve price, results are dramatically different; the successive lowering of PCap results in a subtantial increase in mean Avg LMP, a reflection of the substantial increase in IE frequency.

One aspect of the mean Avg LMP outcomes reported in Table XVI with GenCo learning and w/o IE might appear puzzling. Note that mean Avg LMP with no price cap is 140.30 ($/MWh) whereas mean Avg LMP for PCap=160 ($/MWh) is only 89.65 ($/MWh). This finding indicates that the PCap level 160 is binding on the GenCos' reported marginal costs even though this PCap level is substantially higher than the resulting value 89.65 for mean Avg LMP. A similar comment holds for the remaining four tested PCap levels.

The explanation for this finding is that the distribution of LMPs across the 24 hours of a day can exhibit substantial fluctuations that are obscured when only average hourly LMP outcomes are considered. In particular, the maximum LMP value attained during the peak-demand hour on any given day can be substantially higher than the average hourly LMP attained for this day. Thus, the imposition of a price cap can be a binding constraint on GenCo-reported marginal costs during peak-demand hours even if not in other hours. Since GenCos are only permitted to report one supply offer per day, a binding constraint on reported marginal costs during peak-demand hours translates into a binding supply-offer constraint for every hour.

Figure 16 provides a more disaggregated 24-hour depiction of the mean Avg LMP results reported in Table XVI for GenCo learning with IE reserve prices taken into account. It is now seen more clearly that IE largely occurs around the peak-demand hour 17, and that IE tends to occur in these hours with higher frequency for lower PCap settings.

Figure 17 presents still another way to visualize the mean Avg LMP outcomes reported in Table XVI for GenCo learning with IE reserve prices taken into account. This figure shows that the imposition of a successively lower PCap tends to induce a dramatic increase in LMP spiking and volatility range

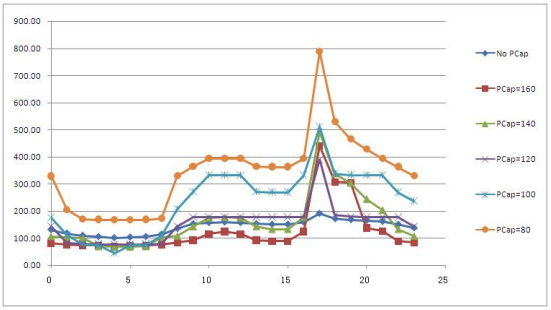

Fig. 16. Mean outcomes for average hourly LMP levels on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and a supply-offer price cap varying from infinitely high (none) to low (80). Inadequacy Event (IE) LMP reserve prices are included in this figure.

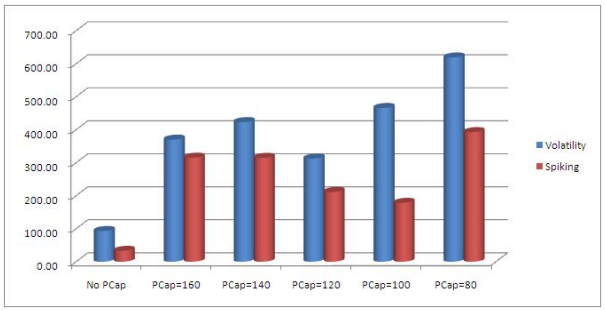

Fig. 17. Mean outcomes for LMP spiking and LMP volatility range on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and a supply-offer price cap varying from infinitely high (none) to low (80). Inadequacy Event (IE) LMP reserve prices are included in this figure.

relative to the no-PCap treatment. As explained more carefully in Appendix D, "spiking" refers to the absolute difference between successive hourly LMPs across all 24 hours of the final (1000th) simulated day, whereas "volatility range" refers to the difference between maximum and minimum LMP across all 24 hours of this final simulated day.

The cautionary bottom line here is that supply-offer price caps can have unintended consequences that outweigh intended benefits. Improperly imposed caps can lead to increased LMP spiking and volatility as well as increased system security issues through inducement of IE, particularly around peak-demand hours, even if LMP values are indeed lowered during other hours.

V. Report of Key Findings for Spatial Cross-Correlations

A. Correlation Experiment Preliminaries

Table VII depicts hourly bus LMP outcomes on day 1000 for a typical run of the benchmark dynamic 5-bus test case extended to include GenCo learning. Table XVII provides mean outcomes for hourly bus LMP outcomes on day 1000 across all runs under the same experimental conditions.

This subsection examines the extent to which these hourly bus LMPs are cross-correlated with GenCo reported marginal costs and with each other. Of particular interest is the extent to which cross-correlations are induced in hourly bus LMPs

either by the marginal status of strategically located and sized GenCos or by network effects.

Three types of experimental findings are reported below: (a) pairwise cross-correlations between reported GenCo marginal costs evaluated at dispatch operating points; (b) pairwise cross-correlations between GenCo reported marginal costs and bus LMPs evaluated at dispatch operating points; and (c) pairwise cross-correlations between bus LMPs evaluated at dispatch operating points. In each case the cross-correlations are calculated at the following four representative hours from the LSE load profiles depicted in Figure 8:

- the off-peak hour 4:00
- the shoulder hour 11:00
- the peak-demand hour 17:00
- the shoulder hour 20:00

Moreover, for each of these four hours the three types of cross-correlations are calculated for three different demand scenarios as characterized by three different settings for R. In total, then, thirty-six distinct cross-correlation treatments ($3 \times 4 \times 3$) are reported below.

Illustrative findings from these treatments are depicted using correlation diagrams as well as tables. Each correlation diagram uses shape, shape direction, and color to convey information about the sign and strength of the resulting pairwise cross-correlations.

The shapes and shape directions in the correlation diagrams are rough indicators of the patterns observed in the underlying scatter plots for the two random variables whose cross-correlation is under examination. Color is used to reinforce shape and shape direction information.

More precisely, if a scatter plot for two random variables X and Y roughly lies along a straight line, this suggests that X and Y are perfectly correlated. If the line is positively sloped, the indication is perfect positive correlation (1.0); if the line is negatively sloped, the indication is perfect negative correlation (-1.0). The correlation diagrams indicate these possible patterns by means of straight lines that are either forward or backward slanted to indicate positive or negative correlation respectively. Conversely, if the scatter plot for X and Y instead consists of a roughly rectangular cloud of points, this indicates that X and Y are independent of each other, implying zero correlation. The correlation diagrams indicate this pattern by means of full circles. Intermediate to this are scatter plots for X and Y that are roughly elliptical in shape, indicating moderate but not perfect correlation between X and Y. The correlation diagrams indicate this pattern by means of oval shapes that point to the right for positive correlation values and to the left for negative correlation values.

Red-colored shapes indicate positive correlation and blue-colored shapes indicate negative correlation. The intensity of the red (blue) color indicates the degree of the positive (negative) correlation. Specifically, the darkest red color corresponds to a positive correlation value between 1.0 and 0.8, whereas the lightest red color corresponds to a positive correlation value between 0.2 and 0.0. Conversely, the darkest blue color corresponds to a negative correlation value between -1.0 and -0.8, whereas the lightest blue color corresponds to a negative correlation value between -0.2 and 0.0.

B. GenCo Cross-Correlations

Table XVIII presents pairwise cross-correlations for GenCo reported marginal costs for the benchmark dynamic 5-bus test case extended to include GenCo learning. The indicated cross-correlations are calculated at the GenCos' dispatch points for the peak-demand hour 17 on the final (1000th) simulated day for 30 different runs.

These GenCo cross-correlations are fairly weak, an indication that the GenCos are not responding in a direct strategic manner to the supply-offer choices of other GenCos. Indeed, the VRE learning algorithm used by the GenCos to determine their daily supply offer choices only takes into account each GenCo's own past net earnings as determined by its own past dispatch and LMP levels. The presence of rival GenCos is not considered.

As will next be shown, stronger patterns are obtained for GenCo-LMP and LMP-LMP cross-correlations.

C. GenCo-LMP Cross Correlations

Table XIX presents pairwise cross-correlations between GenCo reported marginal costs and bus LMPs for the peak-demand hour 17 of day 1000 under the same experimental conditions as in subsection V-B. These cross-correlations indicate a moderately-positive correlation between GenCo 3 and the LMPs at buses 2-4, a negative correlation between GenCo 4 and the LMPs at buses 1 and 5, and a strong positive correlation between GenCo 5 and the LMPs at buses 1 and 5. Note, also, that the final column of values in Table XIX is identical to the final column of values in Table XVIII. What explains these correlation patterns?

One important explanatory factor is branch congestion and direction of branch power flows during hour 17. Recall from section IV that the branch 1-2 connecting bus 1 and bus 2 is typically congested in every hour under learning; an example of this is seen in Table IX. Consequently, buses 2-4 constitute a load-pocket for GenCo 3 located at bus 3. It is therefore not surprising that GenCo 3's reported marginal costs are positively correlated with the LMPs at these load-pocket buses during the peak-demand hour 17.

In addition, the persistent congestion on branch 1-2 results in a negative correlation between the reported marginal cost for GenCo 4 at bus 4 and the LMPs at buses 1 and 5 during the peak-demand hour 17. This happens because the power injected by GenCo 4 during hour 17 substitutes in part for the cheaper power of the marginal GenCos 1 and 5 in servicing load at the load-pocket buses 2-4. This substitution occurs because GenCos 1 and 5 are located at buses 1 and 5 and hence are semi-islanded behind the congested branch 1-2 during hour 17 as dictated by the directions of branch power flows; cf. Table XI.

A second important explanatory factor is limits on generation operating capacities during hour 17, which affect the marginal status of the different GenCos. As previously noted in subsection IV-B, the LMP at each bus with a marginal GenCo is given by the reported marginal cost of this GenCo, whereas the LMP at each bus without a marginal GenCo is given by

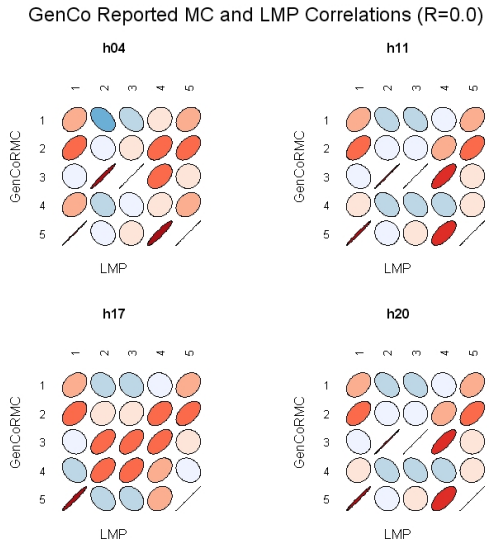

Fig. 18. Pairwise cross-correlations between GenCo reported marginal costs and bus LMPs for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning. Demand for this case is 100% fixed (R=0.0).

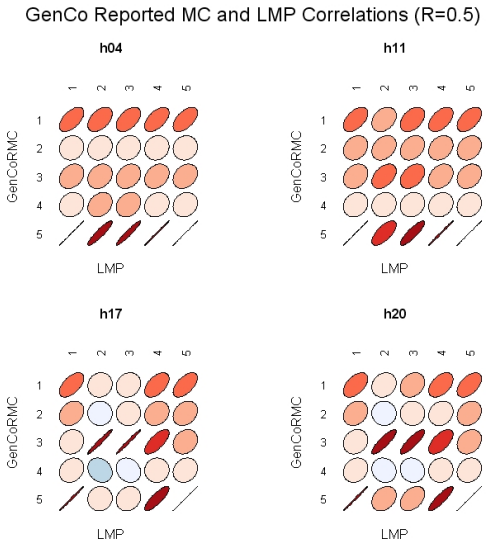

Fig. 19. Pairwise cross-correlations between GenCo reported marginal costs and bus LMP for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and 50% potential price-sensitive demand (R=0.5).

a weighted linear combination of the reported marginal costs of the marginal GenCos.

Table XX reports the frequency (across thirty runs) of each GenCo's marginality during four different hours on day 1000, including the peak-demand hour 17. As indicated, GenCo 5 located at bus 5 is persistently marginal during hour 17, hence the LMP at bus 5 persistently coincides with GenCo 5's reported marginal cost. This explains the finding in Table XIX of a perfect positive correlation of 1.0 between GenCo 5's reported marginal cost and the LMP at bus 5 during hour 17, as well as the appearance of identical final columns of values in Tables XVIII and XIX.

Table XX also indicates that no other GenCo is persistently marginal during hour 17. For example, GenCo 3 is dispatched at maximum operating capacity in 13% of the runs due either to a relatively low reported marginal cost by GenCo 3 or a relatively high reported marginal cost by GenCo 4. This non-marginality of GenCo 3 restrains the positive correlation between GenCo 3's reported marginal costs and the LMPs at the load-pocket buses 2-4 as well as the extent to which power supplied by GenCo 3 can substitute for the power of GenCos 1 and 5 during hour 17.

The correlation diagram in Fig. 18 for the peak-demand hour 17 provides a visualization of the GenCo-LMP cross-correlation findings in Table XIX. In particular, it helps to highlight the importance of GenCos 3 and 4 for the determination of LMPs at the load-pocket buses 2-4, and the importance of GenCo 5 for the determination of LMPs at buses 1 and 5.

The remaining correlation diagrams in Fig. 18 depict the GenCo-LMP cross-correlations that arise in the off-peak hour 4:00, the shoulder hour 11:00, and the shoulder hour 20:00. Comparing these results to the results depicted in Fig. 18 for hour 17, note that GenCo 3's reported marginal cost is now perfectly positively correlated with the LMP at bus 3 and is

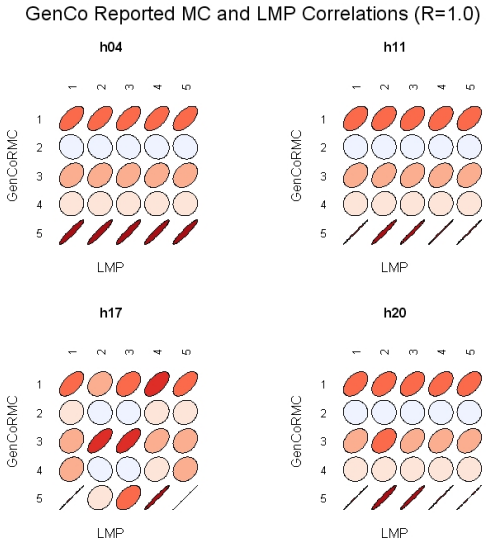

Fig. 20. Pairwise cross-correlations between GenCo reported marginal costs and bus LMPs for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and 100% price-sensitive demand (R=1.0).

strongly positively correlated with the LMPs at its neighboring buses 2 and 4. These changes arise because the substantially lower fixed demand in these three non-peak hours results in the persistent marginality of the relatively large GenCo 3; see Table XX.

Also, in contrast to the peak-demand hour 17, GenCo 4's reported marginal cost is negatively correlated with the LMPs at buses 2 and 3 in the three non-peak hours. This occurs because GenCo 4 is in direct rivalry with the marginal GenCo 3 to supply power to buses 2 and 3 during these non-peak hours. For example, GenCo 4 is dispatched at maximum capacity when its reported marginal cost is relatively low,

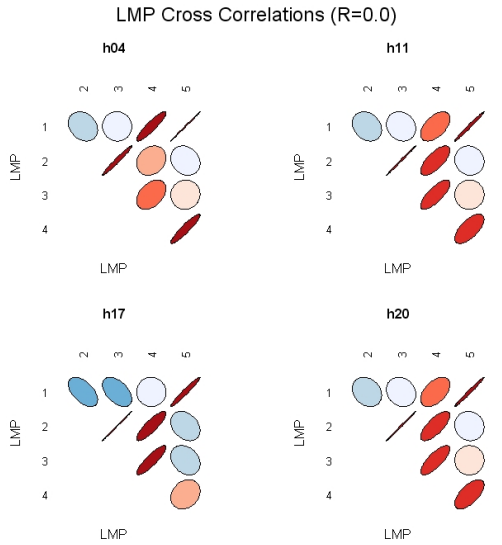

Fig. 21. Pairwise LMP cross-correlations for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning. Demand for this case is 100% fixed (R=0.0).

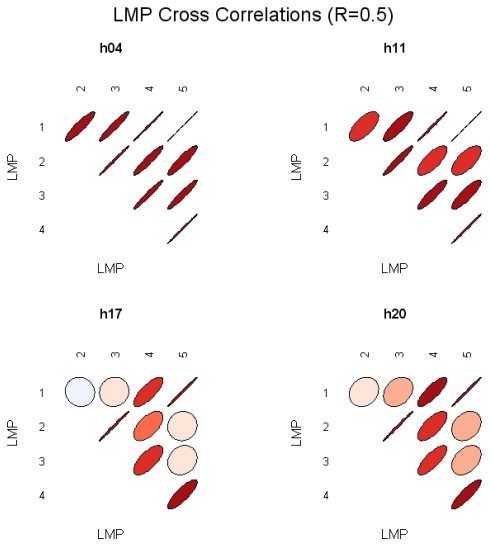

Fig. 22. Pairwise LMP cross-correlations for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and 50% potential price-sensitive demand (R=0.5).

which then permits GenCo 3 to service residual demand at buses 2 and 3 at a relatively high reported marginal cost.

Figures 19 and 20 report the effects on GenCo-LMP cross-correlations when the R ratio measuring the relative price-sensitivity of demand is sytematically increased first to R=0.5 (50% price sensitivity) and then to R=1.0 (100% price sensitivity). As demand becomes more price sensitive, the LSEs more strongly contract their demand in response to price increases and branch congestion becomes less frequent. This limits the ability of the GenCos to profitably exercise economic withholding, which in turn results in dramatically lower reported marginal costs.

In particular, as R increases, the GenCos with relatively low true marginal costs are advantaged and those with relatively high true marginal costs lose out. This can be seen by comparing the correlation diagrams in Figures 18 through 20. As R increases from R=0.0 to R=1.0, the relatively cheap GenCo 5 gains increased influence over each bus LMP while the relatively expensive GenCo 3 loses influence over the load-pocket buses 2 through 4.

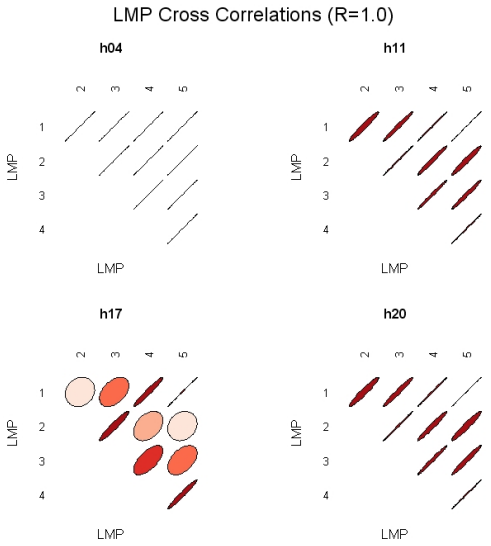

Fig. 23. Pairwise LMP cross-correlations for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and 100% price-sensitive demand (R=1.0).

D. LMP-LMP Cross Correlations

Table XXI reports pairwise cross-correlations for the bus LMPs during the peak-demand hour 17 on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning. Figs. 21 through 23 depict the changes induced in these cross-correlations when the price-sensitivity of demand is systematically increased from R=0.0 (100% fixed) to R=1.0 (100% price sensitive).

The most dominant regularity seen in these LMP correlation results is that the bus LMP cross-correlations become increasingly positive as R increases. This is particularly true for the non-peak hours 04, 11, and 20 with relatively lower LSE fixed demands.

As R increases, a larger portion of LSE total demand is price sensitive. Consequently, the LSEs are able to exercise more resistance to higher prices through demand contraction, which in turn reduces branch congestion. In the current context, bus LMPs are derived from DC OPF solutions with zero losses assumed.[16] Consequently, as congestion diminishes, the LMPs exhibit less separation. In the limit, if all congestion were to disappear, the LMPs would converge to a single uniform price across the grid, which in turn would imply perfect positive correlation among all bus LMPs.

For the non-peak hours 04, 11, and 20, the typical result

[16]See [10] for a rigorous presentation of this LMP derivation.

for the limiting case R=1.0 is no branch congestion. Hence, the bus LMPs during these hours—particularly hour 04—are close to being perfectly positively correlated when R=1.0. For the peak-demand hour 17, however, the branch 1-2 is typically congested even for R=1.0. Consequently, LMP cross-correlations for hour 17 exhibit a strong but not perfect positive correlation.

Another regularity seen in Table XXI and Figs. 21-23 is that the LMP at bus 2 is always strongly positively correlated with the LMP at bus 3. At high R levels, this reflects a lack of branch congestion and hence a lack of LMP separation. At low R levels, however, the branch 1-2 tends to be congested at all hours; cf. Table XI. The congestion on branch 1-2 means that the bulk of the demand at the load-only bus 2 must be supplied along branch 3-2 by the large and frequently marginal GenCo 3. This in turn means that the LMP at bus 2 is most strongly influenced by the LMP at bus 3.

E. Empirical Evidence on LMP Correlations

In this subsection we calculate LMP cross-correlations using real-world price data. In particular, we focus on LMP determination in a neighborhood of the MidAmerican Energy Company (MEC), the largest utility in Iowa.

Through April 2009, MEC was treated as a Balancing Authority (BA) in MISO.[17] A BA is responsible for maintaining load-interchange-generation balance and the support of the Interconnection frequency.

From the geographical map depicted in Fig. 24, we picked four neighboring BAs of MEC in order to study MEC's effect on their LMPs. These BAs are Alliant Energy Corporate Services, Inc. (ALTW), Muscatine Power and Water (MPW), Omaha Public Power District (OPPD), and Nebraska Public Power District (NPPD). We obtained 24-hour historical data from MISO for the real-time and day-ahead LMPs determined for these BAs on August 1, 2, 3 and September 1 in 2008; see [27]. In particular, for ALTW we used the LMP for the loadzone ALTW.MECB, and for the remaining four BAs we used interface LMPs. We then used these data to calculate pairwise cross-correlations between the LMP reported for MEC and the LMPs reported at its four neighboring BAs.

Table XXII reports our LMP cross-correlation findings. All of the LMP cross-correlations are strongly positive. Since MEC is large, and presumably marginal, this suggests that the supply behavior of the MEC could be spilling over to affect the LMPs at neighboring BAs.

On the other hand, as always, care must be taken to recognize potentially confounding effects in real-world data. As noted above, the LMPs reported by MISO for MEC and its four neighboring BAs are load-weighted prices determined for a loadzone and interfaces and not for a single bus. The strong positive LMP cross-correlations in Table XXII could be a statistical artifact arising from the particular load-weighting method employed. Alternatively, they could indicate a lack of branch congestion during the selected days arising either through happenstance or through deliberate ISO planning.

[17]On May 1, 2009, MEC filed an application with the Iowa Utilities Board to become a transmission-owning member of MISO.

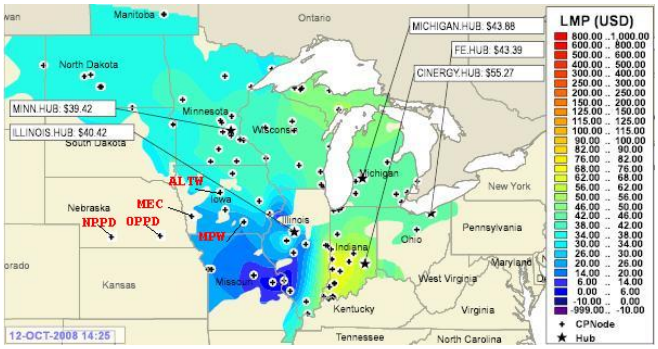

Fig. 24. MidAmerican Energy Company (MEC) Balancing Authority and four neighboring Balancing Authorities in relation to MISO.

To differentiate between these various potential explanations for the strong positive correlations in Table XXII—GenCo spillover effects, statistical artifact, and lack of congestion—we would need to obtain data on MEC supply offers and branch congestion at an hourly level for the selected test days, as well as data giving individuated bus LMPs. To our knowledge, these data are not currently publicly available.

Although agent-based test beds such as AMES can be used to develop interesting hypotheses using simulated scenarios, the real payoff to such development will only come when these hypotheses can be tested more fully against real-world data.

VI. Concluding Remarks

This study has used AMES(V2.02) to conduct a systematic experimental study of LMP separation and volatility under alternative specifications for GenCo learning, demand-bid price sensitivity, and supply-offer price caps. AMES(V2.02) is an agent-based test bed intended to facilitate the study of strategic trading in restructured wholesale power markets for research, teaching, and training purposes.

Particular attention has been focused on dynamic market performance measured in terms of LMPs, dispatch levels, branch flows, LSE payments, GenCo net earnings, and ISO net surplus extractions, and on spatial correlations between differentially situated GenCos and bus LMPs. Future studies will focus on financial and operational risk management issues using AMES(V3.0), an extended version of AMES(V2.02) incorporating a more fully operational two-settlement system, security constrained unit commitment, and enhanced representations of decision-making by traders and market operators.

Appendix A
AMES DC-OPF Problem Formulation

The DC-OPF problem formulation for the AMES(V2.02) day-ahead market outlined below is applicable for any hour H of any day D+1. Reference to these time dimensions is supressed for ease of notation.

The formulation relies heavily on the demand bid, supply offer, and cost function representations developed in Subsections II-B and II-C. An annotated listing of all of the variables

used in the formulation is given in Tables I and II. A more detailed discussion of this formulation can be found in Sun and Tesfatsion ([6], [14], [15]).

A.1 DC-OPF Objective Function Representation

The *gross surplus of LSE* j (\$/h) corresponding to a price-sensitive demand level p^S_{Lj} (MW) is derived from LSE j's price-sensitive demand bid function as follows:[18]

$$GS_j(p^S_{Lj}) \;=\; \int_0^{p^S_{Lj}} D_j(p)dp \;=\; c_j \cdot p^S_{Lj} - d_j \cdot [p^S_{Lj}]^2 \quad (19)$$

The *total gross surplus* (\$/h) of LSEs is then given by

$$TGS(\mathbf{p}^S_L) \;=\; \sum_{j=1}^{J} GS_j(p^S_{Lj}) \;, \quad (20)$$

where

$$\mathbf{p}^S_L \;=\; \left(p^S_{L1}, p^S_{L2}, \cdots, p^S_{LJ}\right) \quad (21)$$

The *reported avoidable cost of GenCo* i (\$/h) corresponding to a generation level p_{Gi} (MW) is derived from GenCo i's reported marginal cost function as follows:

$$VCost^R_i(p_{Gi}) = \int_0^{p_{Gi}} MC^R_i(p)dp = a^R_i \cdot p_{Gi} + b^R_i \cdot [p_{Gi}]^2 \quad (22)$$

The *reported total avoidable cost* (\$/h) corresponding to a vector of GenCo operating levels

$$\mathbf{p}_G \;=\; (p_{G1}, p_{G2}, \cdots, p_{GI}) \quad (23)$$

is then given by

$$TVC^R(\mathbf{p}_G) \;=\; \sum_{i=1}^{I} VC^R_i(p_{Gi}) \quad (24)$$

Reported total net surplus (\$/h) is calculated by the ISO from the LSE price-sensitive demand bids (if any) and the reported GenCo supply offers, as follows:

$$TNS^R(\mathbf{p}^S_L, \mathbf{p}_G) \;=\; TGS(\mathbf{p}^S_L) - TVC^R(\mathbf{p}_G) \quad (25)$$

Reported total net avoidable cost (\$/h) is calculated as the negative of reported total net surplus:

$$TNC^R(\mathbf{p}^S_L, \mathbf{p}_G) \;=\; -\,TNS^R(\mathbf{p}^S_L, \mathbf{p}_G) \quad (26)$$

The standard DC-OPF problem formulation with price-sensitive demand bids involves the minimization of TNC^R, the reported total net avoidable cost of generation, subject to transmission and generation capacity constraints. In this study we augment this standard objective function by inclusion of a penalty function for voltage angle differences. As carefully explained in Sun and Tesfatsion [6], this augmentation provides a number of advantages based on both physical and mathematical considerations.

These advantages can be summarized as follows. First, the validity of the DC-OPF problem as an AC-OPF approximation relies on an assumption of small voltage angle differences, and the augmented objective function permits this assumption to be subjected to systematic sensitivity tests through variations in the penalty weight. Second, solution differences between the non-augmented and augmented forms of the DC-OPF problem can be reduced to arbitrarily small levels by selecting an appropriately small value for the penalty weight. Third, the augmented DC-OPF problem has a numerically desirable strictly convex quadratic programming form permitting the direct determination of solution values for LMPs and voltage angles as well as for price-sensitive demands, generation levels, and branch flows.

A.2 DC-OPF Problem

The SI form of the DC-OPF problem used in this study is as follows:

Minimize

$$TNC^R(\mathbf{p}^S_L, \mathbf{p}_G) + \mu \left[\sum_{km \in BR} [\delta_k - \delta_m]^2\right] \quad (27)$$

with respect to LSE real-power price-sensitive demands, GenCo real-power generation levels, and voltage angles:

$$p^S_{Lj},\ j = 1, ..., J;\ p_{Gi},\ i = 1, ..., I;\ \delta_k,\ k = 1, ..., K \quad (28)$$

subject to

Real-power balance constraint for each bus k=1,...,K:

$$\sum_{i \in I_k} p_{Gi} \;-\; \sum_{j \in J_k} p^S_{Lj} \;-\; \sum_{km \text{ or } mk \;\in\; BR} P_{km} \;=\; \sum_{j \in J_k} p^F_{Lj} \quad (29)$$

where

$$P_{km} \;=\; [V_o]^2 \cdot B_{km}\ [\delta_k - \delta_m] \quad (30)$$

Real-power thermal constraint for each branch km in BR:

$$|P_{km}| \;\leq\; P^U_{km} \quad (31)$$

Reported real-power operating capacity interval for each GenCo i = 1,...,I:

$$Cap^L_i \;\leq\; p_{Gi} \;\leq\; Cap^{RU}_i \quad (32)$$

Real-power purchase capacity interval for price-sensitive demand for each LSE j = 1,...,J:

$$0 \;\leq\; p^S_{Lj} \;\leq\; SLMax_j \quad (33)$$

Voltage angle setting at reference bus 1:

$$\delta_1 \;=\; 0 \quad (34)$$

The shadow price (Lagrange multiplier) solution for the real power balance constraint (29) at bus k, denoted by LMP$_k$, constitutes the *locational marginal price for bus* k. By the well-known envelope theorem, LMP$_k$ (\$/MWh) measures the

[18]The gross surplus of LSE j corresponding to its *fixed* demand bid p^F_{Lj} is infinite if $p^F_{Lj} > 0$; a vertical demand curve literally implies an infinite willingness to pay. For this reason, the DC-OPF objective function used by the ISO to determine the dispatch of generation only takes into account LSE gross surplus corresponding to price-sensitive demand bids.

change in the minimized DC-OPF objective function ($/h) with respect to a change in fixed demand (MW) at bus k; see [10] for a rigorous discussion. Stated less formally, LMP_k essentially measures the cost of efficiently servicing an additional MW of fixed demand at bus k.

The special DC-OPF case in which all LSE demand bids are fixed (no price-sensitive demand) is handled as follows. First, total gross surplus TGB is omitted from (27), so that TNC^R reduces to TVC^R, i.e., to the reported total avoidable costs of generation. Second, the price-sensitive demand variables p^S_{Lj}, j=1,...,J, are removed from the list (28) of choice variables for the DC-OPF problem.

APPENDIX B
ACTION DOMAIN CONSTRUCTION FOR AMES GENCOS

As explained in Section II-B, at the beginning of each day D each GenCo i must choose a supply offer $\text{s}^R_i = (\text{a}^R_i, \text{b}^R_i, \text{Cap}^{RU}_i)$ to report to the ISO for each hour H of the day D+1 day-ahead market. Each supply offer s^R_i characterizes a reported marginal cost function

$$\text{MC}^R_i(p) = a^R_i + 2b^R_i p \tag{35}$$

defined over a reported operating capacity interval

$$\text{Cap}^L_i \;\le\; p \le\; \text{Cap}^{RU}_i \tag{36}$$

Each GenCo i chooses its supply offers s^R_i from an action domain AD_i with finite positive cardinality M_i. In keeping with the modeling goals of empirical sensibility and computational practicality, the action domain AD_i for each GenCo i is constructed under four simplifying assumptions. First, we assume GenCo i only reports upward-sloping marginal cost functions, i.e., $\text{b}^R_i > 0$. Second, we assume GenCo i only reports non-trivial operating capacity intervals, i.e., $\text{Cap}^L_i < \text{Cap}^{RU}_i$. Third, we assume that GenCo i only reports marginal cost functions that lie on or above its true marginal cost function (5) over the range of its reported operating capacity intervals. Fourth, we assume GenCo i always reports an upper operating capacity limit Cap^{RU}_i that is less than or equal to its true upper operating capacity limit Cap^U_i.

Let a supply offer s^R_i for GenCo i be called *admissible* if the corresponding reported marginal cost function $\text{MC}^R_i(p)$ and reported upper operating capacity limit Cap^{RU}_i are in compliance with these four simplifying assumptions. As shown in Sun and Tesfatsion [6, Appendix], given any positive value for a *slope-start* parameter SS_i for GenCo i, any 4-dimensional vector s^A_i consisting of four components in percentage form can be uniquely mapped into an admissible supply offer s^R_i for GenCo i.

Referring to Table I for more precise variable definitions, one can then construct a matrix AD_i for GenCo i characterized by three integer-valued density-control parameters M1_i, M2_i, and M3_i (with $\text{M1}_i \times \text{M2}_i \times \text{M3}_i = \text{M}_i$) and three range-index parameters RIMax^L_i, RIMax^U_i, and RIMin^C_i in percentage form. The three density-control parameters control the *number* of distinct possible ordinate values a^R_i, slope values b^R_i, and upper operating capacity limits Cap^{RU}_i, respectively, that GenCo i can report. The three range-index parameters control the *range* of possible ordinate values, slope values, and upper operating capacity limits, respectively, that GenCo i can report.

The resulting matrix AD_i then has the following property: For any given $\text{SS}_i > 0$, the M_i rows of this matrix constitute M_i distinct vectors s^A_i in percentage form that can be transformed uniquely into M_i distinct admissible supply offers s^R_i for GenCo i. Consequently, the matrix AD_i effectively constitutes an action domain for GenCo i consisting of M_i admissible supply offers s^R_i. Moreover, if the values for the action domain parameters (M1_i,M2_i,M3_i,RIMax^L_i,RIMax^U_i,RIMin^C_i,SS_i) are set identically across the GenCos, and if the above supply-offer construction is then applied for each GenCo i = 1,...,I, the result is a collection $\{\text{AD}_i : i = 1,...,\text{I}\}$ of GenCo-specific action domains that have equal cardinalities and whose supply-offer elements s^R provide similar densities of coverage of the regions lying above the GenCos' true marginal cost curves.

As indicated in Table IV, in this study we set the action domain parameters identically across the GenCos to ensure equal cardinalities and similar densities of their action domains. In addition, we construct the first row of each action domain AD_i to correspond to GenCo i's true cost and capacity attributes (a_i, b_i, Cap^U_i), meaning that GenCo i always has the option of reporting its true marginal cost function and true operating capacity interval to the ISO.

APPENDIX C
AMES GENCO LEARNING

The AMES(V2.02) GenCos are autonomous energy traders with strategic learning capabilities. Each GenCo i adaptively chooses its supply offers ("actions") $\text{s}^R_i = (\text{a}^R_i, \text{b}^R_i, \text{Cap}^{RU}_i)$ from its action domain AD_i on the basis of its own past daily net earnings outcomes. This adaptive choice is implemented by means of a variant of a stochastic reinforcement learning algorithm developed by Roth and Erev ([16], [17]) based on human-subject experiments, hereafter referred to as the *VRE-RL* algorithm. This section describes the implementation of the VRE-RL algorithm for an arbitrary GenCo i starting from the initial day D=1.

Suppose it is the beginning of the initial day D=1, and GenCo i must choose a supply offer from its action domain AD_i to report to the ISO for the day-ahead market in day D+1. As will be seen below, for learning purposes the only relevant attribute of AD_i is that it has finite cardinality $\text{M}_i \ge 1$. Consequently, it suffices to index the supply offers in AD_i by m = 1,...,M_i.

The *initial propensity* of GenCo i to choose supply offer $m \in AD_i$ is given by $\text{q}_{im}(1)$ for m = 1,...,M_i. AMES 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 $\text{q}_i(1)$, as follows:

$$q_{im}(1) \;=\; q_i(1) \;\text{ for all supply offers } m \in AD_i \tag{37}$$

Now consider the beginning of any day D ≥ 1, and suppose the current propensity of GenCo i to choose supply offer m in AD_i is given by $\text{q}_{im}(\text{D})$. The *choice probabilities* that GenCo

i uses to select a supply offer for day D are then constructed from these propensities as follows:[19]

$$p_{im}(D) = \frac{\exp(q_{im}(D)/T_i)}{\sum_{j=1}^{M_i} \exp(q_{ij}(D)/T_i)} \ , \ m \in AD_i \qquad (38)$$

In (38), T_i is a *temperature parameter* 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 (38) become increasingly peaked over the particular supply offers m having the highest propensity values $q_{im}(D)$, thereby increasing the probability that these supply offers will be chosen.

At the end of day D, the current propensity $q_{im}(D)$ that GenCo i associates with each supply offer m in AD_i is updated in accordance with the following rule. Let m' denote the supply offer 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 (11) 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 supply offer m in AD_i,[20]

$$q_{im}(D+1) \ = \ [1-r_i]q_{im}(D) \ + \ Response_{im}(D) \ , \ (39)$$

where

$$Response_{im}(D) = \begin{cases} [1-e_i] \cdot NE_{im'}(D) & \text{if } \ m = m' \\ \\ e_i \cdot q_{im}(D)/[M_i - 1] & \text{if } \ m \neq m', \end{cases} \qquad (40)$$

and $m \neq m'$ implies $M_i \geq 2$. The introduction of the *recency parameter* r_i in (39) acts as a damper on the growth of the propensities over time. The *experimentation parameter* e_i in (40) permits reinforcement to spill over to some extent from a chosen supply offer to other supply offers to encourage continued experimentation with various supply offers in the early stages of the learning process.

In summary, the complete VRE-RL algorithm applied to GenCo i in AMES(V2.02) is fully characterized once user-specified values are set for $(M_i, q_i(1), T_i, r_i, e_i)$, where: M_i denotes the number of supply offer choices available to GenCo i in its action domain AD_i; $q_i(1)$ denotes the initial propensity level in (37); T_i denotes the temperature parameter in (38); r_i denotes the recency parameter in (39); and e_i denotes the experimentation parameter in (40). It is interesting to note, in particular, that this VRE-RL algorithm is well-defined for any action domain AD_i consisting of finitely many elements, regardless of the precise form of these elements.

[19]In the original algorithm developed by Roth and Erev ([16],[17]), the choice probabilities are defined in terms of relative propensity levels. Here, instead, use is made of a "simulated annealing" formulation in terms of exponentials.

[20]The response function appearing in (39) modifies the response function appearing in the original algorithm developed by Roth and Erev ([16],[17]). The modification is introduced to ensure that learning (updating of choice probabilities) occurs even in response to zero-profit outcomes, which are particularly likely to arise in initial periods when GenCo i is just beginning to experiment with different supply offers and the risk of overbidding to the point of non-dispatch is relatively high. See Nicolaisen et al. [28] and Pentapalli [22] for detailed motivation, presentation, and comparative experimental tests of this modified response function.

APPENDIX D
DYNAMIC MARKET PERFORMANCE MEASURES

This appendix section explains more carefully the construction of the dynamic market performance measures reported in the tables and figures for Section IV.

In the absence of GenCo learning, the dynamic 5-bus test case generates a deterministic 24-hour dispatch and LMP schedule for the day-ahead market that is repeated from one day to the next. Table XII presents average hourly outcomes for the dynamic 5-bus test case with no GenCo learning under various demand conditions. Specifically, for each R value in Table XII: (a) *Avg LMP* ($/MWh) denotes LMP averaged across the five buses and the 24 hours of a typical day-ahead market schedule; (b) *Avg Total Demand* (MW) denotes LSE total demand averaged across the 24 hours of a typical day-ahead-market schedule; (c) *Avg True TVCost* ($/h) denotes GenCo true total avoidable cost averaged across the five GenCos and the 24 hours of a typical day-ahead market schedule; and (d) *Avg LI* (unit-free number) denotes the GenCo Lerner Index value averaged across the five GenCos and the 24 hours of a typical day-ahead-market schedule.

For each learning treatment, we calculate mean outcomes for the average hourly measures (a) through (d) across thirty runs. Mean outcomes for measures are indicated by overlines.

More precisely, for each learning treatment we conduct thirty runs corresponding to 30 different random seeds. The length of each run is 1000 days, and only day-1000 data are used for mean-value calculations. Note, however, that these day-1000 calculations refer to scheduled outcomes for the day-ahead market on the subsequent day 1001.

The mean outcomes for *Avg LMP* ($/MWh) presented in Table XIII are calculated from the day-1000 LMP outcomes $LMP_k(H,r)$ conditioned on bus (k), hour (H), and run (r), as follows. First, for each transmission grid bus and each hour of day 1000, determine the average hourly LMP across all 30 runs. Second, for each hour of day 1000, determine the average of these run-averaged hourly LMP values across all five buses. Finally, average these bus-averaged and run-averaged hourly LMP values across all 24 hours of day 1000 to get mean Avg LMP. Thus:

$$\overline{AvgLMP} \ = \ \frac{\left[\sum_{r=1}^{30} \sum_{k=1}^{5} \sum_{H=00}^{23} LMP_k(H,r)\right]}{30 * 5 * 24} \ . \ (41)$$

The corresponding standard deviation is then calculated using the "N" definition (i.e., division by the total number N=[30*5*24] of summed terms rather than N-1), as follows:

$$\sqrt{\frac{\left[\sum_{r=1}^{30} \sum_{k=1}^{5} \sum_{H=00}^{23} [LMP_k(H,r) - AvgLMP]^2\right]}{30 * 5 * 24}} \ . \qquad (42)$$

The mean outcomes for *Avg Total Demand* (MW) presented in Table XIII are calculated from day-1000 data as follows. First, for each of the three LSEs and for each hour of day 1000, determine the LSE's average cleared (satisfied) price-sensitive

demand across all 30 runs. Second, for each LSE and each hour of day 1000, add the LSE's fixed demand and average cleared price-sensitive demand to get the LSE's average total demand. Third, for each hour, sum these LSE average total demands across the three LSEs to get average total demand. Finally, average these hourly average total demands across all 24 hours of day 1000 to get mean Avg Total Demand. The corresponding standard deviation is then calculated in the usual way using the "N" definition.

The mean outcomes for *Avg RepTVCost* ($/h) presented in Table XIII are calculated from day-1000 data as follows. First, for each of the five GenCos and for each hour of day 1000, determine the reported total avoidable costs of the five GenCos averaged across all 30 runs based on the GenCos' *reported* cost and capacity attributes together with their corresponding hourly dispatch levels as determined by the ISO.[21] Second, for each hour of day 1000, determine the average of these run-averaged reported total avoidable cost calculations across all five GenCos. Third, average these GenCo-averaged and run-averaged hourly reported total avoidable cost calculations across all 24 hours of day 1000 to get mean Avg RepTVCost. The corresponding standard deviation is then calculated in the usual way using the "N" definition.

The *Lerner Index (LI)* for any GenCo i supplying a positive amount of (real) power p_{Gi} at bus $k(i)$ during some hour H of some day D is defined as follows:

$$LI_i \;=\; \frac{\left[LMP_{k(i)} - MC_i(p_{Gi})\right]}{LMP_{k(i)}} \;. \tag{43}$$

In (43), $\mathrm{LMP}_{k(i)}$ denotes the LMP at bus $k(i)$, and $\mathrm{MC}_i(\mathrm{p}_{Gi})$ denotes GenCo i's true marginal cost evaluated at p_{Gi}.

The mean outcomes for *Avg LI* (unit-free number) presented in Table XIII are calculated from day-1000 data as follows. First, for each run, for each hour of day 1000, and for each GenCo i with a positive power dispatch level p_{Gi} for this run and hour, determine the GenCo's Lerner Index (43). Second, for each GenCo and each hour of day 1000, determine the average of this GenCo's Lerner Indices across all of the runs for which he had a positive power dispatch level for this hour. Third, for each hour of day 1000, determine the average of these run-averaged Lerner Indices across all GenCos who were dispatched during this hour for at least one run. Finally, determine the average of these GenCo-averaged and run-averaged Lerner Indices across all 24 hours of day 1000 to get mean Avg LI. The corresponding standard deviation is then calculated in the usual way using the "N" definition.

In supply-offer price cap experiments with GenCo learning, an *inadequacy event (IE)* occasionally occurs in some hours in the sense that total GenCo reported capacity is insufficient to meet total fixed demand.[22] For hours in which IEs occur, it is assumed that all fixed demand is met with reserve generation priced at 1000 ($/MWh).

The mean outcomes for LMP spiking ($/MWh) depicted in Figure 17 for learning GenCos under different supply-offer price caps are calculated from day-1000 data with LMP set to the reserve price for hours in which an IE occurs. More precisely, for each run r and for each of the five transmission grid buses k, *LMP spiking* for run r and bus k is first calculated as the maximum absolute difference between successive hourly bus-k LMPs across all 24 hours of day 1000. Next, for each bus k, the average of these LMP spiking measures is determined across all 30 runs r. Finally, the average of these run-averaged LMP spiking measures across all five buses is determined to get mean LMP spiking. The corresponding standard deviation is then calculated in the usual way using the "N" definition.

The mean outcomes for LMP volatility range ($/MWh) depicted in Figure 17 for learning GenCos under different supply-offer price caps are calculated from day-1000 data with LMP set to the reserve price for hours in which an IE occurs. More precisely, for each run r and for each of the five transmission grid buses k, the *LMP volatility range* is calculated as [maxLMP-minLMP] across all 24 hours of day 1000. Second, for each bus k, the average of these LMP volatility range measures is calculated across all 30 runs r. Third, the average of these run-averaged LMP volatility range measures is calculated across all five buses to get the mean LMP volatility range. The corresponding standard deviation is then calculated in the usual way using the "N" definition.

The LMP and IE frequency measures appearing in Table XVI for learning GenCos under different supply-offer price caps are determined for any designated day D as follows:

- *Avg LMP with learning and IE:* This measure reports average hourly LMP for day-D data with IE reserve charges included. Stated more precisely, during any day-D hours in which an IE occurs, i.e., in which offered supply is less than fixed demand, fixed demand is met with reserve generation priced at 1000 ($/MWh). Avg LMP with learning and IE is then calculated across all 24 hours of day D with the LMP for IE hours taken to be 1000 ($/MWh).
- *Avg LMP with learning and w/o IE:* This measure reports average hourly LMP only for those day-D hours in which IEs do not occur. For example, suppose an IE occurs in six of the 24 hours comprising day D, meaning that a well-defined LMP solution is only obtained for each of the remaining 18 hours. Then Avg LMP with learning and w/o IE would be calculated by determining average hourly LMP only for the latter 18 hours of day D.
- *Avg IE with learning:* This measure reports the frequency of IEs for day-D data. For example, suppose that an IE occurs in six of the 24 hours comprising day D. Then Avg IE with learning would be reported as 100% $\times$ [6/24] = 25%.

The mean outcomes for these measures reported in Table XVI are found by averaging their values across all thirty runs using only day D=1000 data.

[21] More precisely, these calculations are made by the ISO using the marginal cost functions *reported* by GenCos to the ISO as part of their reported supply offers, because these are the functions actually used by the ISO in its DC-OPF problems in an attempt to achieve efficient dispatch levels. The ISO does not directly observe the GenCos' true marginal cost functions.

[22] As discussed in Li et al. [7], IEs are an important hidden cost of price caps, since in practice they would require special actions to be taken by the ISO (e.g., reserve procurement, load shedding). AMES(V2.02) reports a tick-count of IEs. Future versions of AMES will incorporate adequacy protection procedures based on empirical ISO practices.

ACKNOWLEDGMENT

The authors are grateful to Chen-Ching Liu and other members of the ISU Electric Energy Economics (E3) Group for helpful discussions on topics related to this paper.

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] A. J. Conejo, E. Castillo, R. Minguez, and F. Milano, "Locational marginal price sensitivities," *IEEE Transactions on Power Systems*, vol. 20, no. 4, pp. 2026–2033, 2005.

[4] T. Orfanogianni and G. Gross, "A general formulation for LMP evaluation," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1163–1173, 2007.

[5] T. Sueyoshi and G. R. Tadiparthi, "Wholesale power price dynamics under transmission line limits: A use of agent-based simulator," *IEEE Transactions on Systems, Man, and Cybernetics*, Part C: Applications and Reviews, vol. 38, no. 2, pp. 229-241, 2008.

[6] 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. www.econ.iastate.edu/tesfatsi/DynTestAMES.JSLT.pdf

[7] H. Li, J. Sun, and L. Tesfatsion, "Dynamic LMP reponse under alternative price-cap and price-sensitive demand scenarios," *Proceedings*, IEEE Power and Energy Society General Meeting, Pittsburgh, PA, July 2008.

[8] 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

[9] MISO. (2008) Homepage. Midwest ISO, Inc. [Online]. www.midwestiso.org/

[10] H. Liu, L. Tesfatsion, and A. A. Chowdhury, "Derivation of locational marginal prices for restructured wholesale power markets," *Journal of Energy Markets*, vol. 2, no. 1, pp. 3–27, 2009.

[11] ISO-NE. (2008) Homepage. ISO New England, Inc. www.iso-ne.com/

[12] L. Tesfatsion, "The AMES wholesale power market test bed as a stochastic dynamic state-space game," Working Paper, Department of Economics, July 2008. http://www.econ.iastate.edu/tesfatsi/AMESStateSpaceModel.pdf

[13] R. Baldick and W. Hogan,"Capacity constrained supply function equilibrium models of electricity markets: Stability, non-decreasing constraints, and function space iterations," University of California Energy Institute Working Paper PWP-089, Revised August 2002. http://www.ucei.berkeley.edu/ucei/PDF/pwp089.pdf

[14] 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. www.econ.iastate.edu/tesfatsi/DC-OPF.JSLT.pdf

[15] J. Sun and L. Tesfatsion, "DC-OPF formulation with price-sensitive demand bids," Working Paper, ISU Economics Department, Revised March 19, 2008. www.econ.iastate.edu/tesfatsi/DC-OPF.PriceSensitiveDemandBids.pdf

[16] 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.

[17] 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.

[18] C. Gieseler, "A java reinforcement learning module for the repast toolkit: Facilitating study and experimentation with reinforcement learning in social science multi-agent simulations," Department of Computer Science, Iowa State University, M.S. Thesis, 2005.

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

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

[21] H. Li and L. Tesfatsion, "Capacity withholding in restructured wholesale power markets: An agent-based test bed study," *Proceedings*, IEEE Power Systems & Exposition Conference, Seattle, WA, March 2009.

[22] 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.

[23] L. Tesfatsion,"Auction basics for restructured wholesale power markets: Objectives and pricing rules," *Proceedings*, IEEE Power and Energy Society General Meeting, Calgary, Alberta, CA, July 2009.

[24] S. Rassenti, V. L. Smith, and B. Wilson, "Controlling market power and price spikes in electricity networks: Demand-side bidding," *Proceedings of the National Academy of Sciences*, vol. 100, no. 5, pp. 2998–3003, 2003.

[25] A. Somani and L. Tesfatsion, "An agent-based test bed study of wholesale power market performance measures," *IEEE Computational Intelligence Magazine*, vol. 3, no. 4, pp. 56-72, 2008.

[26] G. B. Alderete, "Alternative models to analyze market power and financial transmission rights in electricity markets," Ph.D. Thesis, Department of Electrical and Computer Engineering, University of Waterloo, Ontario, 2005.

[27] H. Li and L. Tesfatsion MISO Energy Project homepage, Data section http://www.econ.iastate.edu/tesfatsi/MISOenergygroup.htm#Data

[28] J. Nicolaisen, V. Petrov, and L. Tesfatsion,"Market power and efficiency in a computational electricity market with discriminatory double-auction pricing," *IEEE Transactions on Evolutionary Computation* vol. 5, no. 5, pp. 504–523, 2001.

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.

Junjie Sun received his Ph.D. degree in Economics from Iowa State University in 2006. He is currently a Financial Economist with the Office of the Comptroller of the Currency, U.S. Treasury, Washington D.C. His principal research areas are financial economics, industrial organization, and applied econometrics, with a particular focus on 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, Mathematics, and Electrical and Computer Engineering 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
ADMISSIBLE EXOGENOUS VARIABLES AND FUNCTIONAL FORMS

Variable	Description	Admissibility Restrictions
K	Total number of transmission grid buses	$K > 0$
N	Total number of physically distinct network branches	$N > 0$
J	Total number of LSEs	$J > 0$
I	Total number of GenCos	$I > 0$
J_k	Set of LSEs located at bus k	Card($\cup_{k=1}^{K} J_k$) = J
I_k	Set of GenCos located at bus k	Card($\cup_{k=1}^{K} I_k$) = I
S_o	Base apparent power (in three-phase MVA)	$S_o \geq 1$
V_o	Base voltage (in line-to-line kV)	$V_o > 0$
V_k	Voltage magnitude (kV) at bus k	$V_k = V_o$
km	Branch connecting buses k and m (if one exists)	$k \neq m$
BR	Set of all physically distinct branches km, $k < m$	BR $\neq \emptyset$
x_{km}	Reactance (ohm) for branch km	$x_{km}=x_{mk} > 0$, km in BR
B_{km}	$[1/x_{km}]$ for branch km	$B_{km}=B_{mk} > 0$, $km \in$ BR
P_{km}^{U}	Thermal limit (MW) for real power flow on km	$P_{km}^{U} > 0$, $km \in$ BR
δ_1	Voltage angle (radians) at specified angle reference bus 1	$\delta_1 = 0$
μ	Penalty weight (\$/h·radian) for voltage angle differences in DC-OPF objective function	$\mu > 0$
R_j	Ratio of max potential price-sensitive demand to max potential total demand for LSE j	$0 \leq R_j \leq 1$
BP_{Lj}^{F}(H)	Benchmark-case hour-H fixed demand (MW) for LSE j	$BP_{Lj}^{F}(H) > 0$
p_{Lj}^{F}(H)	Actual hour-H fixed demand (MW) for LSE j	$p_{Lj}^{F}(H) = [1\text{-}R_j]*BP_{Lj}^{F}(H)$
$SLMax_j$(H)	Hour-H upper limit for LSE j's price-sensitive demand (MW)	$SLMax_j(H) = R_j*BP_{Lj}^{F}(H)$
$MPTD_j$(H)	Hour-H maximum potential total demand (MW) for LSE j	$MPTD_j(H)= [p_{Lj}^{F}(H)+SLMax_j(H)]$
c_j(H),d_j(H)	Hour-H demand coefficients (\$/MWh,\$/MW^{2}h) for LSE j	$c_j(H),d_j(H) > 0$
D_{jH}(p)	$D_{jH}(p) = c_j(H) - 2d_j(H)p$ = LSE j's hour-H price-sensitive demand fct for real power p	$D_{jH}(SLMax_j(H)) \geq 0$
$SCost_i$	Hourly pro-rated sunk cost (\$/h) for GenCo i	$SCost_i \geq 0$
Cap_i^{L}	Lower real power operating capacity limit (MW) for GenCo i	$Cap_i^{L} \geq 0$
Cap_i^{U}	Upper real power operating capacity limit (MW) for GenCo i	$Cap_i^{U} > 0$
a_i,b_i	Cost coefficients (\$/MWh,\$/MW^{2}h) for GenCo i	$b_i > 0$
MC_i(p)	$MC_i(p) = a_i+2b_ip$ = GenCo i's true MC function for real power p	$MC_i(Cap_i^{L}) > 0$
$InitMoney_i$	Initial money holdings (\$) of GenCo i	$InitMoney_i > 0$
M_i	Cardinality of the action domain AD_i for GenCo i	$M_i \geq 1$
$M1_i$,$M2_i$,$M3_i$	Integer-valued density-control parameters for AD_i construction	$\prod_{j=1}^{3} Mj_i = M_i$
$RIMax_i^{L}$	Ordinate range-index parameter for AD_i construction	$RIMax_i^{L} \in [0, 1)$
$RIMax_i^{U}$	Slope range-index parameter for AD_i construction	$RIMax_i^{U} \in$ [0,1)
$RIMin_i^{C}$	Capacity-withholding range-index parameter for AD_i construction	$RIMin_i^{C} \in (0, 1]$
SS_i	Slope-start control parameter for AD_i construction	$SS_i > 0$
$MaxDNE_i$	Estimate of maximum possible daily net earnings (\$/D) for GenCo i from AD_i	$MaxDNE_i > 0$
$q_i(1)$	Initial propensity (\$/D) for GenCo i (learning)	$q_i(1) \propto MaxDNE_i$
T_i	Temperature parameter for GenCo i (learning)	$T_i > 0$
r_i	Recency parameter for GenCo i (learning)	$0 \leq r_i \leq 1$
e_i	Experimentation parameter for GenCo i (learning)	$0 \leq e_i < 1$
PCap	Price cap (\$/MWh) imposed on GenCo supply offers by ISO	$PCap > 0$

TABLE II
ENDOGENOUS VARIABLES

Variable	Description
p^S_{Lj}	Real-power price-sensitive demand (MW) by LSE j=1,...,J
a^R_i,b^R_i	Cost coefficients (\$/MWh,\$/MW^{2}h) reported by GenCo i=1,...,I
Cap^{RU}_i	Real-power upper operating capacity limit (MW) reported by GenCo i=1,...,I
p_{Gi}	Real-power generation (MW) supplied by GenCo i=1,...,I
TGS	Total gross surpus (\$/h) of LSEs corresponding to their price-sensitive demands
TVC^R	Reported total avoidable cost (\$/h) of GenCos
TNS^R	Reported total net surplus (TGS - TVC^R)
TNC^R	Reported total net avoidable cost (-1 $\times$ TNS^R)
δ_k	Voltage angle (in radians) at bus k = 2,...,K
P_{km}	Real power (MW) flowing in branch km $\in$ BR
LMP_k	Locational marginal price (\$/MWh) at bus k=1,...,K

TABLE III
NUMERICAL INPUT SPECIFICATIONS FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE: NO GENCO LEARNING, 100% FIXED DEMAND, AND NO SUPPLY-OFFER PRICE CAP

Base Values[a]

S_o	V_o
100	10

K[b]	μ[c]
5	0.05

Branch

From	To	MaxCap[d]	x[e]
1	2	250.0	0.0281
1	4	150.0	0.0304
1	5	400.0	0.0064
2	3	350.0	0.0108
3	4	240.0	0.0297
4	5	240.0	0.0297

GenCo i	at bus	$SCost_i$	a_i	b_i	Cap^L_i	Cap^U_i	$InitMoney_i$
1	1	56.90	14.0	0.005	0.0	110.0	\$1M
2	1	0.11	15.0	0.006	0.0	100.0	\$1M
3	3	2267.53	25.0	0.010	0.0	520.0	\$1M
4	4	5.19	30.0	0.012	0.0	200.0	\$1M
5	5	1391.16	10.0	0.007	0.0	600.0	\$1M

LSE j	at bus	$BP^F(00)$[f]	$BP^F(01)$	$BP^F(02)$	$BP^F(03)$	$BP^F(04)$	$BP^F(05)$	$BP^F(06)$	$BP^F(07)$
1	2	350.00	322.93	305.04	296.02	287.16	291.59	296.02	314.07
2	3	300.00	276.80	261.47	253.73	246.13	249.93	253.73	269.20
3	4	250.00	230.66	217.89	211.44	205.11	208.28	211.44	224.33
LSE j	at bus	$BP^F(08)$	$BP^F(09)$	$BP^F(10)$	$BP^F(11)$	$BP^F(12)$	$BP^F(13)$	$BP^F(14)$	$BP^F(15)$
1	2	358.86	394.80	403.82	408.25	403.82	394.80	390.37	390.37
2	3	307.60	338.40	346.13	349.93	346.13	338.40	334.60	334.60
3	4	256.33	282.00	288.44	291.61	288.44	282.00	278.83	278.83
LSE j	at bus	$BP^F(16)$	$BP^F(17)$	$BP^F(18)$	$BP^F(19)$	$BP^F(20)0$	$BP^F(21)$	$BP^F(22)$	$BP^F(23)$
1	2	408.25	448.62	430.73	426.14	421.71	412.69	390.37	363.46
2	3	349.93	384.53	369.20	365.26	361.47	353.73	334.60	311.53
3	4	291.61	320.44	307.67	304.39	301.22	294.78	278.83	259.61

[a]For simplicity, the base apparent power S_o (MVA) and base voltage V_o (kV) are chosen so base impedance Z_o satisfies $Z_o = V_o^2/S_o = 1$.
[b]Total number of buses
[c]Penalty weight μ (\$/h·radian) for voltage angle differences in DC-OPF objective function
[d]Upper limit P^U_{km} (MW) on the magnitude of real power flow in branch km
[e]Reactance x_{km} (ohm) for branch km
[f]BP^F(H) for LSE j: The benchmark-case fixed demand (MW) for LSE j for each hour H from 00 to 23

TABLE IV

ADDITIONAL NUMERICAL INPUT SPECIFICATIONS FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING: ACTION DOMAIN PARAMETER VALUES, LEARNING PARAMETER VALUES, AND RANDOM SEEDS FOR MULTIPLE RUNS

Action Domain Parameters

GenCo i	$M1_i$	$M2_i$	$M3_i$	$RIMax_i^L$	$RIMax_i^U$	$RIMin_i^C$	SS_i
1	10	10	1	0.75	0.75	1.00	0.001
2	10	10	1	0.75	0.75	1.00	0.001
3	10	10	1	0.75	0.75	1.00	0.001
4	10	10	1	0.75	0.75	1.00	0.001
5	10	10	1	0.75	0.75	1.00	0.001

Learning Parameters

GenCo i	r_i	e_i	$MaxDNE_i$	$\alpha = [q_i(1)/MaxDNE_i]$	$\beta = [q_i(1)/T_i]$
1	0.04	0.96	552,949.06	(1, 1/2, 1/4, 1/10, 1/24)	(100, 50, 10, 2, 1, 1/2)
2	0.04	0.96	538,560.96	(1, 1/2, 1/4, 1/10, 1/24)	(100, 50, 10, 2, 1, 1/2)
3	0.04	0.96	4,615,108.99	(1, 1/2, 1/4, 1/10, 1/24)	(100, 50, 10, 2, 1, 1/2)
4	0.04	0.96	2,148,481.92	(1, 1/2, 1/4, 1/10, 1/24)	(100, 50, 10, 2, 1, 1/2)
5	0.04	0.96	2,099,525.76	(1, 1/2, 1/4, 1/10, 1/24)	(100, 50, 10, 2, 1, 1/2)

Random Seeds for All 30 Runs

RunID	InitialSeed	RunID	InitialSeed	RunID	InitialSeed
01	2096966936	11	736815417	21	1831032783
02	2131965672	12	132292439	22	493464018
03	1235967177	13	207226519	23	930068517
04	511529502	14	1522886012	24	856336506
05	1063330821	15	2000909491	25	1205573239
06	870295371	16	808958575	26	794414294
07	1815184757	17	1150478587	27	1183491260
08	1880683622	18	173232596	28	1846539650
09	122209384	19	999975840	29	437363834
10	220366820	20	1616038132	30	2013640491

TABLE V

ADDITIONAL NUMERICAL SPECIFICATIONS FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE LSE PRICE-SENSITIVE DEMAND FUNCTIONS. THE COLUMN FOR EACH LSE j GIVES THE ORDINATE AND SLOPE VALUES (C,D) FOR LSE j FOR EACH HOUR.

Hour	LSE 1	LSE 2	LSE 3
00	(35.50, 0.40)	(31.65, 0.40)	(21.05, 0.40)
01	(33.95, 0.40)	(30.39, 0.40)	(20.60, 0.40)
02	(32.92, 0.40)	(29.55, 0.40)	(20.30, 0.40)
03	(32.40, 0.40)	(29.13, 0.40)	(20.15, 0.40)
04	(31.89, 0.40)	(28.72, 0.40)	(20.00, 0.40)
05	(32.15, 0.40)	(28.93, 0.40)	(20.07, 0.40)
06	(32.40, 0.40)	(29.13, 0.40)	(20.15, 0.40)
07	(33.44, 0.40)	(29.97, 0.40)	(20.45, 0.40)
08	(36.01, 0.40)	(32.06, 0.40)	(21.20, 0.40)
09	(38.08, 0.40)	(33.74, 0.40)	(21.81, 0.40)
10	(38.60, 0.40)	(34.16, 0.40)	(21.96, 0.40)
11	(38.85, 0.40)	(34.37, 0.40)	(22.03, 0.40)
12	(38.60, 0.40)	(34.16, 0.40)	(21.96, 0.40)
13	(38.08, 0.40)	(33.74, 0.40)	(21.81, 0.40)
14	(37.82, 0.40)	(33.53, 0.40)	(21.73, 0.40)
15	(37.82, 0.40)	(33.53, 0.40)	(21.73, 0.40)
16	(38.85, 0.40)	(34.37, 0.40)	(22.03, 0.40)
17	(78.24, 0.40)	(66.07, 0.40)	(32.61, 0.40)
18	(45.55, 0.40)	(39.78, 0.40)	(23.90, 0.40)
19	(39.88, 0.40)	(35.20, 0.40)	(22.33, 0.40)
20	(39.63, 0.40)	(35.00, 0.40)	(22.26, 0.40)
21	(39.11, 0.40)	(34.57, 0.40)	(22.11, 0.40)
22	(37.82, 0.40)	(33.53, 0.40)	(21.73, 0.40)
23	(36.28, 0.40)	(32.28, 0.40)	(21.28, 0.40)

TABLE VI
HOURLY BUS LMPS DURING A TYPICAL DAY FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE. LMP K DENOTES THE LMP AT BUS K.

Hour	LMP 1	LMP 2	LMP 3	LMP 4	LMP 5
00	15.17	35.50	31.65	21.05	16.21
01	15.16	33.95	30.39	20.60	16.13
02	15.16	32.92	29.55	20.30	16.07
03	15.16	32.40	29.13	20.15	16.04
04	15.15	31.89	28.72	20.00	16.01
05	15.16	32.15	28.93	20.07	16.03
06	15.16	32.40	29.13	20.15	16.04
07	15.16	33.44	29.98	20.45	16.10
08	15.17	36.01	32.06	21.20	16.24
09	15.18	38.08	33.74	21.81	16.35
10	15.18	38.60	34.16	21.96	16.38
11	15.18	38.85	34.37	22.03	16.39
12	15.18	38.60	34.16	21.96	16.38
13	15.18	38.08	33.74	21.81	16.35
14	15.17	37.82	33.53	21.73	16.34
15	15.17	37.82	33.53	21.73	16.34
16	15.18	38.85	34.37	22.03	16.39
17	14.02	78.24	66.07	32.61	17.32
18	15.07	45.56	39.78	23.90	16.64
19	15.18	39.88	35.20	22.33	16.45
20	15.18	39.63	35.00	22.26	16.43
21	15.18	39.11	34.57	22.11	16.41
22	15.17	37.82	33.53	21.73	16.34
23	15.17	36.28	32.28	21.28	16.25

TABLE VII
HOURLY BUS LMPS DURING DAY 1000 FOR A TYPICAL RUN (ID=03) OF THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING.

Hour	LMP 1	LMP 2	LMP 3	LMP 4	LMP 5
00	24.65	222.89	185.33	82.04	34.83
01	25.65	202.76	169.20	76.92	34.74
02	26.31	189.45	158.54	73.54	34.68
03	26.64	182.74	153.17	71.83	34.65
04	26.97	176.15	147.89	70.16	34.62
05	26.80	179.45	150.53	71.00	34.64
06	26.64	182.74	153.17	71.83	34.65
07	25.98	196.17	163.92	75.25	34.71
08	24.32	229.48	190.61	83.72	34.85
09	23.00	256.21	212.03	90.52	34.97
10	22.67	262.92	217.40	92.22	35.00
11	22.50	266.22	220.04	93.06	35.01
12	22.67	262.92	217.40	92.22	35.00
13	23.00	256.21	212.03	90.52	34.97
14	23.16	252.92	209.39	89.68	34.95
15	23.16	252.92	209.39	89.68	34.95
16	22.50	266.22	220.04	93.06	35.01
17	15.65	396.16	324.07	125.81	35.18
18	21.67	282.94	233.44	97.31	35.08
19	21.84	279.52	230.70	96.44	35.07
20	22.01	276.23	228.06	95.61	35.05
21	22.34	269.52	222.69	93.90	35.02
22	23.16	252.92	209.39	89.68	34.95
23	24.15	232.90	193.35	84.59	34.87

TABLE VIII
HOURLY GENCO DISPATCH LEVELS DURING A TYPICAL DAY FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE.

Hour	GenCo 1	GenCo 2	GenCo 3	GenCo 4	GenCo 5
00	110.00	13.87	332.53	0.00	443.59
01	110.00	13.44	269.45	0.00	437.54
02	110.00	13.16	227.71	0.00	433.54
03	110.00	13.02	206.66	0.00	431.52
04	110.00	12.87	185.99	0.00	429.54
05	110.00	12.94	196.39	0.00	430.53
06	110.00	13.02	206.66	0.00	431.52
07	110.00	13.30	248.77	0.00	435.55
08	110.00	14.01	353.20	0.00	445.58
09	110.00	14.60	437.02	0.00	453.63
10	110.00	14.73	458.06	0.00	455.64
11	110.00	14.80	468.39	0.00	456.63
12	110.00	14.73	458.06	0.00	455.64
13	110.00	14.60	437.02	0.00	453.63
14	110.00	14.51	426.67	0.00	452.62
15	110.00	14.51	426.67	0.00	452.62
16	110.00	14.80	468.39	0.00	456.63
17	2.07	0.00	520.00	108.88	522.63
18	107.34	6.11	520.00	0.00	474.15
19	110.00	15.08	510.08	0.00	460.63
20	110.00	15.01	499.83	0.00	459.64
21	110.00	14.88	478.75	0.00	457.63
22	110.00	14.51	426.67	0.00	452.62
23	110.00	14.08	363.95	0.00	446.60
Cap^U	110.00	100.00	520.00	200.00	600.00

TABLE IX
HOURLY GENCO DISPATCH LEVELS DURING DAY 1000 FOR A TYPICAL RUN (ID=03) OF THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING.

Hour	GenCo 1	GenCo 2	GenCo 3	GenCo 4	GenCo 5
00	37.43	26.39	316.95	36.74	482.50
01	40.94	30.35	257.04	28.21	473.86
02	43.26	32.97	217.45	22.57	468.16
03	44.42	34.29	197.48	19.72	465.28
04	45.57	35.59	177.86	16.93	462.45
05	45.00	34.94	187.67	18.33	463.87
06	44.42	34.29	197.48	19.72	465.28
07	42.09	31.65	237.42	25.41	471.03
08	36.28	25.09	336.56	39.53	485.32
09	31.62	19.83	416.10	50.86	496.78
10	30.45	18.51	436.06	53.70	499.66
11	29.88	17.86	445.87	55.10	501.07
12	30.45	18.51	436.06	53.70	499.66
13	31.62	19.83	416.10	50.86	496.78
14	32.20	20.48	406.29	49.46	495.37
15	32.20	20.48	406.29	49.46	495.37
16	29.88	17.86	445.87	55.10	501.07
17	5.80	0.00	520.00	109.69	518.10
18	26.97	14.57	495.63	62.19	508.25
19	27.56	15.25	485.46	60.74	506.78
20	28.13	15.89	475.66	59.34	505.36
21	29.30	17.22	455.69	56.50	502.49
22	32.20	20.48	406.29	49.46	495.37
23	35.68	24.42	346.73	40.98	486.79
Cap^U	110.00	100.00	520.00	200.00	600.00

TABLE X
HOURLY BRANCH POWER FLOWS DURING A TYPICAL DAY FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE.

Hour	1–2	1–4	1–5	2–3	3–4	4–5
00	250.00	129.65	-255.77	-100.00	-67.47	-187.82
01	250.00	126.71	-253.27	-72.95	-80.31	-184.27
02	250.00	124.77	-251.61	-55.05	-88.81	-181.93
03	250.00	123.79	-250.77	-46.02	-93.09	-180.75
04	250.00	122.83	-249.95	-37.16	-97.30	-179.58
05	250.00	123.30	-250.37	-41.62	-95.19	-180.16
06	250.00	123.79	-250.77	-46.02	-93.09	-180.75
07	250.00	125.74	-252.45	-64.09	-84.52	-183.11
08	250.00	130.61	-256.60	-108.86	-63.26	-188.98
09	250.00	134.52	-259.92	-144.80	-46.18	-193.70
10	250.00	135.49	-260.76	-153.83	-41.90	-194.88
11	250.00	135.97	-261.17	-158.25	-39.81	-195.45
12	250.00	135.49	-260.76	-153.83	-41.90	-194.88
13	250.00	134.52	-259.92	-144.80	-46.18	-193.70
14	250.00	134.03	-259.51	-140.37	-48.30	-193.11
15	250.00	134.03	-259.51	-140.37	-48.30	-193.11
16	250.00	135.97	-261.17	-158.25	-39.81	-195.45
17	250.00	98.83	-346.76	-198.62	-63.15	-175.88
18	250.00	137.64	-274.19	-180.73	-29.93	-199.97
19	250.00	137.91	-262.83	-176.14	-31.32	-197.80
20	250.00	137.43	-262.42	-171.73	-33.41	-197.22
21	250.00	136.46	-261.58	-162.70	-37.69	-196.05
22	250.00	134.03	-259.51	-140.37	-48.30	-193.11
23	250.00	131.10	-257.02	-113.48	-61.07	-189.58
Max Cap	250.00	150.00	400.00	350.00	240.00	240.00

TABLE XI
HOURLY BRANCH POWER FLOWS DURING DAY 1000 FOR A TYPICAL RUN (ID=03) OF THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING.

Hour	1–2	1–4	1–5	2–3	3–4	4–5
00	250.00	114.42	-300.66	-100.00	-83.05	-181.91
01	250.00	114.62	-293.38	-72.93	-92.69	-180.54
02	250.00	114.75	-288.57	-55.04	-99.06	-179.64
03	250.00	114.82	-286.15	-46.02	-102.27	-179.18
04	250.00	114.88	-283.77	-37.16	-105.43	-178.74
05	250.00	114.85	-284.96	-41.59	-103.85	-178.96
06	250.00	114.82	-286.15	-46.02	-102.27	-179.18
07	250.00	114.68	-291.00	-64.07	-95.85	-180.09
08	250.00	114.35	-303.04	-108.86	-79.90	-182.35
09	250.00	114.09	-312.70	-144.80	-67.10	-184.16
10	250.00	114.02	-315.12	-153.82	-63.89	-184.62
11	250.00	113.99	-316.32	-158.25	-62.31	-184.84
12	250.00	114.02	-315.12	-153.82	-63.89	-184.62
13	250.00	114.09	-312.70	-144.80	-67.10	-184.16
14	250.00	114.12	-311.51	-140.37	-68.68	-183.94
15	250.00	114.12	-311.51	-140.37	-68.68	-183.94
16	250.00	113.99	-316.32	-158.25	-62.31	-184.84
17	250.00	98.83	-343.09	-198.62	-63.15	-175.09
18	250.00	113.83	-322.36	-180.73	-54.30	-185.97
19	250.00	113.86	-321.12	-176.14	-55.94	-185.74
20	250.00	113.89	-319.93	-171.71	-57.52	-185.52
21	250.00	113.96	-317.51	-162.69	-60.73	-185.06
22	250.00	114.12	-311.51	-140.37	-68.68	-183.94
23	250.00	114.32	-304.27	-113.46	-78.26	-182.58
Max Cap	250.00	150.00	400.00	350.00	240.00	240.00

TABLE XII
AVERAGE HOURLY LMP, TOTAL DEMAND, TRUE TOTAL AVOIDABLE COSTS, AND THE GENCO LERNER INDEX DURING A TYPICAL DAY FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE DEMAND VARYING FROM R=0.0 (100% FIXED) TO R=1.0 (100% PRICE SENSITIVE).

R	Avg LMP	Avg Total Demand	Avg TrueTVCost	Avg LI
0.0	25.18	318.21	3,779.17	0.0056
0.1	24.51	299.19	3,439.32	0.0042
0.2	23.92	279.69	3,100.91	0.0036
0.3	23.33	259.85	2,765.58	0.0032
0.4	22.72	240.18	2,446.54	0.0029
0.5	22.10	220.88	2,143.65	0.0026
0.6	21.35	204.09	1,888.46	0.0022
0.7	20.49	188.67	1,662.19	0.0013
0.8	19.49	175.74	1,481.15	0.0000
0.9	18.27	169.68	1,408.55	0.0000
1.0	17.04	163.87	1,349.49	0.0000

TABLE XIII
MEAN OUTCOMES (WITH STANDARD DEVIATIONS) FOR AVERAGE HOURLY LMP, TOTAL DEMAND, TRUE TOTAL AVOIDABLE COSTS, REPORTED TOTAL AVOIDABLE COSTS, AND THE GENCO LERNER INDEX DURING DAY 1000 FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING AND DEMAND VARYING FROM R=0.0 (100% FIXED) TO R=1.0 (100% PRICE SENSITIVE).

R	$\overline{\text{Avg LMP}}$	$\overline{\text{Avg Total Demand}}$	$\overline{\text{Avg TrueTVCost}}$	$\overline{\text{Avg RepTVCost}}$	$\overline{\text{Avg LI}}$
0.0	140.30 (106.03)	318.21 (0.00)	4,154.01 (3,751.36)	16,045.20 (23,126.74)	0.6347 (0.25)
0.1	128.32 (94.79)	286.39 (0.00)	3,519.94 (3,163.41)	12,492.44 (17,189.76)	0.6092 (0.27)
0.2	58.67 (53.06)	256.23 (8.81)	2,820.88 (3,095.78)	5,641.99 (7,797.29)	0.3792 (0.26)
0.3	42.29 (25.55)	227.52 (15.86)	2,313.91 (2,770.51)	3,895.22 (4,787.31)	0.3485 (0.23)
0.4	37.54 (20.08)	202.89 (24.90)	1,913.74 (2,503.89)	3,056.16 (4,015.67)	0.3231 (0.23)
0.5	32.11 (10.80)	180.25 (32.85)	1,570.93 (2,343.13)	2,469.80 (3,740.36)	0.3054 (0.22)
0.6	29.37 (7.69)	164.55 (42.32)	1,389.99 (2,196.00)	2,133.67 (3,365.10)	0.2769 (0.21)
0.7	28.13 (7.31)	148.94 (52.25)	1,232.78 (2,012.96)	1,857.35 (2,987.34)	0.2654 (0.22)
0.8	26.35 (6.42)	135.30 (62.95)	1,097.52 (1,897.99)	1,604.06 (2,700.79)	0.2425 (0.22)
0.9	24.90 (5.95)	120.56 (73.11)	962.60 (1,753.26)	1,368.97 (2,399.95)	0.2274 (0.21)
1.0	23.34 (5.53)	106.13 (82.63)	832.18 (1,595.86)	1,147.94 (2,099.49)	0.2098 (0.20)

TABLE XIV
TOTAL GENCO DAILY NET EARNINGS, TOTAL GENCO DAILY REVENUES, TOTAL LSE DAILY PAYMENTS, AND ISO DAILY NET SURPLUS ON A TYPICAL DAY FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO PERMIT DEMAND VARYING FROM R=0.0 (100% FIXED) TO R=1.0 (100% PRICE SENSITIVE).

	R=0.0	R=0.2	R=0.4	R=0.6	R=0.8	R=1.0
GenCo 1 DNE	1,556.41	1,412.41	1,316.90	1,239.14	1,193.74	1,145.06
GenCo 2 DNE	26.58	10.93	4.30	1.42	1.21	0.43
GenCo 3 DNE	56,016.09	35,651.85	21,354.23	11,479.86	2,874.96	2,493.13
GenCo 4 DNE	142.27	13.91	0.00	0.00	0.00	0.00
GenCo 5 DNE	34,266.94	32,253.34	30,460.22	28,531.08	26,246.37	23,364.36
Total Gen DNE	92,008.30	69,342.45	53,135.65	41,251.49	30,316.28	27,002.99
GenCo 1 DRev	38,356.90	38,599.53	37,574.43	36,411.03	35,826.48	34,932.06
GenCo 2 DRev	4,801.54	3,082.31	1,912.43	990.22	617.12	303.06
GenCo 3 DRev	321,967.71	229,151.51	144,201.72	74,559.64	24,303.90	17,528.99
GenCo 4 DRev	3,551.07	1,049.37	0.00	0.00	0.00	0.00
GenCo 5 DRev	176,831.32	169,568.69	163,032.42	155,905.91	147,307.23	136,178.18
Total Gen DRev	545,508.54	441,451.41	346,721.00	267,866.80	208,054.73	188,942.29
Total LSE DPay	754,919.61	625,704.76	506,698.47	399,806.50	301,537.97	231,945.71
ISO DNetSurplus	209,411.07	184,253.35	159,977.47	131,939.70	93,483.24	43,003.42

TABLE XV

MEAN OUTCOMES (WITH STANDARD DEVIATIONS) FOR TOTAL GENCO DAILY NET EARNINGS, TOTAL GENCO DAILY REVENUES, TOTAL LSE DAILY PAYMENTS, AND ISO DAILY NET SURPLUS ON DAY 1000 FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING AND DEMAND VARYING FROM R=0.0 (100% FIXED) TO R=1.0 (100% PRICE SENSITIVE).

	R=0.0	R=0.2	R=0.4	R=0.6	R=0.8	R=1.0
GenCo 1 $\overline{\text{DNE}}$	69,219.61 (64,055.42)	21,950.82 (32,888.20)	18,028.37 (20,401.49)	15,317.64 (17,342.48)	11,460.38 (13,341.31)	6,075.72 (8,585.60)
GenCo 2 $\overline{\text{DNE}}$	54,548.72 (57,868.92)	18,919.31 (30,102.78)	13,271.49 (19,648.72)	11,141.69 (15,916.37)	8,368.95 (13,528.49)	5,061.87 (9,487.15)
GenCo 3 $\overline{\text{DNE}}$	1,725,216.72 (389,906.14)	293,743.16 (269,901.79)	41,122.50 (20,776.25)	8,213.84 (7,847.69)	4,059.61 (3,343.84)	2,316.01 (1,775.20)
GenCo 4 $\overline{\text{DNE}}$	321,907.08 (153,782.17)	38,678.95 (73,333.88)	5,589.68 (14,969.93)	66.32 (161.70)	14.11 (51.51)	3.38 (18.22)
GenCo 5 $\overline{\text{DNE}}$	270,754.58 (124,835.20)	167,938.19 (113,128.59)	149,920.04 (85,701.22)	118,535.14 (50,853.37)	83,774.92 (32,392.38)	54,920.77 (20,700.86)
Total Gen $\overline{\text{DNE}}$	2,441,646.71 (153,782.17)	541,230.41 (73,333.88)	227,932.07 (14,969.93)	153,274.62 (161.70)	107,677.99 (51.51)	68,377.76 (18.22)
GenCo 1 $\overline{\text{DRev}}$	93,976.61 (78,884.69)	39,069.66 (46,553.23)	34,172.46 (35,411.75)	30,876.22 (31,825.48)	25,992.32 (27,100.83)	16,407.69 (20,871.22)
GenCo 2 $\overline{\text{DRev}}$	74,751.32 (72,682.55)	32,167.61 (44,047.41)	25,385.25 (34,607.65)	22,521.81 (30,244.82)	18,284.43 (27,425.88)	12,934.50 (22,709.95)
GenCo 3 $\overline{\text{DRev}}$	1,952,910.84 (386,964.13)	432,137.00 (257,439.64)	97,834.40 (37,540.80)	27,830.78 (27,309.84)	13,671.15 (14,680.47)	7,337.80 (6,473.61)
GenCo 4 $\overline{\text{DRev}}$	449,051.68 (195,313.53)	70,968.73 (122,983.56)	14,705.64 (36,263.17)	296.48 (575.30)	77.98 (231.84)	16.88 (90.88)
GenCo 5 $\overline{\text{DRev}}$	372,219.49 (102,726.38)	305,520.62 (106,382.80)	285,483.62 (79,412.59)	238,548.54 (44,763.65)	181,354.58 (47,182.98)	131,542.33 (46,882.23)
Total Gen $\overline{\text{DRev}}$	2,942,909.93 (558,938.79)	879,863.63 (385,241.09)	457,581.36 (113,573.45)	320,073.84 (52,131.79)	239,380.46 (24,347.08)	168,239.20 (25,679.71)
Total LSE $\overline{\text{DPay}}$	5,040,530.89 (1,043,543.03)	1,526,994.60 (975,375.28)	663,801.01 (209,686.70)	377,524.06 (11,366.32)	271,061.40 (26,241.77)	183,118.99 (33,324.23)
ISO $\overline{\text{DNetSurplus}}$	2,097,620.96 (632,303.71)	647,130.97 (633,129.12)	206,219.65 (197,896.93)	57,450.22 (48,696.64)	31,680.94 (30,789.07)	14,879.79 (11,016.23)

TABLE XVI

MEAN OUTCOMES (WITH STANDARD DEVIATIONS) FOR AVERAGE HOURLY LMP AND AVERAGE HOURLY INADEQUACY EVENT (IE) FREQUENCY DURING DAY 1000 FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING AND A SUPPLY-OFFER PRICE CAP VARYING FROM INFINITELY HIGH (NONE) TO LOW (80), BOTH WITH AND WITHOUT IE RESERVE CHARGES.

	No PCap	PCap=160	PCap=140	PCap=120	PCap=100	PCap=80
Avg LMP (w/o learning)	25.18	25.18	25.18	25.18	25.18	25.18
$\overline{\text{Avg LMP}}$ with learning and IE	140.30 (106.03)	126.31 (193.79)	161.77 (273.70)	153.74 (296.70)	238.99 (381.46)	342.05 (442.94)
$\overline{\text{Avg LMP}}$ with learning and w/o IE	140.30 (106.03)	89.65 (75.76)	86.95 (115.17)	50.92 (33.45)	48.31 (28.90)	44.91 (30.54)
$\overline{\text{Avg IE}}$ with learning	0.0%	4.0%	8.2%	7.5%	17.8%	31.1%

TABLE XVII
MEAN OUTCOMES (WITH STANDARD DEVIATIONS) FOR HOURLY BUS LMPS DURING DAY 1000 FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING.

Hour	LMP 1	LMP 2	LMP 3	LMP 4	LMP 5
00	41.72 (25.27)	248.42 (52.52)	209.25 (41.67)	101.56 (20.52)	52.33 (23.46)
01	43.09 (25.19)	216.75 (41.20)	183.85 (32.56)	93.37 (18.93)	52.01 (23.47)
02	44.05 (25.24)	195.93 (34.02)	167.16 (26.82)	88.02 (18.25)	51.85 (23.58)
03	44.55 (25.29)	185.43 (30.52)	158.74 (24.04)	85.34 (18.01)	51.78 (23.66)
04	45.03 (25.36)	175.14 (27.35)	150.49 (21.53)	82.70 (17.86)	51.71 (23.76)
05	44.79 (25.32)	180.29 (28.92)	154.62 (22.76)	84.02 (17.92)	51.74 (23.71)
06	44.55 (25.29)	185.43 (30.52)	158.74 (24.04)	85.34 (18.01)	51.78 (23.66)
07	43.57 (25.21)	206.43 (37.60)	175.58 (29.68)	90.72 (18.56)	51.93 (23.52)
08	41.36 (25.41)	258.93 (56.28)	217.71 (44.70)	104.35 (21.19)	52.53 (23.58)
09	39.90 (25.27)	301.37 (70.80)	251.83 (56.70)	115.60 (24.63)	53.32 (23.57)
10	39.52 (25.41)	311.91 (74.57)	260.30 (59.75)	118.38 (25.51)	53.50 (23.71)
11	39.33 (25.48)	317.10 (76.45)	264.47 (61.28)	119.75 (25.95)	53.59 (23.78)
12	39.52 (25.41)	311.91 (74.57)	260.30 (59.75)	118.38 (25.51)	53.50 (23.71)
13	39.90 (25.27)	301.37 (70.80)	251.83 (56.70)	115.60 (24.63)	53.32 (23.57)
14	40.08 (25.20)	296.20 (68.97)	247.67 (55.21)	114.23 (24.21)	53.23 (23.51)
15	40.08 (25.20)	296.20 (68.97)	247.67 (55.21)	114.23 (24.21)	53.23 (23.51)
16	39.33 (25.48)	317.10 (76.45)	264.47 (61.28)	119.75 (25.95)	53.59 (23.78)
17	35.99 (23.13)	395.18 (95.80)	327.12 (76.66)	139.98 (28.45)	54.42 (21.24)
18	38.28 (25.10)	346.50 (84.16)	288.10 (67.71)	127.51 (28.09)	54.10 (23.59)
19	38.59 (24.99)	338.62 (83.24)	281.77 (67.02)	125.45 (28.04)	53.99 (23.54)
20	38.77 (24.93)	333.28 (81.38)	277.48 (65.54)	124.03 (27.65)	53.89 (23.49)
21	39.15 (25.55)	322.39 (78.32)	268.73 (62.79)	121.15 (26.35)	53.68 (23.84)
22	40.08 (25.20)	296.20 (68.97)	247.67 (55.21)	114.23 (24.21)	53.23 (23.51)
23	41.17 (25.49)	264.39 (58.24)	222.10 (46.27)	105.80 (21.55)	52.63 (23.63)

TABLE XVIII
PAIRWISE CROSS-CORRELATIONS BETWEEN GENCO REPORTED MARGINAL COSTS AT THE PEAK-DEMAND HOUR 17 OF DAY 1000 FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING.

	GenCo 1	GenCo 2	GenCo 3	GenCo 4	GenCo 5
GenCo 1	1.0000	0.1254	-0.3412	-0.0588	0.2879
GenCo 2		1.0000	-0.0355	0.1131	0.5042
GenCo 3			1.0000	-0.3518	0.0163
GenCo 4				1.0000	-0.1718
GenCo 5					1.0000

TABLE XIX
PAIRWISE CROSS-CORRELATIONS BETWEEN GENCO REPORTED MARGINAL COSTS AND BUS LMPS AT THE PEAK-DEMAND HOUR 17 OF DAY 1000 FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING.

	LMP 1	LMP 2	LMP 3	LMP 4	LMP 5
GenCo 1	0.3136	-0.2244	-0.2143	-0.0718	0.2879
GenCo 2	0.4150	0.1344	0.1591	0.4148	0.5042
GenCo 3	-0.1164	0.5147	0.5222	0.5363	0.0163
GenCo 4	-0.2711	0.4641	0.4625	0.3811	-0.1718
GenCo 5	0.9704	-0.3125	-0.2712	0.2293	1.0000

TABLE XX
FREQUENCY OF GENCO MARGINALITY ACROSS 30 RUNS MEASURED AT FOUR DIFFERENT HOURS ON DAY 1000 FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING.

	GenCo 1	GenCo 2	GenCo 3	GenCo 4	GenCo 5
H04	13%	37%	100%	37%	100%
H11	10%	30%	100%	20%	100%
H17	10%	23%	87%	20%	100%
H20	10%	30%	100%	13%	100%

TABLE XXI
PAIRWISE CROSS-CORRELATIONS BETWEEN BUS LMPS AT THE PEAK-DEMAND HOUR 17 OF DAY 1000 FOR THE BENCHMARK DYNAMIC 5-BUS TEST CASE EXTENDED TO INCLUDE GENCO LEARNING.

	LMP 1	LMP 2	LMP 3	LMP 4	LMP 5
LMP 1	1.0000	-0.5328	-0.4957	-0.0127	0.9704
LMP 2		1.0000	0.9991	0.8530	-0.3125
LMP 3			1.0000	0.8747	-0.2712
LMP 4				1.0000	0.2293
LMP 5					1.0000

TABLE XXII
PAIRWISE CROSS-CORRELATIONS BETWEEN REAL-TIME AND DAY-AHEAD MARKET LMPS FOR THE MIDAMERICAN ENERGY CORPORATION (MEC) AND FOUR NEIGHBORING BALANCING AUTHORITIES.

	DA (8/1/08)	DA (8/2/08)	DA (8/3/08)	DA (9/1/08)	RT (8/1/08)	RT (8/2/08)	RT (8/3/08)	RT (9/1/08)
MEC–ALTW	0.998	0.997	0.999	1.000	0.994	0.971	0.974	1.000
MEC–MPW	0.996	0.994	0.998	1.000	0.996	0.970	0.973	1.000
MEC–OPPD	1.000	1.000	0.999	1.000	0.996	0.986	0.973	1.000
MEC–NPPD	0.998	0.998	0.995	0.998	0.983	0.930	0.824	1.000