# Dynamic LMP Response Under Alternative Price-Cap and Price-Sensitive Demand Scenarios

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

***Abstract*—This study investigates the complicated nonlinear effects of demand-bid price sensitivity and supply-offer price caps on Locational Marginal Prices (LMPs) for bulk electric power when profit-seeking generators can learn over time how to strategize their supply offers. Systematic computational experiments are conducted using AMES, an open-source agent-based test bed developed by the authors. AMES models a restructured wholesale power market operating through time over an AC transmission grid subject to line constraints, generation capacity constraints, and strategic trader behaviors.**

***Index Terms*—Restructured wholesale power markets, agent-based test bed, locational marginal prices, demand-bid price sensitivity, supply-offer price caps, learning, strategic pricing, capacity withholding, market power, price spiking, price volatility, AMES, MISO market protocols**

## I. Introduction

IN April 2003 the U.S. Federal Energy Regulatory Commission issued a white paper [2] proposing a template for the restructuring of U.S. wholesale power markets, referred to as the *Wholesale Power Market Platform (WPMP)*. As detailed in [3], versions of the WPMP have been implemented (or scheduled for implementation) in the midwest (MISO), New England (ISO-NE), New York (NYISO), the mid-atlantic states (PJM), California (CAISO), the southwest (SPP), and Texas (ERCOT).

A core design element of the WPMP is a two-settlement system to be managed by an independent system operator (ISO). Roughly, a "two-settlement system" refers to the combined workings of a day-ahead energy market and a real-time energy market that are separately settled each day by means of Locational Marginal Pricing (LMP). Under LMP, a separate price for power is determined at each node of the transmission grid at which power is injected or withdrawn.

As envisioned in the WPMP, and implemented in practice, the day-ahead market is structured as a double auction. Load-serving entities (buyers) are permitted to submit demand bids that include price-sensitive hourly demands, and generators (sellers) are permitted to submit supply offers that include price-sensitive hourly supplies.

In actuality, however, the day-ahead market effectively functions as a one-sided auction because the bulk of the demand takes the form of fixed hourly loads (i.e., load profiles) implying essentially vertical demand curves. A key difficulty is that downstream retail markets are still largely regulated with cost-based pricing, so that load-serving entities in fact have little incentive to submit price-sensitive demand bids. As demonstrated in [4] using human-subject experiments, and in [5] using computational agent experiments, under this scenario electric power generators can easily learn to tacitly collude on reported supply offers higher than true marginal costs. The result is dramatically higher LMPs, hence substantially higher market operation costs.

In this study we use an agent-based test bed – AMES (V2.0) – to investigate how demand-bid price sensitivity and supply-offer price caps affect dynamic wholesale power market performance, with a particular stress on LMP response. AMES implements a wholesale power market operating through time over an AC transmission grid in accordance with core WPMP design features as implemented by the MISO [6].[1]

In particular, the AMES Load-Serving Entities (LSEs) and generators report daily demand bids and supply offers to the AMES ISO for the day-ahead market. The LSEs' demand bids are mixtures of fixed (price-insensitive) demands and price-sensitive demands.[2] The generators' reported supply offers consist of price-sensitive marginal cost functions defined over operating capacity intervals.

The AMES ISO uses these daily reported bids and offers to determine hourly LMPs and commitment levels for the next day as solutions to hourly DC optimal power flow problems. The AMES generators use their daily settlement payments for the day-ahead market to adjust their daily reported supply offers via reinforcement learning. The AMES ISO has the option to impose a price cap on reported supply offers in an attempt to mitigate the exercise of market power by generators.[3]

Section II provides a fuller description of the main features of AMES (V2.0). Section III explains the AMES experimental design used to determine dynamic market performance under systematically varied settings for the following three treatment factors: (a) the degree to which demand bids are price sensitive

Latest revision: 19 April 2008. This work, an abridged version of Li Sun, and Tesfatsion [1], is scheduled to appear in the *Proceedings* of the IEEE Power Engineering Society General Meeting, Pittsburgh, 2008. It has been supported in part by the National Science Foundation under Grant NSF-0527460 and by a grant from the ISU Electric Power Research Center.

Hongyan Li (lihy@iastate.edu), ECpE, 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, Iowa State University, Ames, IA 50011 USA.

[1] AMES is an acronym for *A*gent-based *M*odeling of *E*lectricity *S*ystems. See Sun and Tesfatsion [5] for a detailed description of an earlier version of AMES (V1.3). Downloads, manuals, and tutorial information for all AMES version releases can be accessed at http://www.econ.iastate.edu/tesfatsi/AMESMarketHome.htm .

[2] The actual ratio of cleared price-sensitive demand to cleared fixed demand in the MISO is currently very small (about 1%).

[3] The MISO 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.

(0 to 100%); (b) the level of the supply-offer price cap (infinite, high, moderate, or low); and (c) the absence or presence of generator learning capabilities. Dynamic market performance is measured in terms of six average outcome variables: average LMP; average total demand; average market operational cost; average market power (as measured by the Lerner Index); average LMP spiking; and average LMP volatility range.[4]

Experimental findings for dynamic market performance are reported in Section IV. For example, with generator learning, starting from an all-fixed-demand benchmark, average LMP is shown to increase with small increases in price sensitivity before declining monotonically. Also, with generator learning, starting from a no-price-cap benchmark, it is shown that the imposition of a binding supply-offer price cap can increase average LMP spiking and volatility even though average LMP is reduced.

Concluding remarks are given in Section V. Technical definitions and calculations for average outcome variables are provided in Appendix A.

## II. Overview of the AMES Test Bed

As detailed in Li et al. [1], AMES (V2.0) incorporates in simplified form various core features of the WPMP market design as implemented in the MISO. A summary of these core features is as follows:

- The AMES wholesale power market operates over an AC transmission grid starting on day 1 and continuing through a user-specified maximum day (unless the simulation is 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 generators distributed across the nodes of the transmission grid.
- The ISO undertakes the daily operation of a day-ahead market settled by means of locational marginal prices (LMPs). The binding financial contracts determined in the day-ahead market are carried out as planned (no shocks to the system), hence traders have no need to engage in real-time (spot) market trading.
- 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 linear-affine inverse 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.
- During the morning of each day $D$, each generator reports one supply offer to the ISO to be used for all hours of the day-ahead market for day $D + 1$.[5] Each reported supply offer consists of a price-sensitive linear-affine marginal cost function defined over an operating capacity interval.
- After receipt of these demand bids and supply offers during the morning of day $D$, the ISO determines and publicly reports hourly power supply commitments and LMPs for the day-ahead market for day $D + 1$ as the solution to hourly bid/offer-based DC optimal power flow problems.
- At the end of each day $D$, the ISO settles all of the commitments for the day-ahead market for day $D+1$ on the basis of the LMPs for the day-ahead market for day $D + 1$.
- Each generator uses its day $D$ settlement payment to adjust, via reinforcement learning, its choice of a supply offer to be reported to the ISO on day $D + 1$ for the day-ahead market for day $D + 2$. Generators can adjust the ordinates/slopes of their reported marginal cost functions and/or the upper limits of their reported operating capacity intervals.
- Transmission grid congestion in the day-ahead market is managed via the inclusion of congestion components in LMPs.
- Each LSE and generator has an initial holding of money that changes over time as it accumulates profit earnings and losses.
- There is no entry of traders into, or exit of traders from, the AMES wholesale power market. LSEs and generators are currently allowed to go into debt (negative money holdings) without penalty or forced exit.

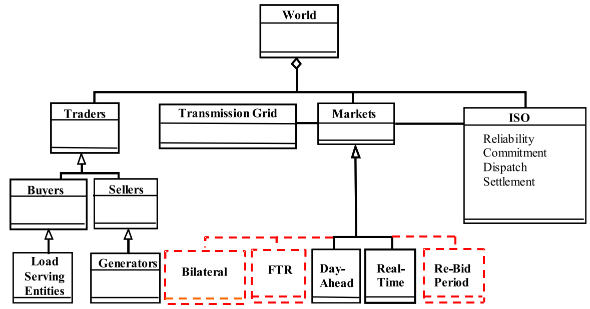


Fig. 1. AMES test bed architecture (agent hierarchy)

Figure 1 schematically depicts the current architecture of the AMES test bed by solid lines. Key elements planned but not yet incorporated are indicated by dashed lines.

As explained more carefully in Sun and Tesfatsion [5], the AMES ISO computes hourly LMPs and power commitments for the day-ahead market by solving bid/offer-based *DC Optimal Power Flow (OPF)* problems that approximate underlying

[4]For brevity, this conference study focuses on average effects. Detailed distributional effects are reported and examined in Li et al. [1].

[5]In the MISO [6], generators 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 generators, the current version of AMES restricts generators to the daily reporting of only one supply offer for the day-ahead market. Interestingly, the latter restriction is imposed on generators by the ISO-NE [7] in its particular implementation of the WPMP. Baldick and Hogan [8, pp. 18-20] conjecture that imposing such limits on the ability of generators to report distinct hourly supply offers could reduce their ability to exercise market power.

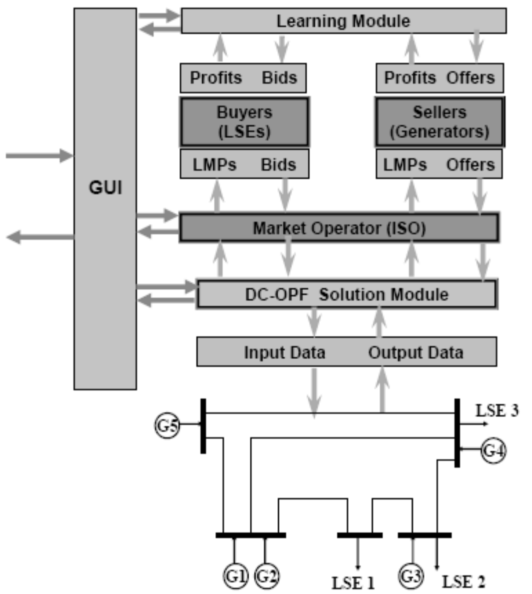


Fig. 2. Schematic depiction of the three core AMES test bed components: Learning module; DC-OPF solver, and Graphical User Interface (GUI)

AC-OPF problems. To handle this computation, we have developed an accurate and efficient DC-OPF solver, *DCOPFJ*, consisting of a strictly convex quadratic programming solver wrapped in an outer SI-pu data conversion shell (Sun and Tesfatsion [9]). The AMES ISO solves its DC-OPF problems by invoking DCOPFJ.

Generator learning is implemented in the AMES test bed by a reinforcement learning module, *JReLM*, developed by Gieseler [10]. JReLM can implement a variety of different reinforcement learning methods, permitting flexible representation of trader learning within this family of methods.

AMES also 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 attributes of the transmission grid; (c) initialization and editing of the attributes of LSEs and generators; (d) specification of the learning method for generators; (e) specification of simulation controls (e.g., the simulation stopping rule); and (f) customizable output reports in the form of both table and chart displays.

DCOPFJ, JReLM, and the GUI are the core components supporting the current implementation of the AMES test bed. This implementation is schematically depicted in Figure 2.

## III. Experimental Design

This section develops an experimental design to explore dynamic market performance under systematically varied settings for demand-bid price sensitivity, the supply-offer price cap, and generator learning capabilities. As the basic foundation for this experimental design, we consider the 5-node transmission grid configuration depicted in Figure 2.

Originally due to John Lally [11], this transmission grid configuration is now used extensively in ISO-NE/PJM training manuals to derive DC-OPF solutions at a given point in time conditional on variously specified generator attributes, LSE loads, and transmission grid reactances and line limits. An implicit assumption in these derivations is that the ISO knows the true attributes of the LSEs and generators. No mention is made of the possibility that LSEs and generators in real-world ISO-managed wholesale power markets might learn to exercise market power over time through strategic reporting of their attributes.

Our experimental design extends these static training cases both dynamically and strategically. The LSEs and generators repeatedly report demand bids and supply offers into the day-ahead market over time. Moreover, we examine what happens when generators are permitted to have learning capabilities enabling them to strategically adjust their reported supply offers on the basis of past profit earnings.

We start by considering the *dynamic 5-node benchmark case* presented in Table I. This benchmark case is characterized by 100% fixed demand (no price sensitivity), the absence of any supply-offer price cap, and the absence of any strategic learning on the part of generators (i.e., reported supply offers convey true cost and capacity attributes). The transmission grid configuration, reactances, locations of the generators and LSEs, and initial hour-0 load levels in Table I are taken from Lally [11]. The general shape of the LSE load profiles is adopted from a 3-node example presented in Shahidehpour et al. [12, p. 296-297].

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 $p^F_{Lj}(H)$ (in MWs) and a price-sensitive demand bid function

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

defined over a *purchase capacity interval*

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

In (1), the term $D_j(p^S_{Lj}(H))$ denotes LSE $j$'s *true reservation value* for $p^S_{Lj}(H)$, i.e., the maximum dollar amount it is truly willing to pay (per MWh) for the additional power $p^S_{Lj}(H)$ (in MWs). The parameter values $c_j(H)$ and $d_j(H)$ in (1) are required to be nonnegative.

Also, for each day $D$, the supply offer reported by generator $i$ for use in every hour of the day-ahead market for day $D+1$ consists of a reported marginal cost function

$$\mathrm{MC}^R_i(p_{Gi}(D)) \;=\; a^R_i(D) \;+\; 2b^R_i(D)\cdot p_{Gi}(D) \qquad (3)$$

defined over a reported *operating capacity interval*

$$0 \;\leq\; p_{Gi}(D) \;\leq\; Cap^{RU}_i(D)\;. \qquad (4)$$

In (3) the term $MC^R_i(p_{Gi}(D))$ denotes generator $i$'s *reported reservation value* for $p_{Gi}(D)$, i.e., the minimum dollar payment it reports it is willing to accept (per MWh) for the power supply $p_{Gi}(D)$ (in MWs). The parameter values $a^R_j(D)$ and $b^R_j(D)$ in (3) are required to be nonnegative.

Generator $i$ learns over time how to strategically report its daily supply offers based on the profit earnings it has obtained from its past supply offer choices. In particular, the parameter values ($a^R_i(D)$, $b^R_i(D)$, $\mathrm{Cap}^{RU}_i(D)$) that generator $i$ reports

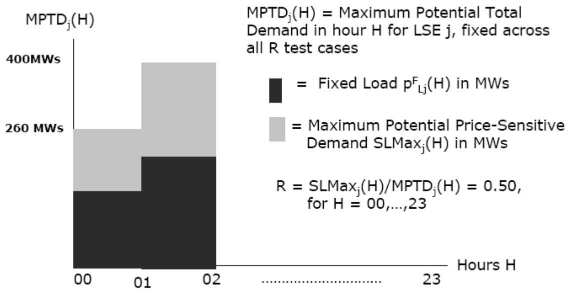


Fig. 3. Illustration of the construction of the R ratio for measuring relative demand-bid price sensitivity

on any given day D can deviate from its true supply offer parameter values ($a_i$, $b_i$, Cap$_i^U$).

In this study the primary 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

$$R_j(H) \ = \ \frac{SLMax_j(H)}{MPTD_j(H)} \tag{5}$$

where $SLMax_j(H)$ 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

$$MPTD_j(H) \ = \ [p^F_{Lj}(H) + SLMax_j(H)] \tag{6}$$

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 3.

We start with an experimental treatment in which all of the $R$-values in (5) are set equal to $R = 0.0$ for each LSE $j$ and each hour $H$ (the pure fixed-demand case). We then systematically increase $R$ by tenths, ending with the value $R = 1.0$ (the pure price-sensitive demand case). A positive R value indicates that the LSEs are able to exercise at least some degree of price resistance.

The maximum potential price-sensitive demands $SLMax_j(H)$ for each LSE are thus systematically increased across experiments. However, we control for confounding effects arising from changes in overall demand capacity as follows: For each LSE $j$ and each hour $H$, the denominator value $MPTD_j(H)$ in (6) is held constant across experiments by appropriate reductions in the fixed demand $p^F_{Lj}(H)$ as $SLMax_j(H)$ is increased. Specifically, $MPTJ_j(H)$ is set equal across all experiments to the hour-$H$ fixed-load level L-H reported for LSE $j$ in Table I.

Moreover, we also control for confounding effects arising from changes in demand-bid functional forms. The demand-bid ordinate and slope parameter values $\{(c_j(H), d_j(H)) : H = 00, \ldots, 23\}$ for each LSE $j$ are held fixed across experiments; see Li et al. [1] for the precise values used.

The second treatment factor explored in this study is PCap ($/MWh), an ISO-imposed supply-offer price cap. In experiments in which PCap is imposed, generators are not permitted to report marginal costs (reservation values) that rise above PCap. Consequently, the generators report supply offers such that their reported marginal costs at the upper limits $Cap^{RU}$ of their reported operating capacity intervals do not exceed PCap. The following four values are tested for PCap: (a) an effectively infinite value (no PCap); (b) a "high" value ($\$120$/MWh); (c) a "moderate" value ($\$100$/MWh); and (d) a "low" value ($\$80$/MWh).

The third treatment factor briefly explored in this study is generator learning capabilities. In each experiment we impose one of two treatments: (a) generators have no learning capabilities, hence they always report supply offers to the ISO that reflect their true cost and capacity conditions; or (b) generators have reinforcement learning capabilities that they use to strategically adjust their reported supply offers over time. For a careful description of the precise algorithmic representation used for generator learning, see Li et al. [1].

## IV. Report of Key Findings

This section uses the experimental design outlined in Section III to test the effects of changes in demand-bid price-sensitivity, supply-offer price caps, and generator learning on dynamic market performance as measured by the following six average outcome variables: Avg LMP; Avg Total Demand; Avg Op Cost (operational cost); Avg LI (Lerner Index); Avg LMP Spiking; and Avg LMP Volatility Range. The technical definitions and calculations of these average outcome variables are discussed in Appendix A.

Table II reports experimental findings for average outcomes under alternative settings for R (relative demand-bid price sensitivity) in the absence of a supply-offer price cap and with no generator learning. Table III repeats these experiments for the case in which generators have learning capabilities and hence learn to report strategic supply offers to the ISO over time.

TABLE II
Average effects of R changes with no supply-offer price cap (No PCap) and no generator learning

| R | Avg LMP | Avg Total Demand | Avg Op Cost | Avg LI |
|---|---|---|---|---|
| 0.0 | 25.18 | 318.21 | 3779.17 | 0.0056 |
| 0.1 | 24.51 | 299.19 | 3439.32 | 0.0042 |
| 0.2 | 23.92 | 279.69 | 3100.91 | 0.0036 |
| 0.3 | 23.33 | 259.85 | 2765.58 | 0.0032 |
| 0.4 | 22.72 | 240.18 | 2446.54 | 0.0029 |
| 0.5 | 22.10 | 220.88 | 2143.65 | 0.0026 |
| 0.6 | 21.35 | 204.09 | 1888.46 | 0.0022 |
| 0.7 | 20.49 | 188.67 | 1662.19 | 0.0013 |
| 0.8 | 19.49 | 175.74 | 1481.15 | 0.0000 |
| 0.9 | 18.27 | 169.68 | 1408.55 | 0.0000 |
| 1.0 | 17.04 | 163.87 | 1349.49 | 0.0000 |

As seen in Table II, in the absence of generator learning an incremental increase in R starting from the benchmark case R=0.0 (no price-sensitive demand) has the usual intuitively-expected effects: Avg LMP, Avg Total Demand, Avg Op Cost, and Avg LI all monotonically decline with increases in R.

Indeed, except for the presence of grid congestion between node 1 and node 2 and a binding operating-capacity constraint on generator 3 for the cases in which Avg Total Demand is relatively high, all of the Avg LI outcomes in Table II

TABLE III
AVERAGE EFFECTS (WITH STANDARD DEVIATIONS) OF R CHANGES WITH NO PCAP AND WITH GENERATOR LEARNING

| R | Avg LMP | Avg Total Demand | Avg Op Cost | Avg LI |
|---|---|---|---|---|
| 0.0 | 70.10<br>(3.14) | 318.21<br>(0.00) | 9198.63<br>(125.88) | 0.5692<br>(0.01) |
| 0.1 | 73.84<br>(3.24) | 286.39<br>(0.00) | 8450.26<br>(444.20) | 0.5755<br>(0.01) |
| 0.2 | 81.46<br>(2.85) | 254.57<br>(0.05) | 7629.94<br>(298.22) | 0.5933<br>(0.01) |
| 0.3 | 72.67<br>(3.02) | 223.84<br>(1.14) | 5501.09<br>(228.62) | 0.5433<br>(0.01) |
| 0.4 | 39.43<br>(1.16) | 198.70<br>(2.03) | 3300.37<br>(172.36) | 0.4341<br>(0.01) |
| 0.5 | 35.75<br>(0.48) | 170.75<br>(2.42) | 2717.73<br>(157.73) | 0.4185<br>(0.01) |
| 0.6 | 33.52<br>(0.41) | 155.47<br>(2.86) | 2259.65<br>(135.01) | 0.3660<br>(0.01) |
| 0.7 | 28.73<br>(0.60) | 145.84<br>(4.23) | 1877.91<br>(151.64) | 0.2815<br>(0.01) |
| 0.8 | 26.75<br>(0.54) | 133.99<br>(4.96) | 1627.45<br>(157.23) | 0.2547<br>(0.01) |
| 0.9 | 25.09<br>(0.51) | 120.17<br>(5.43) | 1388.31<br>(132.60) | 0.2342<br>(0.01) |
| 1.0 | 23.23<br>(0.48) | 108.51<br>(5.80) | 1184.18<br>(125.88) | 0.2078<br>(0.01) |

would be zero. Generators have no learning capabilities and are reporting their true cost and capacity conditions to the ISO each day, hence they are not making any deliberate efforts to exercise market power. Rather, as explained more carefully in Li et al. [1], the grid congestion is causing some separation of LMP values both from each other (cross-sectionally across the grid) and from generator marginal costs, and the binding operating-capacity constraint causes separation of generator marginal costs from each other. Both effects result in non-zero values for Avg LI.

Comparing the no-learning Table II results to the results with generator learning reported in Table III, it is seen that generator learning has strong effects on average outcomes. With generator learning, Avg LMP, Avg Op Cost, and Avg LI are all dramatically higher for every level of R even though Avg Total Demand is lower. The reason is that the profit-seeking generators quickly learn to tacitly collude on higher-than-true reported marginal costs even when demand bids are fully price sensitive (R=1.0) and the generators are competing for limited demand.

Moreover, with generator learning, Avg LMP and Avg LI exhibit a counterintuitive behavior: as R is incrementally increased from R=0.0 (no price-sensitive demand) to R=0.2 (some price-sensitive demand), both Avg LMP and Avg LI actually increase. These initial increases occur even though Avg Total Demand and Avg Op Cost are monotonically declining. From a policy standpoint, it is interesting to note that the current R ratio is about 0.01 for the MISO [6].

As explored more fully in Li et al. [1], these initial increases in Avg LMP and Avg LI appear to be robust phenomena that arise from complicated interactions between learning and network effects. The "critical R value" R* at which Avg LMP and Avg LI exhibit a turning point from increasing to decreasing depends on other maintained parameter value settings. For example, R* varies systematically with changes in the c-values (ordinates) of the LSEs' price-sensitive demand-bid functions.

TABLE IV
AVERAGE LMP EFFECTS (WITH STANDARD DEVIATIONS) OF CHANGES IN PCAP WITH NO DEMAND-BID PRICE SENSITIVITY (R=0.0)

| | No PCap | PCap=120 | PCap=100 | PCap=80 |
|---|---|---|---|---|
| Avg LMP with no gen learning | 25.18 | 25.18 | 25.18 | 25.18 |
| Avg LMP with gen learning | 70.10<br>(3.14) | 65.72<br>(4.01) | 58.00<br>(1.51) | 54.96<br>(2.41) |

Table IV reports Avg LMP outcomes under four alternative scenarios for PCap, the supply-offer price cap. For the subsequent interpretation of these findings, it is important to recall from Section III that PCap is a price cap on generator-reported marginal costs and *not* on LMPs per se. In the presence of congestive grid conditions, LMPs can separate from generator-reported marginal costs and hence from PCap.

None of the three numerical PCap values in Table IV is binding in the absence of generator learning, hence the Avg LMP outcome $25.18/MWh with no generator learning provides a common benchmark value.[6] However, with generator learning, each of these three PCap values results in an Avg LMP outcome that differs from the Avg LMP outcome with no price cap. This indicates that these three PCap levels are binding on generator-reported marginal costs. More precisely, a binding PCap level means that one or more generators have been forced to reduce the ordinate/slope values and/or the upper operating capacity limits of the supply offers they report to the ISO.

As intuitively expected, Avg LMP monotonically decreases as PCap is decreased in increments from an effectively infinite value (No Price Cap) to a low value ($80/MWh). Due to learning and network effects, however, the relationship between PCap and LMP outcomes is more complicated than indicated by this Avg LMP effect.

In particular, note in Table IV that Avg LMP with no price cap is $70.10/MWh whereas Avg LMP for PCap=$120/MWh is only $65.72./MWh. This finding indicates that the high PCap level $120/MWh is binding on the generators' reported marginal costs even though this PCap level is substantially higher than the resulting value $65.72/MWh for Avg LMP. A similar comment holds for the remaining two PCap levels.

The explanation for this finding is that the distribution of LMPs across the 24 hours of a day can exhibit substantial spiking and volatility that are obscured when only daily Avg LMP outcomes are considered. For example, as shown in Figure 4, the maximum LMP value attained during peak demand hours can be sustantially higher than Avg LMP calculated across all 24 hours. Thus, the imposition of a price cap can be a binding constraint on generator-reported marginal costs during peak demand hours even if not in other hours. Since generators are only permitted to report one supply offer per day, a binding constraint on reported marginal costs during peak demand

[6] As shown in Li et al. [1], in the no-learning case the price cap level PCap only becomes binding on generator-reported marginal costs when it drops below $35.40/MWh.

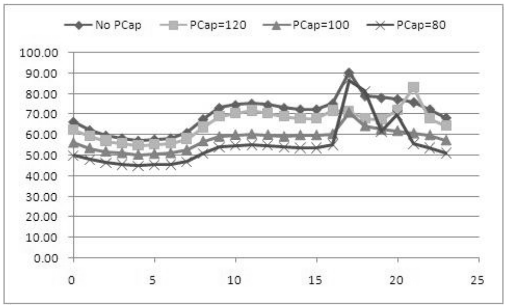


Fig. 4. Average hourly LMPs during the final market day under alternative supply-offer price caps with no demand-bid price sensitivity (R=0.0) and with generator learning

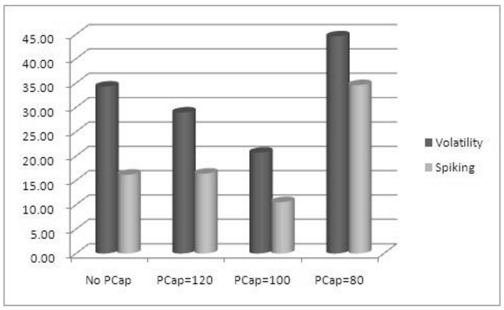


Fig. 5. Average LMP spiking and volatility range effects of supply-offer price caps with no demand-bid price sensitivity (R=0.0) and with generator learning

hours translates into a binding supply-offer constraint for every hour.

Moreover, as examined at much greater length in Li et al. [1], the introduction of a binding PCap level can in some cases worsen LMP spiking and volatility while in other cases LMP spiking and volatility are dampened. For example, as seen in Figure 5 for the tested scenario with no demand-bid price sensitivity (R=0.0) and with generator learning, the introduction of the strongly binding PCap level $80/MWh increases both spiking and volatility whereas the introduction of the more moderately binding PCap level $100/MWh has the opposite effect.

## V. Concluding Remarks

Restructured wholesale power markets are complex systems encompassing nonlinear physical constraints, complicated institutional arrangements, and the behavioral dispositions of potentially strategic human participants. To be compelling and useful, studies of such systems must take all three elements into proper account. As carefully discussed in [8], this has proved to be extremely difficult to accomplish using standard analytical and statistical tools.

In this study we demonstrate the potential usefulness of an agent-based test bed, AMES, for the exploratory study of restructured wholesale power markets operating through time over an AC transmission grid subject to line constraints, generator capacity constraints, and strategic trader behaviors. AMES was first released as open-source software at the IEEE PES General Meeting (June 2007) and is available at the website of the IEEE Task Force on Open-Source Software (http://ewh.ieee.org/cmte/psace/CAMS_taskforce/index.htm).

In particular, we conduct systematic experiments with AMES to determine the complex effects of changes in demand-bid price sensititivity, supply-offer price caps, and generator learning on wholesale power market performance – in particular, on dynamic LMP response — due to potentially congestive grid conditions and to potentially binding capacity constraints on power generation. A more detailed study of LMP spiking and volatility patterns in response to systematic changes in these treatment factors is provided in Li et al. [1].

## Appendix A
## Calculation of Average Effects

This appendix explains the definition and calculation of the average outcome effects reported in Section IV.

Thirty runs were conducted for each treatment factor configuration corresponding to 30 different random seeds. These seeds were generated using the standard Java "random" class. See Li et al. [1] for a listing of these seed values.

Each run terminates at a "final day" determined in accordance with the following stopping rule: Either end at day 100 or end at the earliest day for which each generator has converged to the choice of a single reported supply offer with probability at least 0.999. For each calculation below, only the final-day data for each run are used. Note, however, that this final-day data consists of scheduled hourly outcomes for the 24 hours of the day-ahead market on the subsequent day.

Also, for the reported findings in Table IV and Figs. 4 and 5 with an imposed PCap, with no demand-bid price sensitivity (R=0.0), and with generator learning, an occasional *inadequacy event* occurred in a small number of runs around the peak-demand hour 17 in that total generator reported capacity was insufficient to meet total fixed demand.[7] These few runs/hours are excluded from the averages calculated below, but for expositional simplicity we ignore this complication in the calculation descriptions.

Avg LMP ($/MWh) is calculated as follows. First, for each transmission grid node and each hour, determine the average hourly LMP across all 30 runs. Second, for each hour, determine the average of these run-averaged hourly LMP values across all five nodes. Finally, average these node-averaged and run-averaged hourly LMP values across all 24 hours to get Avg LMP.

Avg Total Demand (MWs) is calculated as follows. First, for each of the three LSEs and for each hour, determine the LSE's average cleared (satisfied) price-sensitive demand across all 30 runs. Second, for each LSE and each hour, 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 to get Avg Total Demand.

[7] As discussed in Li et al. [1], inadequacy events 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 (V2.0) reports a tick-count of inadequacy events. Future versions of AMES will incorporate adequacy protection procedures based on empirical ISO practices.

Avg Op Cost (ISO market operational cost, \$/h) is calculated as follows. First, for each of the five generators and for each hour, determine the average total variable cost of the generator across all 30 runs based on the generator's *reported* supply offer and subsequent hourly power commitment levels as determined by the ISO.[8] Second, for each hour, determine the average of these run-averaged total variable cost calculations across all five generators. Third, average these generator-averaged and run-averaged hourly total variable cost calculations across all 24 hours to get Avg Op Cost.

The Lerner Index for any generator $i$ supplying a positive amount of power $\mathrm{P}_{Gi}$ is defined as follows:

$$\mathrm{LI}_i(P_{Gi}) \;=\; \frac{\left[\mathrm{LMP}_{k(i)} - \mathrm{MC}_i(\mathrm{P}_{Gi})\right]}{\mathrm{LMP}_{k(i)}} \;. \tag{7}$$

In (7), $k(i)$ denotes the nodal location of generator $i$, $\mathrm{LMP}_{k(i)}$ denotes the LMP at node $k(i)$, and $\mathrm{MC}_i(P_{Gi})$ denotes generator $i$'s true marginal cost of supplying $\mathrm{P}_{Gi}$.

Avg LI (pure number) is calculated as follows. First, for each run, for each hour, and for each generator $i$ with a positive power commitment $\mathrm{P}_{Gi}$ for this run and hour, determine the generator's Lerner Index (7). Second, for each hour and for each generator, determine the average of this generator's Lerner Indices across all of the runs for which he had a positive power commitment for this hour. Third, for each hour, determine the average of these run-averaged Lerner Indices across all generators who were committed during this hour for at least one run. Finally, determine the average of these generator-averaged and run-averaged Lerner Indices across all 24 hours to get Avg LI.

The Lerner Index is a simple commonly-used measure of market power in electricity and other markets, but it has severe drawbacks as well. A more careful examination of market power in relation to demand-bid price sensitivity and supply-offer price caps is provided in Li et al. [1].

The Avg LMP Spiking (\$/MWh) depicted in Figure 5 is calculated as follows. First, for each run and for each of the five transmission grid nodes, calculate nodal LMP spiking as the maximum absolute difference between successive hourly LMPs across all 24 hours. Second, for each node, determine the average of these nodal LMP spiking measures across all 30 runs. Third, determine the average of these run-averaged nodal LMP spiking measures across all five nodes to get Avg LMP Spiking.

The Avg LMP Volatility Range (\$/MWh) depicted in Figure 5 is calculated as follows. First, for each run and for each of the five transmission grid nodes, calculate the nodal LMP volatility range as [maxLMP-minLMP] across all 24 hours. Second, for each node, determine the average of these nodal LMP volatility range measures across all 30 runs. Third, determine the average of these run-averaged nodal LMP volatility range measures across all five nodes to get the Avg LMP Volatility Range.

[8]That is, these average total variable cost calculations are made by the ISO using the marginal cost functions *reported* by generators 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 power commitment levels. The ISO does not know the generators' true marginal cost functions.

## REFERENCES


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

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

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

[4] 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.

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

[6] MISO. (2008) Home page. Midwest ISO, Inc. [Online]. Available: www.midwestiso.org/

[7] ISO-NE. (2008) Home page. ISO New England, Inc. [Online]. Available: www.iso-ne.com/

[8] R. Baldick and W. Hogan,"Capacity constrained supply function equilibrium models of electricity markets: Stability, non-decreasing constraints, and function space iterations," Working Paper Series, Program on Workable Energy Regulation (POWER), University of California Energy Institute, Revised August 2002.

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

[10] 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.

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

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



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

TABLE I

INPUT DATA FOR THE DYNAMIC 5-NODE BENCHMARK CASE: NO DEMAND-BID PRICE SENSITIVITY (R=0.0); NO SUPPLY-OFFER PRICE CAP (NO PCAP); AND NO GENERATOR LEARNING

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Base Values[a] | | | | | | | | | |
| $S_o$ | $V_o$ | | | | | | | | |
| 100 | 10 | | | | | | | | |
| $K^b$ | $\pi^c$ | | | | | | | | |
| 5 | 0.05 | | | | | | | | |
| Branch From | To | lineCap[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 | | | | | | |
| Gen ID | atNode | FCost | $a$ | $b$ | $\text{Cap}^L$ | $\text{Cap}^U$ | Init$ | | |
| 1 | 1 | 1600.0 | 14.0 | 0.005 | 0.0 | 110.0 | $1M$ | | |
| 2 | 1 | 1200.0 | 15.0 | 0.006 | 0.0 | 100.0 | $1M$ | | |
| 3 | 3 | 8500.0 | 25.0 | 0.010 | 0.0 | 520.0 | $1M$ | | |
| 4 | 4 | 1000.0 | 30.0 | 0.012 | 0.0 | 200.0 | $1M$ | | |
| 5 | 5 | 5400.0 | 10.0 | 0.007 | 0.0 | 600.0 | $1M$ | | |
| LSE ID | atNode | L-00[f] | L-01 | L-02 | L-03 | L-04 | L-05 | L-06 | L-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 |
| ID | atNode | L-08 | L-09 | L-10 | L-11 | L-12 | L-13 | L-14 | L-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 |
| ID | atNode | L-16 | L-17 | L-18 | L-19 | L-20 | L-21 | L-22 | L-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 nodes
[c]Soft penalty weight $\pi$ for voltage angle differences
[d]Upper limit $P_{km}^U$ (in MWs) on the magnitude of real power flow in branch $km$
[e]Reactance $X_{km}$ (in ohms) for branch $km$
[f]L-H: Load (in MWs) for hour H, where H=00,01,...,23