2009

Dynamic performance of restructured wholesale power markets with learning generation companies: an agent-based test bed study

Hongyan Li
Iowa State University

Dynamic performance of restructured wholesale power markets with learning generation companies: an agent-based test bed study

by

Hongyan Li

A thesis submitted to the graduate faculty

in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Major: Electric Engineering

Program of Study Committee:
Leigh Tesfatsion, Co-major Professor
James D. McCalley, Co-major Professor
Venkataramana Ajjarapu
David Hennessy
Lizhi Wang

Iowa State University

Ames, Iowa

2009

DEDICATION

I would like to dedicate this thesis to my wife Yunhui and to my daughter Angelina without whose support I would not have been able to complete this work.

I would also like to thank my friends and family for their loving guidance and financial assistance during the PhD study.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGEMENTS

I would like to take this opportunity to express my thanks to those who helped me with various aspects of conducting research and the writing of this thesis.

First and foremost, Dr. Leigh Tesfatsion for her guidance, patience and support throughout this research and the writing of this thesis. Her insights and words of encouragement have often inspired me and renewed my hopes for completing my graduate education.

Next I want to thank Dr. James D. McCalley for his guidance throughout my graduate career.

I would also like to thank my committee members for their efforts and contributions to this work: Dr. Venkataramana Ajjarapu, Dr. David Hennessy and Dr. Lizhi Wang.

ABSTRACT

In April 2003, the U.S. Federal Energy Regulatory Commission (FERC) proposed a new market design for U.S. wholesale power markets. Core features of this design include oversight of operations by some form of Independent System Operator (ISO), a two-settlement system consisting of a day-ahead market supported by a parallel real-time market to ensure continual balancing of supply and demand for power, and management of grid congestion by means of locational marginal pricing. Seven U.S. energy regions are now operating under a variant of FERC's market design. This dissertation undertakes the systematic study of core features of FERC's market design by means of intensive simulation experiments. Specific studied issues include: the effects of generator learning behaviors on market efficiency and supply adequacy; the effects of changes in generator learning parameters, demand-bid price sensitivities, and generator supply-offer price caps on locational marginal price separation and volatility over time; market efficiency implications of ISO net surplus (congestion rent) collections and redistributions; and the effects of generator economic and physical capacity withholding on generator net earnings and market efficiency. To carry out this research, major extensions of the AMES wholesale power market test bed have been developed. To encourage the accumulation of further research findings, these extended versions of AMES have been released as open-source software.

CHAPTER 1. OVERVIEW

1.1 Research Motivation and Background

In April 2003, the U.S. Federal Energy Regulatory Commission (FERC) proposed a new market design for U.S. wholesale power markets. This new market design contains the following core features: central oversight by some form of *Independent System Operator (ISO)*; a two-settlement system consisting of a day-ahead market supported by a parallel real-time market to ensure continual balancing of supply and demand for power; and management of grid congestion by means of locational marginal pricing (LMP), i.e., the pricing of power by the location and timing of its injection into, or withdrawal from, the transmission grid. Versions of FERC's market design have now been implemented (or adopted 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).

Wholesale power markets have special characteristics that distinguish them from other commodity markets, as follows: (1) electricity is difficult to store in large quantities; (2) from an operational point of view, electricity requires demand to be always met by supply; and (3) transmission lines can become congested due to transmission limits. In addition, generation units have operating capacity limits that must be respected.

Prior to restructuring, generation, transmission and distribution were handled jointly by vertically integrated utility companies. Consequently, these wholesale power market characteristics were relatively easy to manage. As part of restructuring, however, an effort was made to divest generation from the management of transmission in order to introduce a greater role for market forces in generation, hence new ways of managing these characteristics had to be found. In each restructured energy region this has resulted in a complex set of market protocols and regulations that are conditioned on physical constraints and the anticipated behaviors of participant traders. Even today these protocols and regulations are still being modified and revised, accompanied by new offerings of market products and services. Clearly a great deal of testing is needed to understand and evaluate the extent to which these protocols and regulations ensure efficient and reliable system operations.

The general goal of my dissertation research is the systematic study and evaluation of U.S. restructured wholesale power markets. More precisely, the following objectives have been pursued:

- (1) To study the potential learning behaviors of *Generation Companies (GenCos)* participating in U.S. restructured wholesale power markets.

- (2) To investigate whether the complicated rules and regulations governing market operations encourage GenCos to engage in strategic supply offer selection that reduces overall system performance measured in terms of market efficiency (i.e., non-wastage of resources) and supply adequacy (i.e., sufficiency of offered supply to meet demand).

- (3) To investigate 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 GenCos.

- (4) To explore the market efficiency implications of the net surplus (congestion rents) collected and redistributed by ISOs in restructured wholesale power markets with grid congestion managed by LMP.

- (5) To investigate the market power and market efficiency implications of strategic *economic* and *physical* capacity withholding by GenCos.

Given the complicated nature of restructured wholesale power markets, I have chosen simulation as my main investigative tool. Several companies provide wholesale power markets simulation software packages, but it is not possible to fully access this proprietary commercial software for study and research purposes. In addition, various academic researchers have used power market simulations, but it is difficult to duplicate their test cases and verify their results because their code has not been publicly released.

Consequently, in order to carry out my dissertation research, I have developed appropriately extended versions of AMES, an agent-based test bed that captures in simplified form some of the core features of U.S. restructured wholesale power markets. These versions have been released as open source software to facilitate understanding and replication of my results as well as to encourage the accumulation of further research findings.

1.2 Organization of the Dissertation

The remainder of this dissertation is organized as follows. Chapter 2 provides a review of existing related literature. Specifically, this review covers: (1) restructured wholesale power markets background; (2) history and usage of agent-based modeling, which is the main modeling method used in this dissertation; (3) learning methods introduction; (4) review of other agent-based electricity market simulation research; and (5) additional motivation for my decision to develop special open-source software for my dissertation research.

Chapter 3 describes the main components of the AMES wholesale power market test bed as developed for this dissertation. AMES includes three types of decision-making agents: an ISO; a collection of generation companies (GenCos); and a collection of load-serving entities (LSEs). The ISO manages two types of markets - a day-ahead market and a real-time market - operating over an AC transmission grid. The GenCos are learning agents that can learn over time to report strategic supply offers to the ISO for the day-ahead market. The LSEs submit demand bids to the ISO for the day-ahead market that consist of both fixed and price-sensitive parts. This chapter describes these different components and their interactions in relation to real-world restructured wholesale power markets.

Chapter 4 uses the AMES wholesale power market test bed introduced in Chapter 3 to systematically investigate the effects of changes in GenCo learning parameters, demand-bid price sensitivities, and GenCo supply-offer price caps on LMP separation and volatility over time. 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).

Chapter 5 explores the market efficiency implications of the net surplus (congestion rents) collected and redistributed by ISOs in restructured wholesale power markets with grid congestion managed by LMP. The AMES simulation findings suggest that these ISO net surplus collections can be substantial and tend to increase in conditions unfavorable to market efficiency. ISO yearly reports for PJM and other energy regions indicate that actual ISO net surplus collections are in fact substantial. A practical implication is that a more transparent public oversight of all net surplus collections and uses in wholesale power markets operating under LMP would be publicly prudent because these collections are not structurally well-aligned with market efficiency objectives.

Chapter 6 investigates strategic capacity withholding by GenCos in restructured wholesale power markets. Real-world restructured wholesale power markets are sequential open-ended games for GenCos, so they have a learning-to-learn issue. Consequently, as a preliminary step, the GenCos' learning methods are calibrated to their decision environment. Experiments are then conducted to investigate GenCo economic and physical capacity withholding both separately and in combination. Results indicate that economic capacity strongly dominates physical capacity withholding in terms of permitting GenCos to substantially increase their net earnings.

Chapter 7 provides concluding remarks as well as planned future work. Appendix A gives more details about the AMES agent-based wholesale power market test bed. Appendix B and C give complete input data for 5-bus and 30-bus test cases used in various parts of this dissertation study.

CHAPTER 2. LITERATURE REVIEW

2.1 Restructured Wholesale Power Markets Background

In the 1990s before power industry restructuring in the U.S., electricity retail rates were regulated by state commission, and utility companies were vertically-integrated, which means they owned generation, transmission, and the distribution network. The Energy Policy Act of 1992 had two main effects: (1) it encouraged power market competition, and (2) it gave FERC the authority to grant access to existing transmission lines.

In order to lower costs by improving competition and efficiency, FERC issued Orders 888 and 889 in 1996. These orders provided guidance regarding the formation of ISOs, required ISOs to provide unbiased access to transmission, permitted to charge a rate set by FERC, and required posting data to Open Access Same-Time Information System (OASIS). Order 888 and 889 are the foundations for creating competitive wholesale power markets.

FERC issued Order 2000 in 1999. This order required each public utility that owns, operates, or controls facilities for the transmission of electric energy in interstate commerce to make certain filings with respect to forming and participating in a Regional Transmission Organization (RTO). It also encouraged transmission owners to join an RTO to end any discrimination by transmission owners, and it established four characteristics (independence, regional configuration, operational authority, short-term reliability) and eight key functions (tariff administration, congestion management, parallel path flows, ancillary services, OASIS and capability calculations, market monitoring, planning and expansion, interregional coordination) of RTOs.

The wholesale power market design proposed by FERC FERC (2003) contains the following core features: central oversight by an independent system operator (ISO); a two-settlement system consisting of a day-ahead market supported by a parallel real-time market to ensure continual balancing of supply and demand for power; and management of grid congestion by means of locational marginal pricing (LMP), i.e., the pricing of power by the location and timing of its injection into, or withdrawal from, the transmission grid.

Different versions of FERC's market design have 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).

These restructured wholesale power markets are extremely complicated, involving (1) physical constraints; (2) institutional arrangements; (3) behavioral dispositions of human participants. The complexity of FERCs market design, together with the relative recency of its adoption in many regions of the U.S. (implying a short data series), make it extremely difficult to undertake adequate efficiency and reliability studies using standard analytical and statistical modeling tools.

One key problem for participants in restructured wholesale power markets operating under FERC's design is a lack of full transparency regarding market operations. Due in great part to the complexity of the market design in its various implementations, business practices manuals and other public documents provided by ISOs are daunting to read and difficult to comprehend. Moreover, in many ISO websites (e.g., MISO), data is only posted in partial and masked form with a significant time delay. The result is that many participants are wary regarding market efficiency, market reliability, and market fairness (e.g., settlement practices and market power mitigation rules). Moreover, it is more difficult for outsiders (e.g., university researchers) to study and test the design systematically, completely and with an open and unbiased opinion.

2.2 Agent-Based Computational Modeling

However, powerful new agent-based modeling tools have been developed to analyze this degree of complexity and have fruitfully been applied to the study of complex economic systems, see Tesfatsion, L. and Judd, K. L. (2006). In particular, some researchers have demonstrated the potential of agent-based modeling tools for the study of power systems.

Tesfatsion, L. (2009b) defines that ACE is the computational study of economic processes modeled as dynamic systems of interacting agents. Here "agent" refers broadly to a bundle of data and behavioral methods representing an entity constituting part of a computationally constructed world. Usually, ACE tools have computational framework, graphical user interface and modular capabilities for expansion.

Current ACE research can be divided into four categories:

(1) Empirical understanding, to study why particular observed regularities have evolved and persisted despite the absence of top-down planning and control. One examples of such regularities is market protocols. The aim of empirical understanding is try to understand whether particular types of observed regularities can be reliably generated from particular types of agent-based worlds.

(2) Normative understanding, to study computational laboratories for the discovery of good economic designs using ACE models. One example is market design. The main aim is to evaluate whether designs proposed for economic policies, institutions, or processes will result in socially desirable system performance over time. The key issue is does a proposed or actual market design ensure efficient, fair, and orderly market outcomes over time despite repeated attempts by traders to game the design for their own personal advantage?

(3) Qualitative insight and theory generation, using ACE models to gain a better understanding of economic systems through a better understanding of their full range of potential behaviors over time. Such understanding would help to clarify not only why certain types of regularities have evolved and persisted but also why others have not.

(4) Methodological advancement, to provide best ACE methods and tools to undertake theoretical studies of economic systems through systematic computational experiments, and to examine the compatibility of experimentally-generated theories with real-world data. A variety of ways ranging from careful consideration of methodological principles to the practical development of programming, visualization, and validation tools are explored by researchers.

In ACE, the core part is how to model different agents with different attributes. Currently in software programming implementation, agents are coded using object-oriented programming technique. These agent objects have their private data and methods, having communications with other agent objects. Here are some featured capabilities of agents: (1) adaptation to environmental conditions, (2) social communication with other agents, (3) goal-directed learning abilities, and (4) autonomy (self-activation and self-determinism based on private internal processes).

The typical experiment control process for ACE has the following steps: (1) researcher builds a virtual world to simulate the target environment, (2) researcher creates different types of agents which have the desired behaviors to simulate target institutions, organizations or human, (3) researcher sets initial conditions for the virtual world and agents, (4) researcher then has no interventions for the experiment environment and only observe how agents and the virtual world evolve over time.

In order to simulate realistically, the virtual world should reflect the main aspects of the studied system. In particular, the agents should reflect the main aspects of the institutions, physical features, and behaviors of the system. The most difficult part of ACE is to model both the virtual world and agents in a proper way, grasping the essential parts for study.

One of the most active research fields of ACE is normative understanding, that is, how to use ACE models as computational tools to propose or verify good economic designs. This is especially useful when new designs are proposed for economic policies, institutions, or processes. ACE tools can be used to evaluate if the proposed designs can achieve desired goals when agents have strategic behaviors and interact with each other and the virtual world over time.

There are at least two potential advantages of ACE for dynamic market modeling: (1) permits systematic experimental study of empirical regularities, economic institutions, and dynamic behaviors of complex market processes. (2) facilitates creative experimentation with realistically modeled market processes by using ACE test beds, researchers can evaluate interesting conjectures of their own devising, with immediate feedback and no original programming required, ACE software permits relatively easy modification/extension of features.

As will be demonstrated below, it is feasible to develop a useful agent-based tool to study FERC's market design for restructured wholesale power markets.

2.3 Learning Methods

A key aspect of learning for agents is the amount of anticipation (look-ahead) that agents employ. The general learning process can be expressed as: at beginning, an agent is at a state in a general environment. When a stimulus occurs, the agent reacts to this stimulus by choosing a particular action (response). The agent then observes an outcome, and it uses this outcome to either weaken or strengthen the association between the state and the action in the future.

There are three types of learning: (1) unsupervised learning updates structure based on agent intrinsic motivation (such as curiosity, enjoyment, moral duty); (2) reinforcement learning (RL) updates structure in response to successive rewards attained through actions taken; (3) supervised learning updates structure on basis of examples of desired (or required) state-action associations provided by an expert external supervisor.

Below is short introduction of different learning methods:

(1) Reinforcement Learning (RL). The basic intuition underlying reinforcement learning is that the tendency to implement an action should be strengthened (reinforced) if it results in favorable outcomes and weakened if it results in unfavorable outcomes, see Sutton, R. S. and Barto, A. G. (2000). RL is a relatively straight-forward type of learning method in which an agent constructs associations between states and actions. If the outcome is relatively good, the probability of future choice of this action is increased; if the outcome is relatively bad, the probability of future choice of this action is decreased.

RL choice problems can be divided into non sequential and sequential choice problems. For non sequential choice problems, an agent must learn a mapping from states to actions that maximizes expected immediate reward. In sequential choice problems, the agent must again learn a mapping from states to actions, but now the actions selected by the agent may influence future situations and hence future rewards as well as immediate rewards. Consequently, it might be advantageous for the agent to engage in anticipatory evaluation of the future possible consequences of its current actions.

(2) Stochastic Reactive RL Roth-Erev Algorithms. This method is developed by Roth and Erev, see Roth, A. E. and Ido, E. (1995) and Erev, I. and Roth, A. E. (1998). The form of this method is based on observations of people's behavior in iterated game play with multiple strategically interacting players in various game contexts. There are two extensions found necessary relative to RL methods developed earlier by psychologists for individuals learning in fixed environments: need to "forget" rewards received in distant past and need for "spillover" of reward attributions across actions in early game play to encourage experimentation and avoid premature fixation on a suboptimal chosen action.

(3) Q-Learning. Q-learning is a RL method developed by Watkins, see Watkins, C. (1989). Q-learning does not need a model of the environment and can be used on-line in contexts where multiple

agents are engaging in repeated non-zero sum games against unknown rivals and choosing their actions in an anticipatory way. Q-learning works by estimating the values of state-action pairs. The Q-value Q(s, a) is defined to be the expected discounted sum of future returns obtained by taking action a starting from state s and following an optimal action decision rule thereafter. Once these values have been learned, the optimal action from any state is the one with the highest Q-value.

(4) Genetic Algorithms (GA). GA uses directed search algorithm based on the mechanics of biological evolution. It was first developed by John Holland in the 1960s, and the GA remains one of the most prominent types of methods used in evolutionary computation. A GA is an abstraction of biological evolution. It is a method for evolving a new population of entities from an existing population of entities, with evolution biased in favor of more fit entities . The evolution proceeds via genetic operations (recombination, mutation,...) that act directly upon the structural characteristics of the population members.

(5) Artificial Neural Networks (ANNs). The ANNS method is inspirited from neurobiology. It has a collection of interconnected processing units working together. The structure contains unit configuration (numbers of input units, hidden units, and output units); unit connections and connection weights; and the structure can be updated using unsupervised learning, RL, or supervised learning. For example, ANNs by back propagation has such steps: first training examples are used to get desired input-output associations. During this process, error equals difference between desired and actual output for any given input, and connection weights is updated relative to error size. In summary, ANNs by back propagation starts by calculating output layer error and weight correction, then "propagate back" through previous layers.

For its easily implementation and understanding, stochastic reactive RL Roth-Erev algorithms and Q-learning is most widely used for agent-based electricity market simulation.

2.4 Agent-based Electricity Market Simulation

There are different traditional approaches in electricity modeling, such as competitive equilibrium models, Nash equilibrium models, supply function equilibrium models and experimental approaches. Here is more about competitive equilibrium models: the assumption is no player attempts to game the market, all generators report their marginal generating costs as supply offers. The results is the traditional market performance benchmark. But equilibrium models have drawbacks: normally strategic bidding behavior are not considered; assume that players have enough information such as other players' characteristics and behavior. Also in reality, it is hard to get equilibrium results due to environment changing everyday and hard to verify results.

To study the complexity of electricity markets, different modeling techniques are needed to understand market dynamics and to observe results for the appropriate design features. Comparing to traditional analytical methods widely used in power system, Agent-Based Computational Economics (ACE) is a fairly young research field that offers methods for realistic electricity market modeling. In the past few years, more and more researchers have developed different kinds of agent-based models to simulate electricity markets. In Weidlich, A. and Veit, D. (2008), Sensfuß, F. et al. (2007), and Zhou, Z. , Chan, W. and Chow, J. (2007), surveys about agent-based simulation of wholesale electricity markets give a lot insights of start-of-art ACE application.

In Bower, J. and Bunn, D.W. (2000), an ACE simulation model of the England and Wales electricity market is used to compare different market mechanisms, i.e. daily versus hourly bidding and uniform versus discriminatory pricing. This is the first ACE simulation for electricity markets.

In Zhou, Z. , Chan, W. and Chow, J. (2007), there are four agent-based electricity market simulation tools are mentioned:

(1) Simulator for electric power industry agents (SEPIA), is developed by Honeywell Technology Center and the University of Minnesota. Generation companies, consumers and transmission operator are modeled as agents. Generation company agents can use both a Q-learning module and a genetic classifier learning module to make decisions. As one of the earliest agent-based simulator for power markets, it is a good example. And it contains two learning methods. But it lacks ISO agent, and has some practical limitations of model, both restrict its usage.

(2) Electricity market complex adaptive systems (EMCAS), see Conzelmann, G. et al. (2004) and EMCAS (2009), is developed by center for energy, environmental and economic systems analysis at the Argonne national lab. Generation companies (GenCos), transmission companies operators (ISOs) or regional transmission organizations (RTOs), demand companies (DemCos), consumers, and regula-

tors are modeled as agents. An EMCAS simulation runs over six decision levels, ranging from hourly dispatching to long-term planning. At each decision level, agents make certain decisions, including determining electricity consumption (customer agents), unit commitment (generation companies), bilateral contracting (generation and demand companies), and unit dispatch (ISO/RTO agent). EMCAS is used by regulatory institutions interested in market design and consumer impact issues, transmission companies and market operators interested in system and market performance, and generation companies for strategic company issues.

(3) Short-term electricity market simulator-real time (STEMS-RT), is developed by Electric Power Research Institute (EPRI). STEMS-RT has features: agents rely on mathematical programming for bidding decisions and the latest techniques and strategies for bidding and realistic market rules can be added and their effects can be tested.

(4) National electricity market simulation system (NEMSIM), is developed particularly for the Australia National Electricity Market (NEM). Interestingly, the latest NEM model is a modified and extended version of the Agent-Based Modeling of Electricity System (AMES) model.

In later Chapters of this thesis, AMES model will be introduced and related research work will be demonstrated separately with details. In Zhou, Z. , Chan, W. and Chow, J. (2007), AMES is the only one open-source software for agent-based power market simulation. And even today, AMES is still the only one open-source agent-based power market simulation package by Google search. AMES got attention from other researchers. For example in Europe, London Metropolitan Business School proposed a ACEGES project - Agent-based Computational Economics of the Global Energy System. Here is a quote from its website: "The ACEGES laboratory, which will be open-source to encourage and stimulate its cumulative development over time. This is consistent with complementary research initiatives in USA such as the 'AMES Wholesale Power Market Test Bed' developed at Iowa State University".

2.5 Open-Source Software

Open-Source Software (OSS) expresses the idea that developers can publish their software with an open-source license, enabling anyone to use and modify this software. Today, OSS is widely used in the software industry, such as for language development tools (e.g., NetBeans for Java), office document processors (e.g., OpenOffice), and operating systems (e.g., Linux, OpenSolaris).

The OSS idea is especially useful for researchers. By studying current OSS work of peers, new researchers can catch up easily and focus on verifying the latest new ideas. This process is much faster than traditional research process.

The advantage of OSS is that persons interested in the software can read the detailed code, understand the related design, and modify the code to suit their special needs. People do not have to "reinvent the wheel" each time they need a specific code design for a specific problem. Instead, they can start from a previously established foundation subjected to open peer review, and they can do so without having to pay anything.

The exponential growth of OSS from 1993 through early 2008 is documented by Deshpande, A. and Riehle, D. (2008); see Fig. 2.1. OSS is now widely used in the software industry, such as for language development tools (e.g., NetBeans for Java), office document processors (e.g., OpenOffice), and operating systems (e.g., Linux, OpenSolaris). Increasingly OSS is being written as commercial-grade software, which could represent a significant change in the traditional proprietary approach to software development.

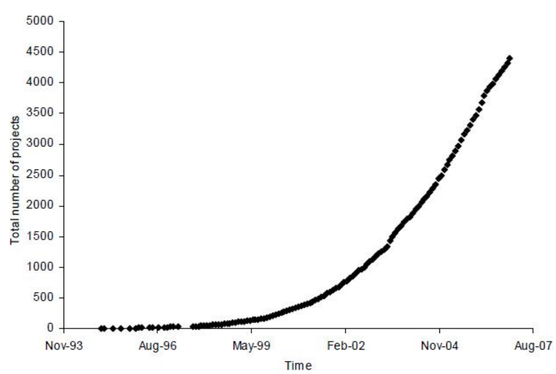

Figure 2.1 OSS Project Growth from November 1993 through August 2007. Source: Deshpande and Riehle (Deshpande, A. and Riehle, D. , 2008, Fig.1)

In an April 2003 white paper FERC (2003), the U.S. Federal Energy Regulatory Commission (FERC) proposed a new market design for U.S. wholesale power markets. Over 50% of U.S. generating capacity is now operating within the footprint of a wholesale power market restructured in compliance with the basic provisions of FERC's design.

These restructured wholesale power markets are complex, involving physical constraints, complicated market protocols, and behavioral dispositions of human participants. Moreover, time series are short due to the relative recency of the restructuring efforts, and the data that are available are often released only with a delay and only in partially masked form. Consequently, it is difficult to model and study these markets using standard analytical and statistical tools.

An additional complicating factor is that many economists are not familiar with transmission grid aspects of power systems, so they often focus on highly simplified two-bus or three-bus systems. Conversely, many power engineers are not familiar with basic economic market concepts, let alone the complicated design of restructured wholesale power markets. Modeling efforts by interdisciplinary teams capable of addressing both engineering and economic concerns would therefore be highly desirable. Until recently, there has not been OSS for restructured wholesale power markets.

In response to these concerns, an interdisciplinary group of researchers at Iowa State University has undertaken the OSS development of a wholesale power market test bed, referred to as AMES (*A*gent-based *M*odeling of *E*lectricity *S*ystems). The AMES test bed permits the systematic experimental study of strategic trading behaviors within restructured wholesale power markets operating over realistically rendered AC transmission grids. In addition, AMES facilitates augmentation of empirical input data with simulated input data to permit the study of a broader array of scenarios.

From the beginning, AMES was designed for research and teaching purposes rather than for commercial-grade application. AMES is entirely developed in the widely used Java programming language in order to facilitate readability and use. AMES is entirely OSS, combining together a collection of basic OSS modules for learning representation, optimal power flow solution, graphic display, and other functions. The modular and extensible OSS architecture of AMES permits users to modify and extend the code with relative ease to suit their special needs.

The first version of AMES was released as OSS at the IEEE Power and Energy Society General Meeting (PES GM) in 2007, and a substantially expanded second version was released as OSS at the IEEE PES GM in 2008. Downloads, manuals, and tutorial information for all AMES version releases to date are accessible at the AMES homepage, see Tesfatsion, L. (2009d). AMES is also available for downloading at the software site of the IEEE Taskforce on Open Source Software for Power Systems; see IEEE OSS (2009).

The release of AMES as OSS is intended to encourage the cumulative development of this test bed by multiple researchers in directions appropriate for their specific needs. It is also intended to encourage continual dialog with market stakeholders and regulators leading to successive refinements and improvements of the test bed.

CHAPTER 3. AGENT-BASED WHOLESALE POWER MARKET TEST BED CONSTRUCTION

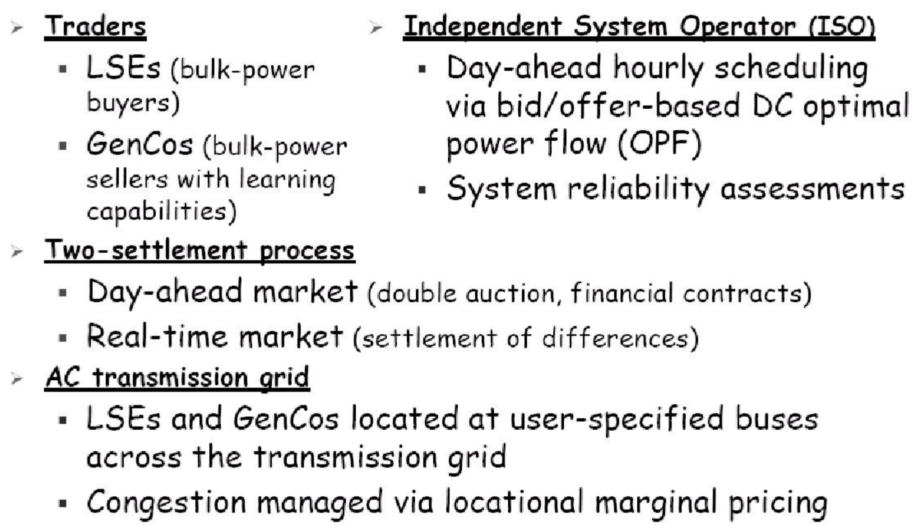

Figure 3.1 AMES test bed architecture.

The AMES Wholesale Power Market Test Bed incorporates in simplified form the core features of the wholesale power market design proposed by the U.S. FERC (2003); see Figure 3.1.

The simulated 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 simulated 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.

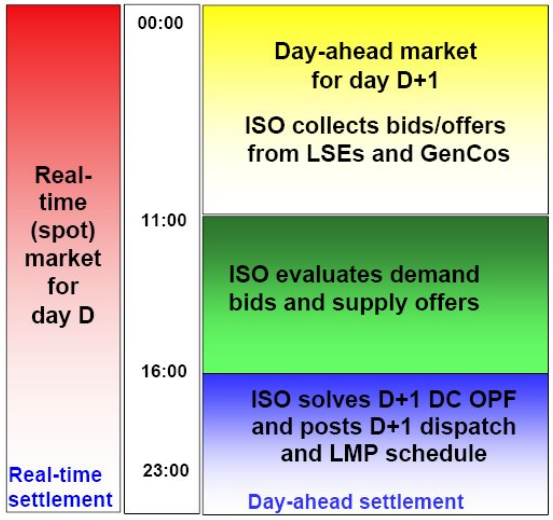

Figure 3.2 AMES ISO activities during a typical day D.

3.1 Independent System Operator Agent Model

3.1.1 Introduction

The ISO undertakes the daily operation of the transmission grid within a two-settlement system consisting of a Real-Time Market and a Day-Ahead Market. Each market is separately 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.[1] 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.

3.1.2 ISO Activities

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

[1]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, H. et al. (2009) for a careful discussion of LMP derivation from optimal power flow solutions.

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 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.[2] 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.

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.

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.

The activities of the ISO on a typical day D are depicted in Fig. 3.2. Fig. 3.3 provides a schematic depiction of simulated day-ahead market activities during a typical day D.

[2] In the MISO (2009), 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 (2009) in its particular implementation of FERC's market design. Baldick and Hogan (Baldick, R. and Hogan, W. , 2002, 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.

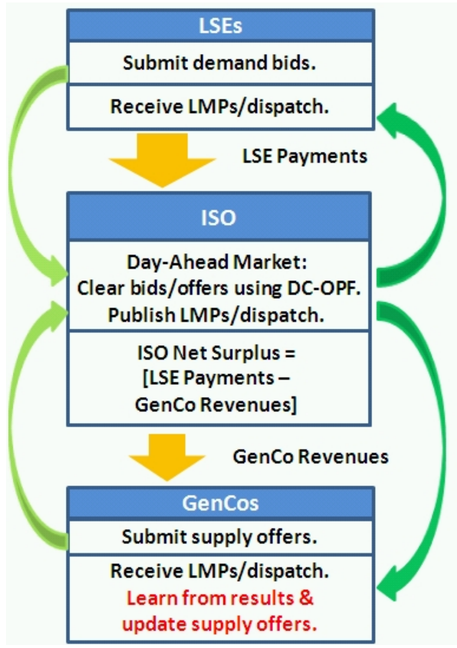

Figure 3.3 AMES day-ahead market activities during a typical day D.

3.1.3 ISO Net Surplus

The *ISO net surplus* extracted by the ISO during any given day D is the difference between LSE payments and GenCo revenues for the day-ahead market for day D+1:

$$\text{ISONetSur}(D) = \sum_{j=1}^{J} \text{Pay}_j(D) - \sum_{i=1}^{I} \text{Rev}_i(D) \quad (\$/h) \tag{3.1}$$

An illustration of ISO net surplus collection for a simple 2-bus system is depicted in Fig. 3.4. The LSE at bus 2 pays $\text{LMP}_2 > \text{LMP}_1$ for each unit of the cleared load p_L^F. However, M units of this cleared load are supplied by GenCo G1 at bus 1, who receives only LMP_1 per unit. The ISO net surplus collection is then given by M $\times$ [LMP_2 - LMP_1].

The standard ISO DC-OPF objective used on day D for deriving LMP and dispatch solutions for day D+1 is the maximization of *total net surplus*, measured by the area TotNetSur(D) between the total system demand and supply curves for day D+1. As carefully shown in Somani, A. and Tesfatsion, L. (2008), TotNetSur(D) can be expressed as the following sum of component surpluses:

$$\text{LSENetSur}(D) + \text{GenNetSur}(D) + \text{ISONetSur}(D) \tag{3.2}$$

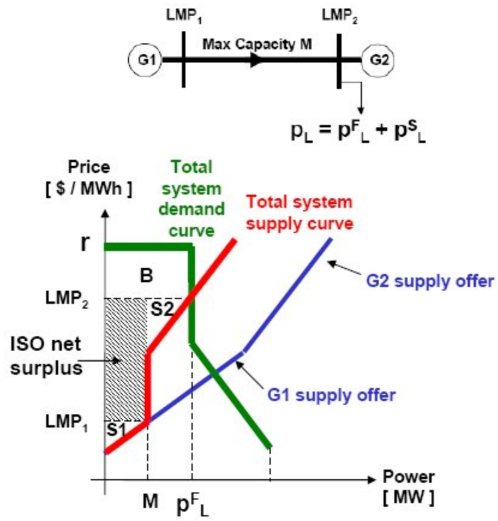

Figure 3.4 Illustration of ISO net surplus collection for a simple 2-bus system with a thermal branch limit M restricting power flow from the cheaper GenCo G1 at bus 1 to the load p_L at bus 2. (Figure adapted from Salazar, H. (2008).)

For example, TotNetSur in Fig. 3.4 is the sum of the LSE net surplus B, the GenCo net surpluses S1 and S2 extracted by G1 and G2, and the ISO net surplus. When GenCos are learners, the ISO constructs GenNetSur(D) in (3.2) using GenCo *reported* marginal costs (3.6) rather than *true* marginal costs (3.8).

3.1.4 Market Performance Measures

This section defines and explains the construction of the dynamic market performance measures reported in the tables and figures used for later sections.

In the absence of GenCo learning, the dynamic 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.

With GenCo learning under various demand conditions, there are different average measures:

- *Avg LMP* ($/MWh) denotes LMP averaged across all the buses and the 24 hours of a typical day-ahead market schedule;
- *Avg Total Demand* (MW) denotes LSE total demand averaged across the 24 hours of a typical day-ahead-market schedule;

- *Avg True TVCost* (\$/h) denotes GenCo true total avoidable cost averaged across total GenCos and the 24 hours of a typical day-ahead market schedule;

- *Avg LI* (unit-free number) denotes the GenCo Lerner Index value averaged across total GenCos and the 24 hours of a typical day-ahead-market schedule;

- *LMP spiking* (\$/MWh) denotes the maximum absolute difference between successive hourly LMPs averaged across all 24 hours when GenCos have different supply-offer price caps;

- *LMP volatility range* (\$/MWh) denotes the as [maxLMP-minLMP] averaged across all 24 hours when GenCos have different supply-offer price caps.

For each learning treatment, mean outcomes for the average hourly measures across thirty runs is calculated. Mean outcomes for measures are indicated by overlines.

(a) The mean outcomes for *Avg LMP* (\$/MWh) are calculated from a calculation day 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 the calculation day, determine the average hourly LMP across all 30 runs. Second, for each hour of the calculation day, determine the average of these run-averaged hourly LMP values across all buses. Finally, average these bus-averaged and run-averaged hourly LMP values across all 24 hours of the calculation day to get mean Avg LMP. Using a 5-bus test case as an example,

$$\overline{AvgLMP} = \frac{\left[\sum_{r=1}^{30}\sum_{k=1}^{5}\sum_{H=00}^{23} LMP_k(H,r)\right]}{30*5*24} . \tag{3.3}$$

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:

$$\sqrt{\frac{\left[\sum_{r=1}^{30}\sum_{k=1}^{5}\sum_{H=00}^{23}[LMP_k(H,r) - AvgLMP]^2\right]}{30*5*24}} . \tag{3.4}$$

(b) The mean outcomes for *Avg Total Demand* (MW) are calculated from a calculation day data as follows. First, for each LSEs and for each hour of the calculation day, determine the LSE’s average cleared (satisfied) price-sensitive demand across all 30 runs. Second, for each LSE and each hour of the calculation day, 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 total

LSEs to get average total demand. Finally, average these hourly average total demands across all 24 hours of the calculation day to get mean Avg Total Demand. The corresponding standard deviation is then calculated in the usual way using the “N” definition.

(c) The mean outcomes for *Avg RepTVCost* (\$/h) are calculated from a calculation day data as follows. First, for each GenCo and for each hour of the calculation day, determine the reported total avoidable costs of total 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. Second, for each hour of the calculation day, determine the average of these run-averaged reported total avoidable cost calculations across all GenCos. Third, average these GenCo-averaged and run-averaged hourly reported total avoidable cost calculations across all 24 hours of the calculation day to get mean Avg RepTVCost. The corresponding standard deviation is then calculated in the usual way using the “N” definition.

(d) 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:

$$LI_i \; = \; \frac{\left[LMP_{k(i)} - MC_i(p_{Gi})\right]}{LMP_{k(i)}} \; . \tag{3.5}$$

In (3.5), $\mathrm{LMP}_{k(i)}$ 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) are calculated from a calculation day data as follows. First, for each run, for each hour of the calculation day, and for each GenCo i with a positive power dispatch level p_{Gi} for this run and hour, determine the GenCo’s Lerner Index (3.5). Second, for each GenCo and each hour of the calculation day, 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 the calculation day, 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 the calculation day to get mean Avg LI. The corresponding standard deviation is then calculated in the usual way using the “N” definition.

(e) 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. 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) for learning GenCos under different supply-offer price caps are calculated from a calculation day 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 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 the calculation day. 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.

(f) The mean outcomes for LMP volatility range ($/MWh) for learning GenCos under different supply-offer price caps are calculated from a calculation day 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 transmission grid buses k, the *LMP volatility range* is calculated as [maxLMP-minLMP] across all 24 hours of the calculation day. 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.

(g) The LMP and IE frequency measures 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%.

3.2 GenCo Agent Model

3.2.1 Introduction

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.

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.

3.2.2 GenCo Supply Offers

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*

$$MC_i^R(p_{Gi}) \;=\; a_i^R + 2b_i^R p_{Gi} \quad (\$/MWh) \tag{3.6}$$

defined over a *reported operating capacity interval*

$$Cap_i^L \;\leq\; p_{Gi} \;\leq\; Cap_i^{RU} \quad (MW) \tag{3.7}$$

for real power p_{Gi}. The expression $MC_i^R(p_{Gi})$ 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*

$$MC_i(p_{Gi}) \;=\; a_i + 2b_i p_{Gi} \quad (\$/MWh)\ . \tag{3.8}$$

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

$$Cap_i^L \;\leq\; p_{Gi} \;\leq\; Cap_i^U \quad (MW)\ . \tag{3.9}$$

Note from the above discussion that each reported supply offer for GenCo i can be summarized in the form of a vector (a_i^R,b_i^R,Cap_i^{RU}).

3.2.3 GenCo Supply Offer Price Cap

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.

3.2.4 GenCo 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:

$$VCost_i(p_{Gi}) = \int_0^{p_{Gi}} MC_i(p)dp = a_i p_{Gi} + b_i[p_{Gi}]^2 \ (\$/h), \tag{3.10}$$

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

$$TC_i(p_{Gi}) \ = \ [VCost_i(p_{Gi}) + SCost_i] \ \ (\$/h), \tag{3.11}$$

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

$$\pi_i(H, D) \ = \ LMP_{k(i)} * p_{Gi} - TC_i(p_{Gi}) \ \ (\$/h). \tag{3.12}$$

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

$$NE_i(H, D) \ = \ LMP_{k(i)} * p_{Gi} - VCost_i(p_{Gi}) \ \ (\$/h). \tag{3.13}$$

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

$$NE_i(D) \ = \ \sum_{H=00}^{H=23} NE_i(H, D) \ \ (\$). \tag{3.14}$$

As will be seen in Chapter 4 Section 4.4, the estimates of 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) are used. Specifically,[3]

$$MaxDNE_i \;=\; 24 * \left(\max_{s_i^R \in AD_i} \left[HNE(s_i^R) \right] \right) \;(\$), \tag{3.15}$$

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

$$HNE(s_i^R) \;=\; MC_i^R(Cap_i^{RU}) * Cap_i^{RU} - VCost_i(Cap_i^{RU})\,. \tag{3.16}$$

3.2.5 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.

GenCos are autonomous energy traders with strategic learning capabilities. Each GenCo i adaptively chooses its supply offers ("actions") s$_i^R$ = (a$_i^R$,b$_i^R$,Cap$_i^{RU}$) 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 (Roth, A. E. and Ido, E. (1995),Erev, I. and Roth, A. E. (1998)) 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 M$_i$ ≥ 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 q$_{im}$(1) for m = 1,...,M$_i$. AMES permits the user to set these initial propensity levels to any real numbers. However,

[3]Compare (3.15) with definition (3.14) 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.

the assumption used in this study is that GenCo i's initial propensity levels are all set equal to some common value $q_i(1)$, as follows:

$$q_{im}(1) \;=\; q_i(1) \;\text{ for all supply offers } m \in AD_i \tag{3.17}$$

Now consider the beginning of any day $D \geq 1$, and suppose the current propensity of GenCo i to choose supply offer m in AD_i is given by $q_{im}(D)$. The *choice probabilities* that GenCo i uses to select a supply offer for day D are then constructed from these propensities as follows:[4]

$$p_{im}(D) = \frac{\exp(q_{im}(D)/T_i)}{\sum_{j=1}^{M_i} \exp(q_{ij}(D)/T_i)} \;,\; m \in AD_i \tag{3.18}$$

In (3.18), 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 (3.18) 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 (3.14) 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,[5]

$$q_{im}(D+1) \;=\; [1-\rho_i]q_{im}(D) \;+\; Response_{im}(D) \;, \tag{3.19}$$

[4] In the original algorithm developed by Roth and Erev (Roth, A. E. and Ido, E. (1995),Erev, I. and Roth, A. E. (1998)), 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.

[5] The response function appearing in (3.19) modifies the response function appearing in the original algorithm developed by Roth and Erev (Roth, A. E. and Ido, E. (1995),Erev, I. and Roth, A. E. (1998)). 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, J. , Petrov, V. and Tesfatsion, L. (2001) and Pentapalli, M. (2008) for detailed motivation, presentation, and comparative experimental tests of this modified response function.

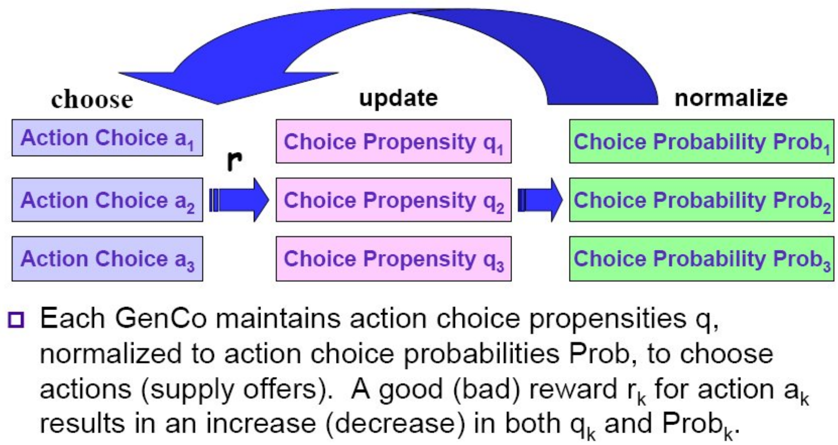

Figure 3.5 AMES GenCos use stochastic reinforcement learning to determine the supply offers they report to the ISO for the day-ahead market.

where

$$Response_{im}(D) = \begin{cases} [1 - e_i] \cdot NE_{im'}(D) & \text{if } m = m' \\ \\ e_i \cdot q_{im}(D)/[M_i - 1] & \text{if } m \neq m', \end{cases} \tag{3.20}$$

and $m \neq m'$ implies $M_i \geq 2$. The introduction of the *recency parameter* ρ_i in (3.19) acts as a damper on the growth of the propensities over time. The *experimentation parameter* e_i in (3.20) 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 is fully characterized once user-specified values are set for (M_i,$q_i(1)$,T_i,ρ_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 (3.17); T_i denotes the temperature parameter in (3.18); ρ_i denotes the recency parameter in (3.19); and e_i denotes the experimentation parameter in (3.20). 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.

Each GenCo's learning is implemented by means of a Java reinforcement learning module, *JReLM*, developed by Gieseler, see Gieseler, C. (2005). 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.

3.2.6 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.

As explained in Section 3.2.2, at the beginning of each day D each GenCo i must choose a supply offer $\mathrm{s}_i^R = (\mathrm{a}_i^R, \mathrm{b}_i^R, \mathrm{Cap}_i^{RU})$ to report to the ISO for each hour H of the day D+1 day-ahead market. Each supply offer s_i^R characterizes a reported marginal cost function

$$\mathrm{MC}_i^R(p) = a_i^R + 2b_i^R p \tag{3.21}$$

defined over a reported operating capacity interval

$$\mathrm{Cap}_i^L \ \leq \ p \leq \ \ \mathrm{Cap}_i^{RU} \tag{3.22}$$

Each GenCo i chooses its supply offers s_i^R 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, assume GenCo i only reports upward-sloping marginal cost functions, i.e., $\mathrm{b}_i^R > 0$. Second, assume GenCo i only reports non-trivial operating capacity intervals, i.e., $\mathrm{Cap}_i^L < \mathrm{Cap}_i^{RU}$. Third, assume that GenCo i only reports marginal cost functions that lie on or above its true marginal cost function (3.8) over the range of its reported operating capacity intervals. Fourth, assume GenCo i always reports an upper operating capacity limit Cap_i^{RU} that is less than or equal to its true upper operating capacity limit Cap_i^U.

Let a supply offer s_i^R for GenCo i be called *admissible* if the corresponding reported marginal cost function $\mathrm{MC}_i^R(p)$ and reported upper operating capacity limit Cap_i^{RU} are in compliance with these four simplifying assumptions. As shown in Sun and Tesfatsion (Sun, J. and Tesfatsion, L. , 2007a, Appendix), given any positive value for a *slope-start* parameter SS_i for GenCo i, any 4-dimensional vector s_i^A consisting of four components in percentage form can be uniquely mapped into an admissible supply offer s_i^R for GenCo i.

Referring to Table A.1 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 $M1_i \times M2_i \times M3_i = M_i$) and three range-index parameters $RIMax_i^L$, $RIMax_i^U$, and $RIMin_i^C$ in percentage form. The three density-control parameters control the *number* of distinct possible ordinate values a_i^R, slope values b_i^R, and upper operating capacity limits Cap_i^{RU}, 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 $SS_i > 0$, the M_i rows of this matrix constitute M_i distinct vectors s_i^A in percentage form that can be transformed uniquely into M_i distinct admissible supply offers s_i^R for GenCo i. Consequently, the matrix AD_i effectively constitutes an action domain for GenCo i consisting of M_i admissible supply offers s_i^R. Moreover, if the values for the action domain parameters ($M1_i$,$M2_i$,$M3_i$,$RIMax_i^L$,$RIMax_i^U$,$RIMin_i^C$,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 $\{AD_i : i = 1,...,I\}$ 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 B.2, in this dissertation set the action domain parameters identically across the GenCos to ensure equal cardinalities and similar densities of their action domains. In addition, 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_i^U), meaning that GenCo i always has the option of reporting its true marginal cost function and true operating capacity interval to the ISO.

3.3 LSE Agent Model

3.3.1 Introduction

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.

3.3.2 LSE Demand Bids

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

$$D_{jH}(p^S_{Lj}(H)) \;=\; c_j(H) - 2d_j(H)p^S_{Lj}(H) \quad (\$/MWh) \tag{3.23}$$

defined over a *true purchase capacity interval*

$$0 \;\leq\; p^S_{Lj}(H) \;\leq\; SLMax_j(H) \quad (MW) \tag{3.24}$$

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.

3.3.3 Relative Demand-Bid Price Sensitivity Measure

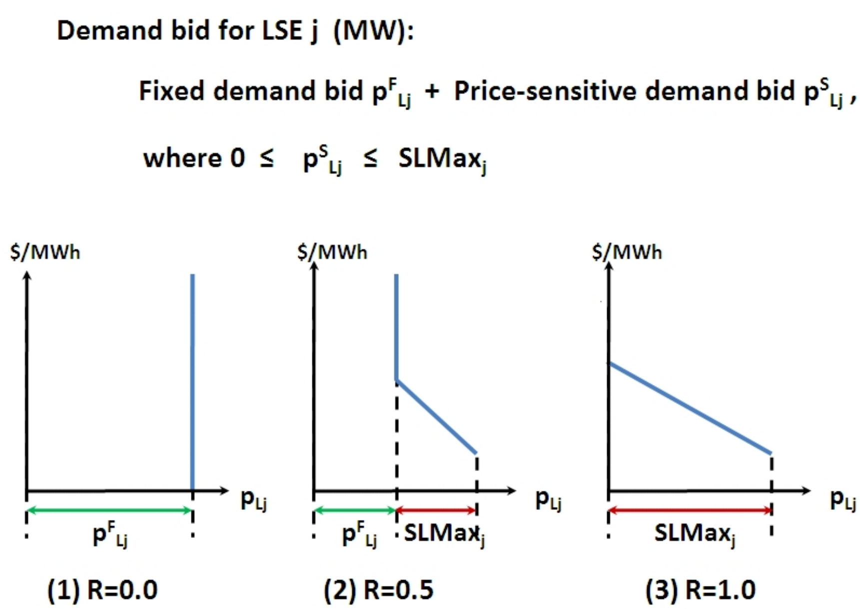

Figure 3.6 Illustration of the R ratio construction for the experimental control of relative demand-bid price sensitivity in each hour H.

R ratio is defined as maximum potential price-sensitive demand to maximum potential total demand. More precisely, for each LSE j and each hour H, let

$$R_j(H) \;=\; \frac{SLMax_j(H)}{MPTD_j(H)}\;. \tag{3.25}$$

In (3.25) 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 (3.24), and

$$MPTD_j(H) \;=\; [p^F_{Lj}(H) + SLMax_j(H)]\;\;(MW) \tag{3.26}$$

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 3.6 for the special cases R=0.0, R=0.5, and R=1.0.

To investigate LSE demand-bid price sensitivity in later experiments, the ratio R is systematically varied starting from R=0.0 (100% fixed demand) and ending with R=1.0 (100% price-sensitive demand). A positive R value indicates that the LSEs are able to exercise at least some degree of price resistance.

3.3.4 LSE Net Surplus

Suppose LSE j, located at bus $k(j)$, is cleared at a load level p_{Lj}(H,D) = [p^F_{Lj}(H,D)+p^S_{Lj}(H,D)] at price $\text{LMP}_{k(j)}$(H,D) for hour H of the day-ahead market for day D+1. The *payments* of LSE j over all 24 hours of day D+1, incurred at the end of day D, are

$$\text{Pay}_j(D) = \sum_{H=00}^{23} LMP_{k(j)}(H,D)\cdot p_{Lj}(H,D)\;(\$) \tag{3.27}$$

Using standard market efficiency analysis for buyers, see Tesfatsion, L. (2009a), the gross surplus GS_j(D) for LSE j for day D+1, incurred on day D, is then given by the revenue amount

$$\sum_{H=00}^{23}\left[r\cdot p^F_{Lj}(H,D) + \int_0^{p^S_{Lj}(H,D)} \text{D}_{jH}(p)dp\right] \tag{3.28}$$

and the *LSE net surplus* for day D+1, incurred on day D, is

$$\text{LSENetSur}(D) = \sum_{j=1}^{J}\left[\text{GS}_j(D) - \text{Pay}_j(D)\right] \tag{3.29}$$

3.4 Day-Ahead Market Agent Model

3.4.1 Introduction

Day-ahead market is a forward market where energy is sold prior to the real operating day, in which hourly LMPs are calculated for the next operating day based on supply offers and demand bids.

3.4.2 Day-Ahead Market Activities

- 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.

- 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.

- During the morning of each day D, each GenCo i uses its current action choice probabilities to choose a *supply offer* from its action domain AD_i to report to the ISO for use in all 24 hours of the day-ahead market for day D+1.

  Each supply offer in AD_i consists of a linear marginal cost function defined over an operating capacity interval. GenCo i's ability to vary its choice of a supply offer from 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. 3.5, 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.

The activities of the ISO on a typical day D are depicted in Fig. 3.2.

3.4.3 Day-Ahead Market Clearing Mechanism

3.4.3.1 DC-OPF for Day-ahead Market Clearing Mechanism

The DC-OPF problem formulation for 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 3.2.2, 3.2.4 and 3.3.2. An annotated listing of all of the variables used in the formulation is given in Tables A.1 and A.2. A more detailed discussion of this formulation can be found in Sun and Tesfatsion (Sun, J. and Tesfatsion, L. (2007a,b, 2008)).

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:[6]

$$GS_j(p^S_{Lj}) \;=\; \int_0^{p^S_{Lj}} D_j(p)dp \;=\; c_j \cdot p^S_{Lj} - d_j \cdot [p^S_{Lj}]^2 \tag{3.30}$$

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

$$TGS(\mathbf{p}^S_L) \;=\; \sum_{j=1}^{J} GS_j(p^S_{Lj}) \;, \tag{3.31}$$

where

$$\mathbf{p}^S_L \;=\; \left(p^S_{L1}, p^S_{L2}, \cdots, p^S_{LJ}\right) \tag{3.32}$$

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:

$$VCost^R_i(p_{Gi}) = \int_0^{p_{Gi}} MC^R_i(p)dp = a^R_i \cdot p_{Gi} + b^R_i \cdot [p_{Gi}]^2 \tag{3.33}$$

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

$$\mathbf{p}_G \;=\; (p_{G1}, p_{G2}, \cdots, p_{GI}) \tag{3.34}$$

[6]The gross surplus of LSE j corresponding to its *fixed* demand bid p^F_{Lj} is infinite if $\mathrm{p}^F_{Lj} > 0$; 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.

is then given by

$$TVC^R(\mathbf{p}_G) \;=\; \sum_{i=1}^{I} VC_i^R(p_{Gi}) \tag{3.35}$$

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:

$$TNS^R(\mathbf{p}_L^S, \mathbf{p}_G) \;=\; TGS(\mathbf{p}_L^S) - TVC^R(\mathbf{p}_G) \tag{3.36}$$

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

$$TNC^R(\mathbf{p}_L^S, \mathbf{p}_G) \;=\; -\,TNS^R(\mathbf{p}_L^S, \mathbf{p}_G) \tag{3.37}$$

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 objective function the standard objective function is augmented by inclusion of a penalty function for voltage angle differences. As carefully explained in Sun, J. and Tesfatsion, L. (2007a), 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.

DC-OPF Problem

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

Minimize

$$TNC^R(\mathbf{p}_L^S, \mathbf{p}_G) + \mu \left[\sum_{km \in BR} [\delta_k - \delta_m]^2 \right] \tag{3.38}$$

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

$$p_{Lj}^{S},\ j = 1,...,J;\ p_{Gi},\ i = 1,...,I;\ \delta_k,\ k = 1,...,K \tag{3.39}$$

subject to

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

$$\sum_{i \in I_k} p_{Gi} \ - \ \sum_{j \in J_k} p_{Lj}^{S} \ - \sum_{km \text{ or } mk \ \in \ BR} P_{km} \ = \ \sum_{j \in J_k} p_{Lj}^{F} \tag{3.40}$$

where

$$P_{km} \ = \ B_{km}\left[\delta_k - \delta_m\right] \tag{3.41}$$

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

$$|P_{km}| \ \leq \ P_{km}^{U} \tag{3.42}$$

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

$$Cap_i^{L} \ \leq \ p_{Gi} \ \leq \ Cap_i^{RU} \tag{3.43}$$

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

$$0 \ \leq \ p_{Lj}^{S} \ \leq \ SLMax_j \tag{3.44}$$

Voltage angle setting at angle reference bus 1:

$$\delta_1 \ = \ 0 \tag{3.45}$$

The shadow price (Lagrange multiplier) solution for the real power balance constraint (3.40) 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 change in the minimized DC-OPF objective function (\$/h) with respect to a change in fixed demand (MW) at bus k; see Liu, H. et al. (2009) 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 (3.38), 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 (3.39) of choice variables for the DC-OPF problem.

3.4.3.2 SCUC for Day-Ahead Market Clearing Mechanism

For current ISOs, the day-ahead energy market is cleared using *Security-Constrained Unit Commitment (SCUC)* and *Security-Constrained Economic Dispatch (SCED)* programs to satisfy energy demand requirements. The results of the day-ahead market clearing include hourly LMP values as well as hourly demand and supply quantities.

Lagrangian relaxation and Benders decomposition methods are most commonly used to solve SCUC optimization problems. The Lagrangian relaxation method uses a direct technique to get SCUC results. Benders decomposition is used to decompose SCUC into a UC master problem and hourly security checking sub-problems, which are coordinated through Benders cuts. Recently, mixed integer programming (MIP) has been successfully used to solve large-scale SCUC problems; see Hobbs, B. et al. (2001), Streiffert, D. et al. (2005), and Pinto, H. et al. (2006).

J.F. Benders introduced the Benders decomposition algorithm for solving large-scale, mixed-integer linear programming problems (MILP); see Benders, J. F. (1962). Geoffrion generalized this method, and made it applicable to nonlinear problems; see Geoffrion, A. M. (1972). When applying Benders decomposition, the original problem is decomposed into a master problem and several subproblems. Generally, the master-program is an integer problem and the subproblems are linear programs. The final solution based on the Benders decomposition algorithm can require iterations between the master problem and subproblems.

Fig. 3.7 presents a flowchart depicting how the Benders decomposition method is used to solve SCUC problems. Unit commitment is treated as the master problem, which contains only binary variables representing unit on/off and start/stop states. A feasibility problem is used for hourly security checks using outputs of the master problem. If some violations arise, then Benders feasibility cuts are added to the original master problem and the master problem is re-solved. If there are no violations for the feasibility problem, the optimality problem is tested for economic dispatch. If the optimality check is not satisfied, then Benders optimality cuts as well as feasibility cuts are added to the original master problem, and the master problem is re-solved for another iteration.

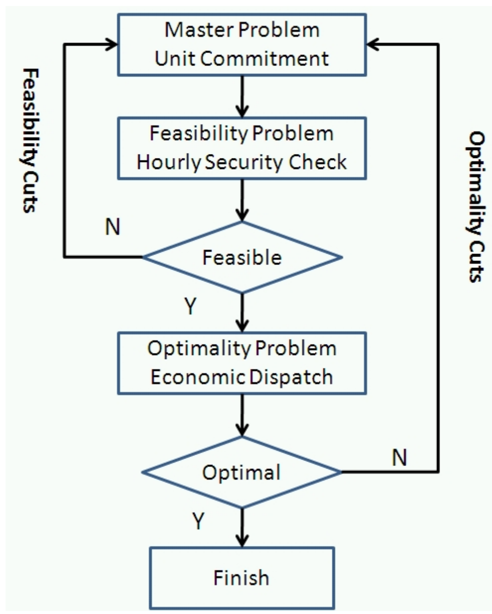

Figure 3.7 SCUC flowchart depicting the Benders decomposition method

For SCUC implementation, one Mixed Integer Linear Programming (MILP) solver - lp_solve (V5.5) is used. The reason is (1) it is open-source software that is easy to assess, (2) although coded in C language, it has Java language interface and easy to program.

SCUC Problem

The form of the SCUC problem is as follows:

Minimize

$$Z = \sum_{t=1}^{T} \sum_{i=1}^{I} \mathrm{st}_i \alpha_{it} + \mathrm{sd}_i \beta_{it} + \mathrm{c}_i \mathrm{p}_{it} \tag{3.46}$$

with respect to LSE real-power fixed demands, GenCo commitment status (on/off), start/stop status, real-power generation levels, and voltage angles:

$$p^F_{Ljt},\ j = 1, ..., J;\ I_{it},\ ST_{it},\ SD_{it},\ p_{Git},\ i = 1, ..., I;\ \delta_{kt},\ k = 1, ..., K; t = 1, ..., T \tag{3.47}$$

subject to

GenCo on/off and start status (binary) requirements:

$$I_{it}\ -\ I_{i(t-1)}\ -\ ST_{it}\ \leq\ 0 \tag{3.48}$$

GenCo on/off and stop status (binary) requirements:

$$I_{it}\ -\ I_{i(t-1)}\ +\ SD_{it}\ \geq\ 0 \tag{3.49}$$

System reserve requirements:

$$\sum_{i=1}^{I} Cap_i^{RU} I_{it}\ \geq\ D_t\ +\ R_t \tag{3.50}$$

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

$$\sum_{i\in I_k} p_{Git}\ -\ \sum_{km\,\mathrm{or}\,mk\ \in\ BR} P_{kmt}\ =\ \sum_{j\in J_k} p^F_{Ljt} \tag{3.51}$$

where

$$P_{kmt}\ =\ B_{km}\,[\delta_{kt} - \delta_{mt}] \tag{3.52}$$

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

$$|P_{kmt}|\ \leq\ P^U_{km} \tag{3.53}$$

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

$$Cap_i^L\ \leq\ p_{Git}\ \leq\ Cap_i^{RU} \tag{3.54}$$

Reported real-power ramping rates for each GenCo i = 1,...,I:

$$p_{Git} \; - \; p_{Gi(t-1)} \leq Ramp_{Gi} \tag{3.55}$$

Voltage angle setting at angle reference bus 1:

$$\delta_{1t} \;\; = \;\; 0 \tag{3.56}$$

The shadow price (Lagrange multiplier) solution for the real power balance constraint (3.51) at bus k, denoted by LMP$_{kt}$, constitutes the locational marginal price for bus k.

(1) SCUC Master Problem

The form of the SCUC master problem is as follows:

Minimize Zlower

$$Zlower \geq \sum_{t=1}^{T} \sum_{i=1}^{I} \mathrm{st}_i \alpha_{it} + \mathrm{sd}_i \beta_{it} \tag{3.57}$$

with respect to GenCo commitment status (on/off), start/stop status:

$$I_{it}, \; ST_{it}, \; SD_{it}, \; i = 1, ..., I; \; t = 1, ..., T \tag{3.58}$$

subject to

GenCo on/off and start status (binary) requirements:

$$I_{it} \; - \; I_{i(t-1)} \; - \; ST_{it} \; \leq \; 0 \tag{3.59}$$

GenCo on/off and stop status (binary) requirements:

$$I_{it} \; - \; I_{i(t-1)} \; + \; SD_{it} \; \geq \; 0 \tag{3.60}$$

System reserve requirements:

$$\sum_{i=1}^{I} Cap_i^{RU} I_{it} \; \geq \; D_t \; + \; R_t \tag{3.61}$$

(2) SCUC Hourly Feasibility Check Problem

The form of the SCUC hourly feasibility check problem is as follows:

Minimize

$$LoadSheding_t = \sum_{j=1}^{J} \mathrm{P}_{Ljt} \tag{3.62}$$

with respect to LSE real-power fixed demands, LSE real-power load shedings, GenCo real-power generation levels, and voltage angles:

$$p^F_{Ljt},\ p_{Ljt},\ j = 1, ..., J;\ p_{Git},\ i = 1, ..., I;\ \delta_{kt},\ k = 1, ..., K;\ t = 1, ..., T \tag{3.63}$$

subject to

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

$$\sum_{i \in I_k} p_{Git} - \sum_{j \in J_k} p_{Ljt} - \sum_{km \text{ or } mk \in BR} P_{kmt} = \sum_{j \in J_k} p^F_{Ljt} \tag{3.64}$$

where

$$P_{kmt} = B_{km}\left[\delta_{kt} - \delta_{mt}\right] \tag{3.65}$$

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

$$|P_{kmt}| \leq P^U_{km} \tag{3.66}$$

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

$$Cap^L_i I_{it} \leq p_{Git} \leq Cap^{RU}_i I_{it} \tag{3.67}$$

Voltage angle setting at angle reference bus 1:

$$\delta_{1t} = 0 \tag{3.68}$$

If $LoadSheding_t > 0$, then the corresponding infeasibility cut is generated as

$$LoadSheding_t + \sum_{i=1}^{I} \overline{\lambda_{it}} Cap_i^{RU} (I_{it} - \widehat{I_{it}}) - \sum_{i=1}^{I} \underline{\lambda_{it}} Cap_i^{L} (I_{it} - \widehat{I_{it}}) \leq 0 \tag{3.69}$$

where $LoadSheding_t$ is the objective of SCUC hourly feasibility check problem, the $\overline{\lambda_{it}}$ and $\underline{\lambda_{it}}$ are multipliers for reported real-power operating capacity interval constraints (3.67), and $\widehat{I_{it}}$ are GenCos on/off status reqults from SCUC master problem.

(3) SCUC Optimality Check Problem

The form of the SCUC optimality check (economic dispatch) problem is as follows:

Minimize

$$Z2 = \sum_{t=1}^{T} \sum_{i=1}^{I} \mathrm{c}_i \mathrm{p}_{it} \tag{3.70}$$

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

$$p_{Ljt}^F,\ j = 1, ..., J;\ p_{Git},\ i = 1, ..., I;\ \delta_{kt},\ k = 1, ..., K; t = 1, ..., T \tag{3.71}$$

subject to

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

$$\sum_{i \in I_k} p_{Git} - \sum_{km \,\mathrm{or}\, mk \,\in\, BR} P_{kmt} = \sum_{j \in J_k} p_{Ljt}^F \tag{3.72}$$

where

$$P_{kmt} = B_{km} \left[\delta_{kt} - \delta_{mt}\right] \tag{3.73}$$

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

$$|P_{kmt}| \leq P_{km}^U \tag{3.74}$$

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

$$Cap_i^L I_{it} \leq p_{Git} \leq Cap_i^{RU} I_{it} \tag{3.75}$$

Reported real-power ramping rates for each GenCo i = 1,...,I:

$$p_{Git} - p_{Gi(t-1)} \leq Ramp_{Gi} \tag{3.76}$$

Voltage angle setting at reference bus 1:

$$\delta_{1t} = 0 \tag{3.77}$$

If $Zupper = Zlower + Z2$, and Zlower and Zupper don't satisfy converge creteria, for example,

$$(Zupper - Zlower)/(Zupper + Zlower) \leq \epsilon \tag{3.78}$$

where ϵ is the specified converge creteria number,

then the corresponding feasibility cut is generated as

$$Zlower \geq \sum_{t=1}^{T}\sum_{i=1}^{I} \mathrm{st}_i\alpha_{it} + \mathrm{sd}_i\beta_{it} + Z2 + \sum_{t=1}^{T}\sum_{i=1}^{I}\overline{\pi_{it}}Cap_i^{RU}(I_{it} - \widehat{I_{it}}) - \sum_{i=1}^{I}\underline{\pi_{it}}Cap_i^{L}(I_{it} - \widehat{I_{it}}) \tag{3.79}$$

where Z2 is the objective of SCUC optimality check (economic dispatch) problem, the $\overline{\pi_{it}}$ and $\underline{\pi_{it}}$ are multipliers for reported real-power operating capacity interval constraints (3.75), and $\widehat{I_{it}}$ are the GenCo on/off status results from the SCUC master problem.

3.5 Real-Time Market Agent Model

Based on the day-ahead clearing of offers, bids and schedules, the real-time market becomes the balancing market between what was cleared in the day-ahead and what is required to meet real-time energy needs. The real-time energy market provides a continuous process on a five-minute basis for least cost balancing of supply and demand while recognizing current operating conditions, forecasted conditions, and generator offers.

For the real-time market, a real-time SCED optimization program is used to simultaneously balance injections and withdrawals, manage congestion, and produce LMPs. The SCED runs every five minutes. The objective of the SCED algorithm is to minimize the cost of real-time energy procurement over the next dispatch interval, subject to transmission network constraints and generator operating constraints.

The real-time settlement is the difference between the day-ahead quantity and the real-time quantity multiplied by the real-time price. The real-time quantities are based on actual load and generation meter data. The day-ahead quantities are based on cleared day-ahead supply offers and demand bids. The real-time prices are calculated based on actual market conditions. The goal of a two-settlement system is to reduce the uncertainty of the difference between the day-ahead quantity and the real-time quantity, as well as to reduce the volatility of real-time market, since the majority of the quantity is settled in the day-ahead market.

To date, the AMES real-time market model has not been activated because shocks and disturbances to the system have not yet been considered, meaning all day-ahead market contracts are carried out as planned. Future work will more fully consider the parallel operation in AMES of day-ahead and real-time markets in the presence of shocks and disturbances.

CHAPTER 4. KEY FINDINGS OF LMP SEPARATION AND VOLATILITY STUDY

4.1 Introduction

LMP is widely used in FERC's market design. For example, for day-ahead market settlement, LMP is used to calculate payment to GenCos and charges to LSEs by using LMP multiplying cleared quantity of GenCos power supply and LSEs demand; for real-time market settlement, LMP is used to calculate the difference between day-ahead settlement by using LMP multiplying the difference between the day-ahead quantity and the real-time quantity of GenCos power supply and LSEs demand. So LMP is a core part of FERC's market design.

From FERC's market design and ISOs' Business Practice Manuals (BPM), LSE demand bids for day-ahead market are mixtures of fixed (price-insensitive) demands and price-sensitive demands. When look at real data for day-ahead market, using MISO's data as an example, currently the price-sensitive demand is only about 1% of the total bid-in demand for the day-ahead market. So a question is arising, what is LMP response if systematically change the percentage of price-sensitive demand from 0% to 100%?

In addition, another related question is what if the ISO imposes a price cap on GenCo supply offers for the day-ahead market? Do LMPs have a controllable upper limit? Will this give profit-seeking GenCos an incentive to report smaller-than-true max capacities?

Previous studies have derived analytical expressions for LMPs at a point in time, conditional on given grid, demand, and supply conditions; see, for example, Conejo, A. J. et al. (2005) and Orfanogianni, T. and Gross, G. (2007). 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.

For example, Sueyoshi, T. and Tadiparthi, G. R. (2008) 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 inter-tie 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 inter-tie lines is systematically increased.

This chapter undertakes 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).[1] 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.[2]

4.2 Experimental Design

This section explains the 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 later sections.

[1] 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 (2009) is currently only about 1%.

[2] Price-cap policies differ widely across U.S. restructured wholesale power markets. For example, MISO (2009) 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.

The experimental design for this study is based on the *benchmark dynamic 5-bus test case* presented in Table B.1. 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. 4.1, with branch reactances, locations of LSEs and GenCos, and initial hour-0 LSE fixed demand levels adopted from Lally, J. (2002).[3]

- True GenCo cost and capacity attributes are as depicted in Figure 4.2. GenCos range from GenCo 5, a relatively large coal-fired base load 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, M. , Yamin, H. and Li, Z. , 2002, p. 296-297); see Fig. 4.3. 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. 4.4.

- GenCos are non-learners, meaning they report supply offers to the ISO that convey their true marginal cost functions (3.8) and true operating capacity limits (3.9).

- There is no supply-offer price cap.

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 function (3.6) in pursuit of increased net earnings.[4]

[3]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.

[4]Recall that a detailed description of the VRE-RL algorithm is provided in Chapter 3 Section 3.2.5. In this study the GenCos are only allowed to exercise economic withholding of capacity through strategic marginal cost reporting. In separate

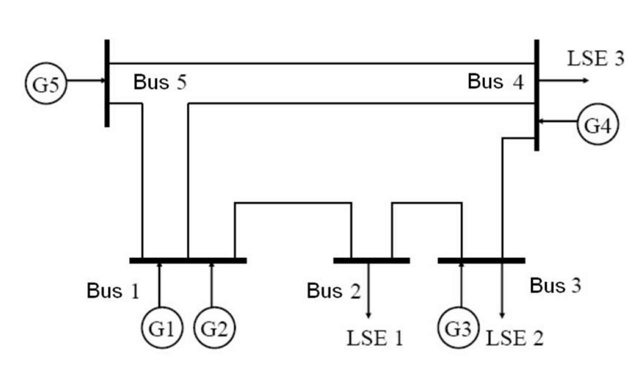

Figure 4.1 Transmission grid for the benchmark dynamic 5-bus test case.

For the learning treatments, the action domain AD_i for each GenCo i is constructed as in Li, H. , Sun, J. and Tesfatsion, L. (2008) 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 Pentapalli, M. (2008). 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 B.2 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, thirty seed values is generated via the standard Java class "random;" see Table B.2 for a listing of these seed values. For each learning treatment these thirty seed values are then used to implement thirty distinct runs, each 1000 simulated days in length.

The second treatment factor is the ratio R of maximum potential price-sensitive demand to maximum potential total demand. The construction of the R ratio is illustrated in Figure 3.6 for the special case R=0.0, R=0.5, and R=1.0.

For price-sensitive demand experiments to start by setting all of the R values (3.25) for each LSE j and each hour H equal to R=0.0 (the 100% fixed-demand case). 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. 4.4 with 100% fixed demand (R=0.0) to the true total

studies Li, H. and Tesfatsion, L. (2009d) and Chapter 6 the consequences of permitting GenCos to engage in physical withholding of capacity through strategic reporting of their operating capacity limits is explored.

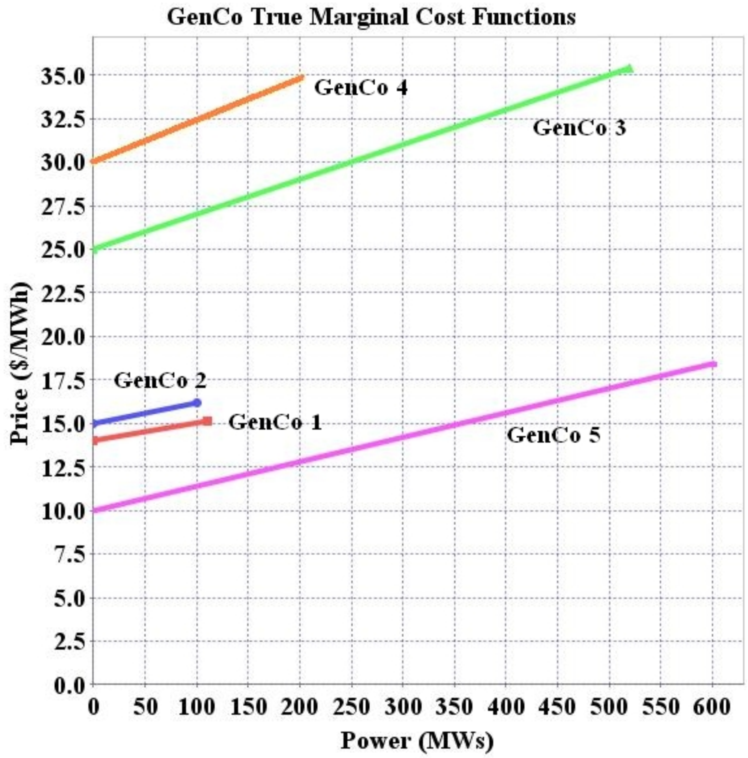

Figure 4.2 GenCo true marginal cost functions and true capacity attributes for the benchmark dynamic 5-bus test case.

demand curves in Fig. 4.5 with 20% potential price-sensitive demand (R=0.2).

The maximum potential price-sensitive hourly demands SLMax_j(H) for each LSE j are thus systematically increased across experiments. However, to 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 (3.26) is held constant across experiments by appropriate reductions in the fixed demand p^F_{Lj}(H) as SLMax_j(H) is increased. Specifically, MPTD_j(H) is set equal across all experiments to BP^F_{Lj}(H), the hour-H fixed-demand level BP^F(H) for LSE j depicted in Table B.1 for the benchmark dynamic 5-bus test case. Consequently, for each tested R value,

$$p^F_{Lj}(H) = [1 - R] * BP^F_{Lj}(H)\,; \tag{4.1}$$

$$SLMax_j(H) = R * BP^F_{Lj}(H). \tag{4.2}$$

Moreover, as R is incrementally increased from R=0.0 to R=1.0, to control for confounding effects arising from changes in the LSEs' price-sensitive demand bids by holding fixed the ordinate and slope

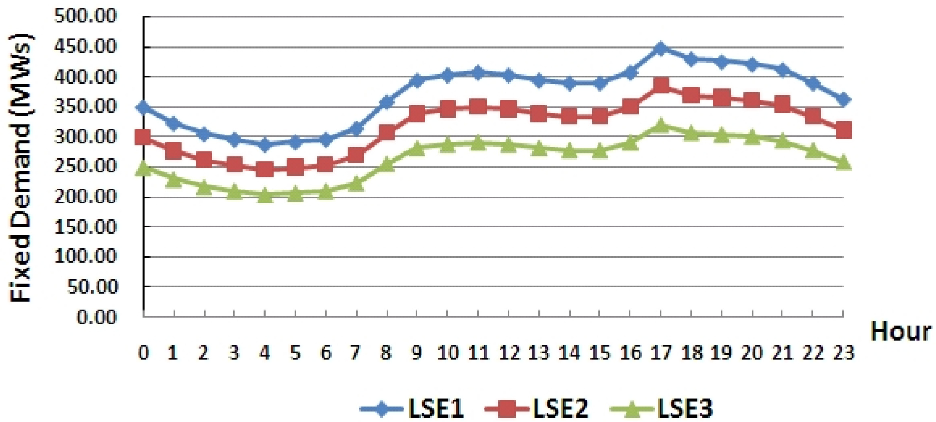

Figure 4.3 Daily LSE fixed demand (load) profiles for the benchmark dynamic 5-bus test case.

values $\{(c_j(H),d_j(H))$: H=00,...,23$\}$ for each LSE j. A listing of the specific numerical values used can be found in Table B.3.

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 c_j(H) is set equal to the hour-H LMP solution $LMP_{k(j)}$(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 c_j(H)/2d_j(H) for each LSE j is set to ensure that it is greater or equal to the $SLMax_j$(H) value determined by (4.2), as required by the admissibility restrictions imposed on LSE j's demand function D_{Hj}(p) in Table A.1.

The third treatment factor 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, $MC_i^R(Cap_i^{RU})$, does not exceed PCap.

As will be seen in the following Section 4.5 – in particular Table 4.8 – the mean outcome for average hourly LMP with GenCo learning and with 100% fixed demand (R=0.0) is approximately 140 ($/MWh). Therefore six PCap settings are tested 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

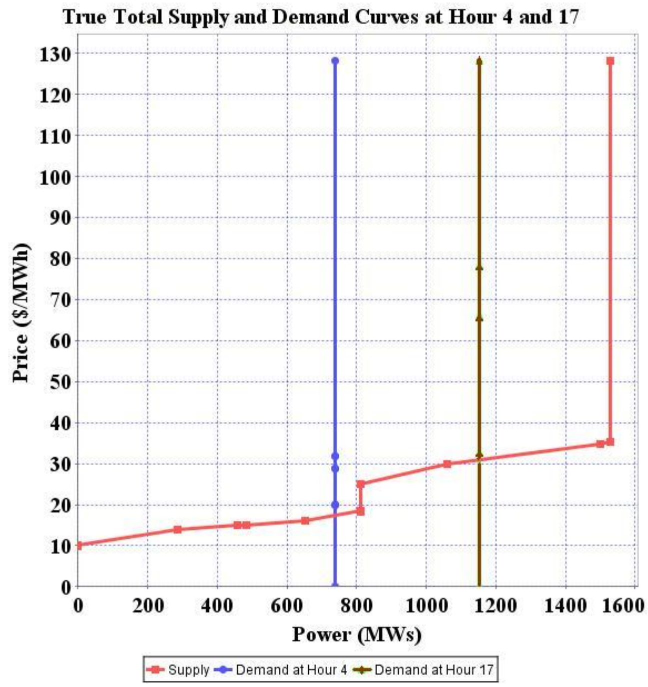

Figure 4.4 True total supply and demand curves for hours 4:00 and 17:00 for the benchmark dynamic 5-bus test case. Demand for this benchmark case is 100% fixed (R=0.0).

100; and (f) a very low value 80.

The latter sections use the experimental design 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

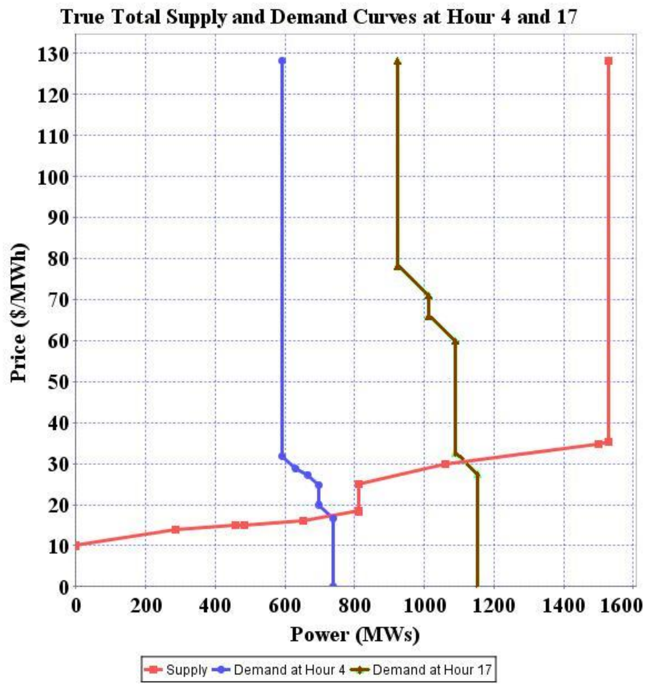

Figure 4.5 True total supply and demand curves for hours 4:00 and 17:00 for the benchmark dynamic 5-bus test case extended to include 20% potential price-sensitive demand (R=0.2).

For no-learning treatments, each of these measures for a typical day is calculated. For learning treatments, the mean of each of these measures across thirty runs for the final (1000th) simulated day are calculated. The precise definitions and calculations of these measures (with accompanying standard deviations) are provided in Section 3.1.4.

4.3 Without GenCo Learning Benchmark Case

Fig. 4.6 depicts hourly bus LMP levels, GenCo dispatch levels, and branch power flows during a typical day for the benchmark dynamic 5-bus test case. This benchmark case involves non-learning GenCos, 100% fixed demand, and no supply-offer price cap. Tables 4.1 through 4.3 provide more detailed numerical data on hourly bus LMPs, GenCo dispatch levels and branch power flows.

Table 4.1 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 4.2 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

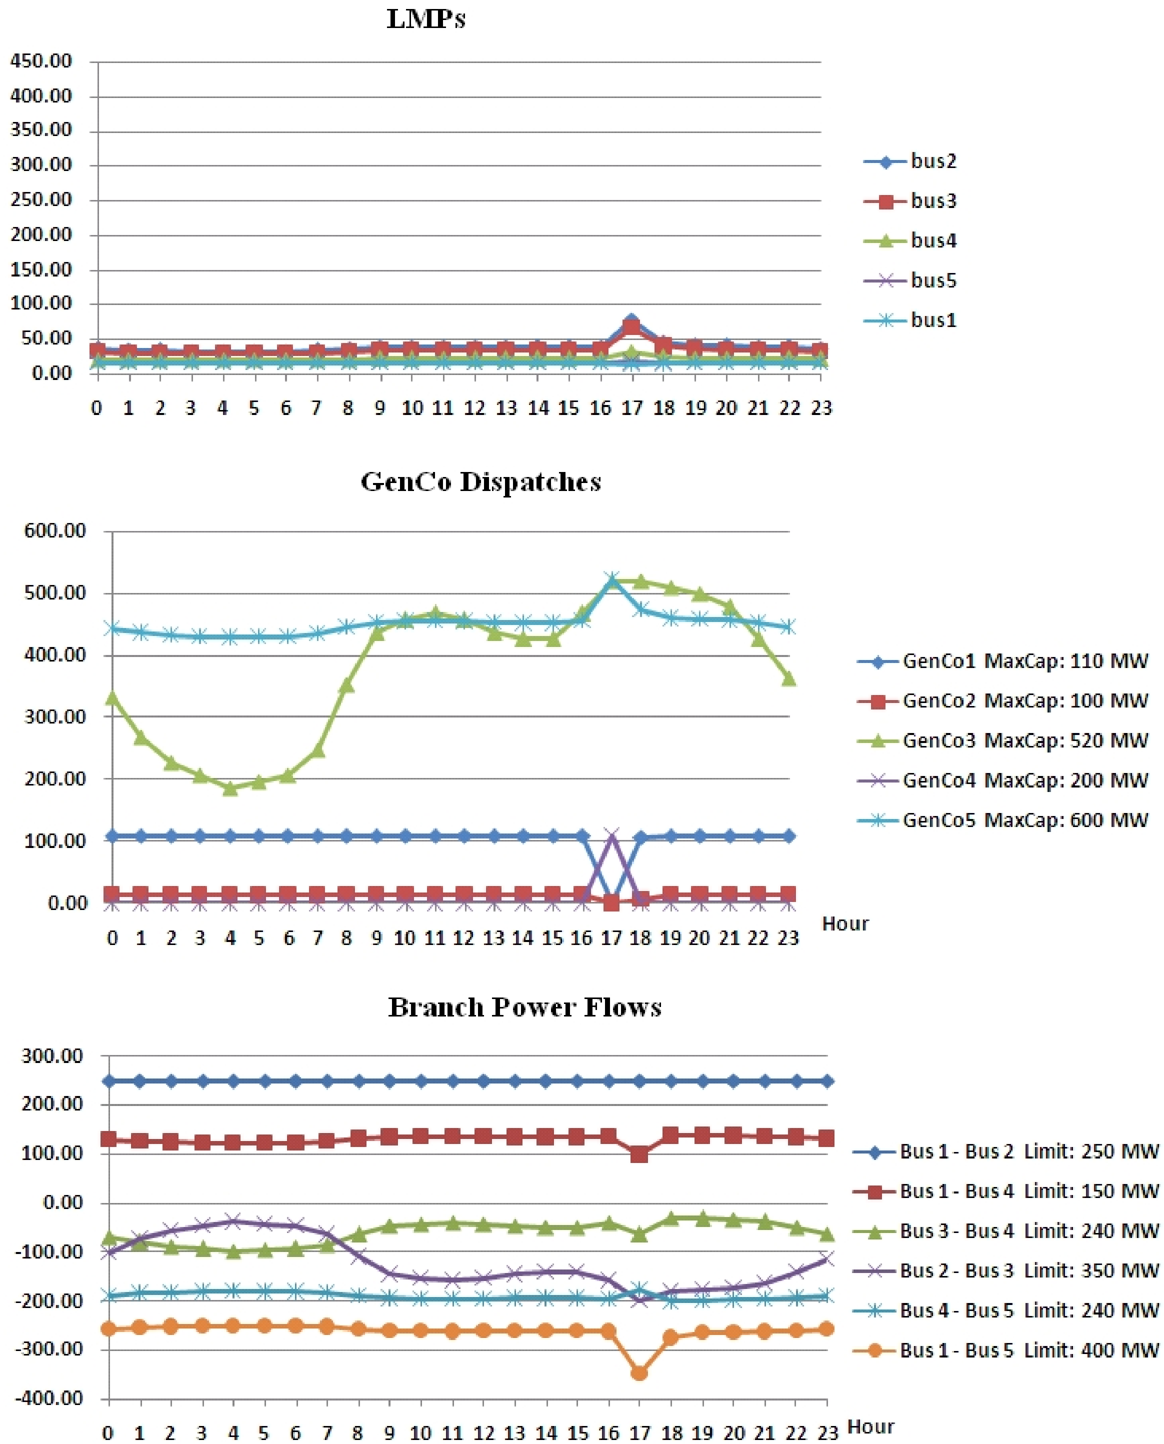

Figure 4.6 Hourly bus LMP levels, GenCo dispatch levels, and branch power flows during a typical day for the benchmark dynamic 5-bus test case (no GenCo learning).

Table 4.3 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

4.4 GenCo Learning Calibration for Economical Capacity Withholding

As a prelude to conducting experiments with GenCo learning, 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 B.2, a range of values for α_i and β_i are systematically tested, 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 (3.15);
- 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 4.7 depicts experimental findings for mean Total Gen DNE outcomes under alternative settings for

$$\alpha = \frac{q_i(1)}{MaxDNE_i}\ , \quad \beta = \frac{q_i(1)}{T_i}\ , \ i = 1, \ldots, I \tag{4.3}$$

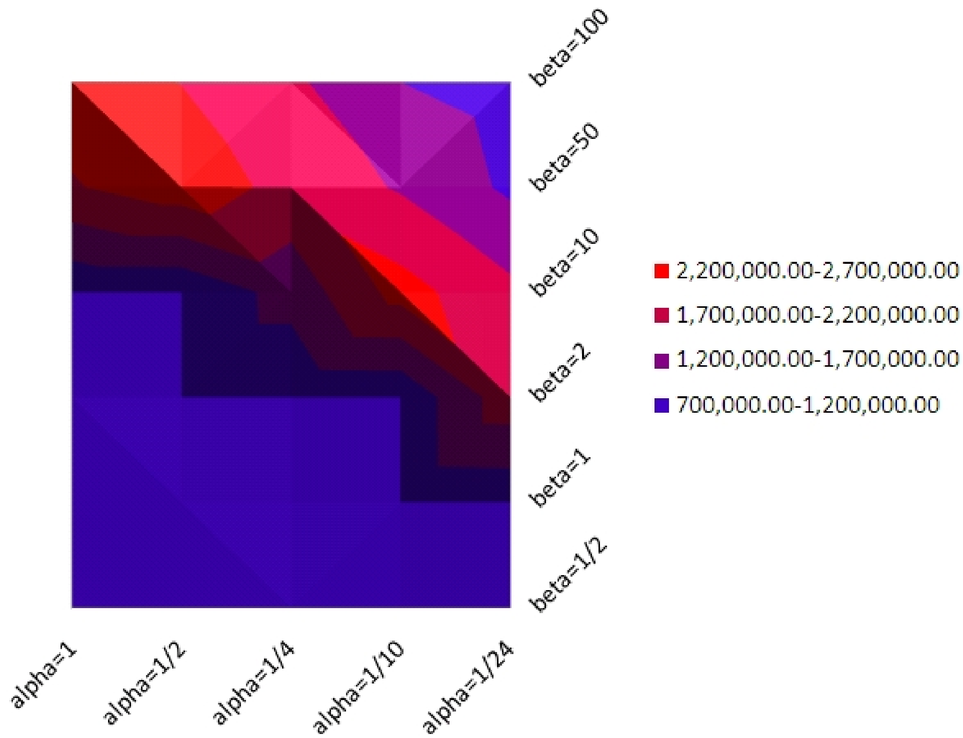

Figure 4.7 A 2D depiction of mean outcomes for total GenCo daily net earnings on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning under alternative settings for the two key stochastic reinforcement learning parameters (α, β).

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 (3.6). 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. 4.7 for the lowest tested β levels 1 and 1/2.

Figure 4.7 depicts mean total GenCo daily net earnings on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning under alternative settings for α and β, assuming these parameter values are set commonly across all GenCos. Two interesting findings are immediately evident. First, learning parameter specification substantially affects outcomes. Second, the highest outcomes are associated with “sweet spot” (α,β) combinations that lie along a nonlinear ridge line ranging from (α,β)=(1,100) in the northwest corner to (α,β)=(1/24,2) in the south-central region. The particular sweet-spot settings (α, β) = (1,100) are used in all of the learning treatments reported in the remainder of this study.

4.5 Pure GenCo Learning Experiments

Fig. 4.8 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. Tables 4.4 through 4.6 provide more detailed numerical data regarding the effects of learning on hourly bus LMPs, GenCo dispatch levels and branch power flows.

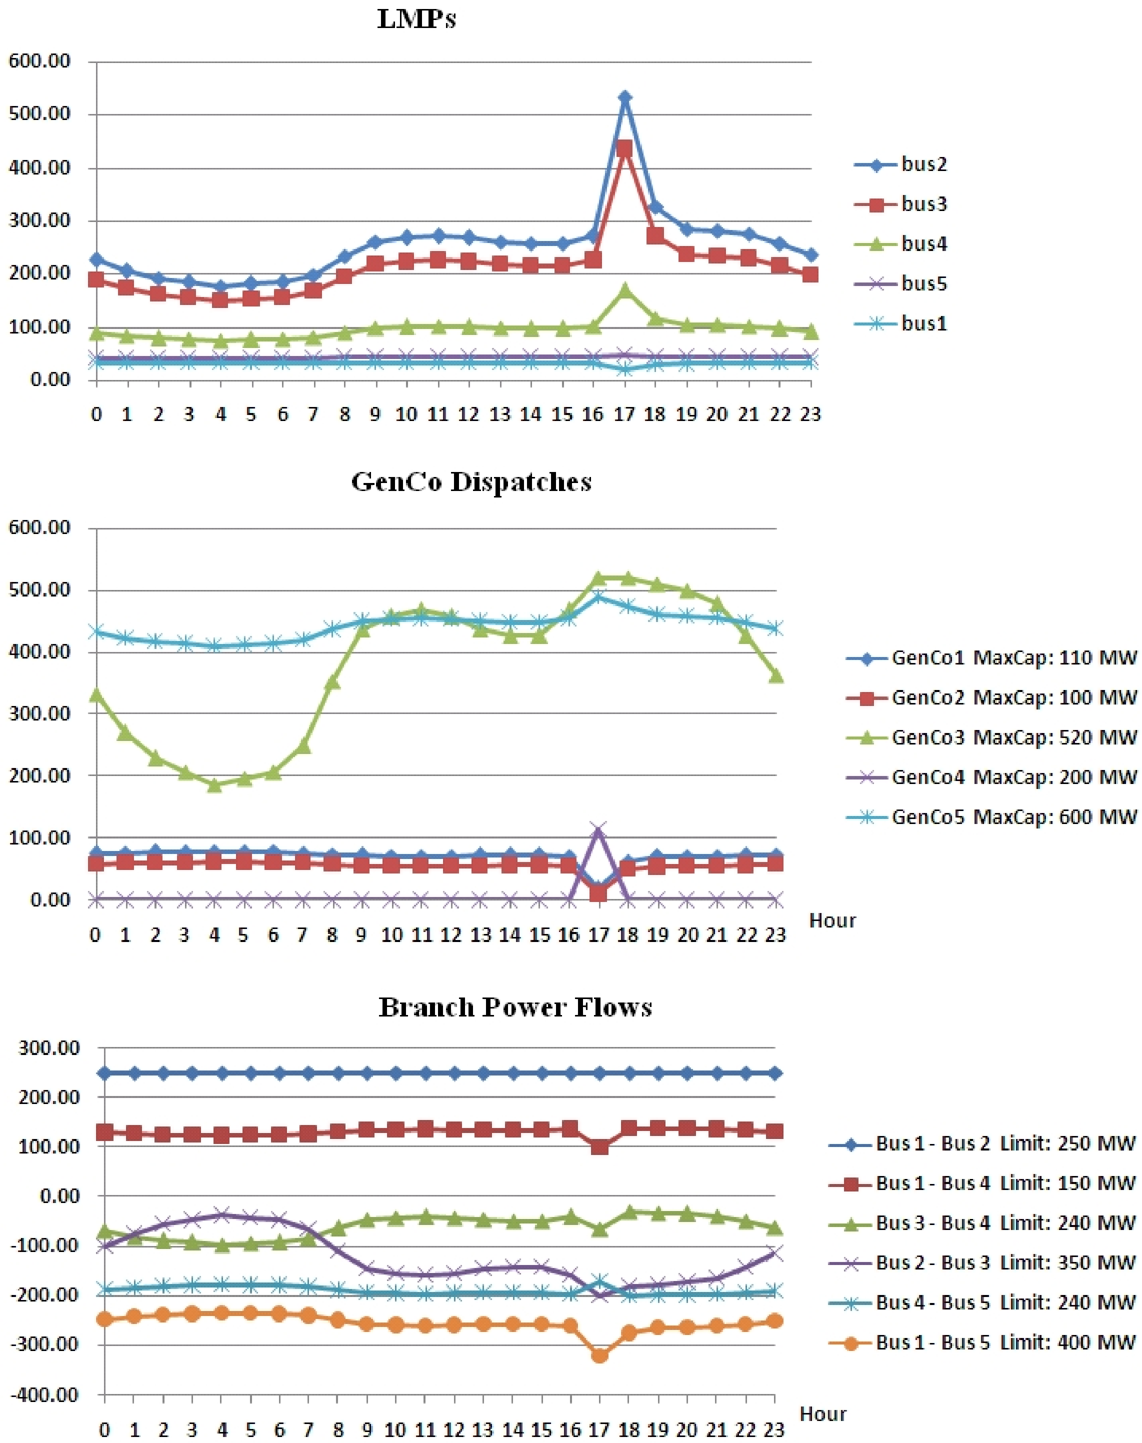

Figure 4.8 Hourly bus LMP levels, GenCo dispatch levels, and branch power flows for a typical run (ID=03) of the benchmark dynamic 5-bus test case extended to include GenCo learning.

Table 4.4 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

Comparing Figures 4.6 and 4.8, 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 4.1 through 4.6 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 4.6 and 4.8.

Consider, first, the benchmark no-learning case. To understand the pattern of LMPs reported in Table 4.1 for this benchmark case, it is important to understand congestion effects. Note from Table 4.3 that the branch 1-2 connecting bus 1 and bus 2 is congested in every hour; all other branches are uncongested in every hour.

Table 4.5 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

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 4.1, 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 4.1 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 (3.9). 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., Orfanogianni, T. and Gross, G. (2007).

Table 4.6 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

From Table 4.2, 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 4.2 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 4.1, 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 4.1, 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 4.4 for the benchmark case extended to include learning GenCos. Comparing these outcomes to the outcomes reported in Table 4.1 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 4.6, 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 4.5 for learning GenCos with the dispatch outcomes reported in Table 4.2 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 IS0. 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 withholding 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.

4.6 Price-Sensitivity Experiments without GenCo Learning

Table 4.7 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 4.7 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

As seen in Table 4.7, 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.[5]

4.7 Price-Sensitivity Experiments With GenCo Learning

Table 4.8 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.

Table 4.8 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)

Comparing the no-learning results presented in Table 4.7 to the results with GenCo learning presented in Table 4.8, 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

[5] 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 Tesfatsion, L. (2009a), 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.

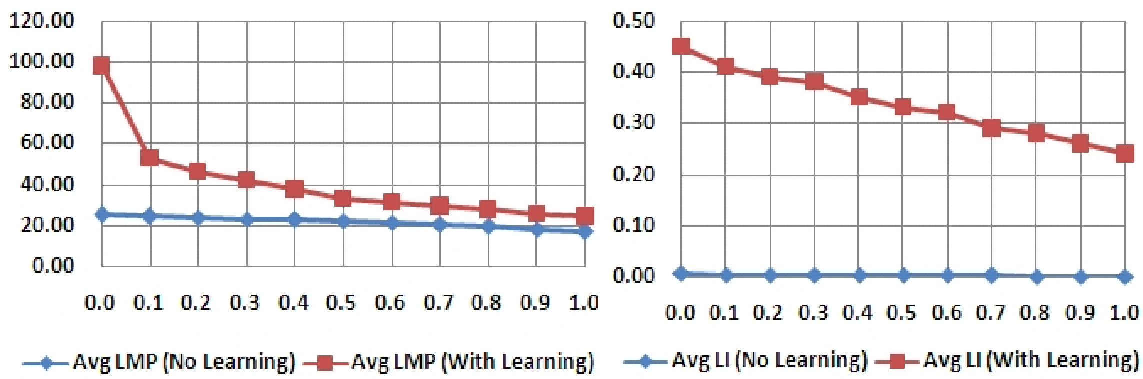

Figure 4.9 Mean outcomes for average hourly LMPs and LI levels on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and demand varying from R=0.0 (100% fixed) to R=1.0 (100% price sensitive).

before. However, as highlighted in Fig. 4.9, 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 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 Tesfatsion, L. (2009a), Rassenti, S. , Smith, V. L. and Wilson, B. (2003), 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.

4.8 GenCo Price Cap With and Without Learning

Table 4.9 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).[6] None of the tested PCap settings is binding on supply offers in the absence of GenCo learning. To see this, note in Figure 4.2 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.

Table 4.9 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%

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

[6]For the subsequent interpretation of these findings, it is important to recall from Chapter 3 Section 3.2.3 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 Tesfatsion, L. (2009a), 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.

capacities until their reported marginal costs at their reported maximum capacities are no greater than PCap.

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 4.9 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 4.9 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 substantial 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 4.9 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.

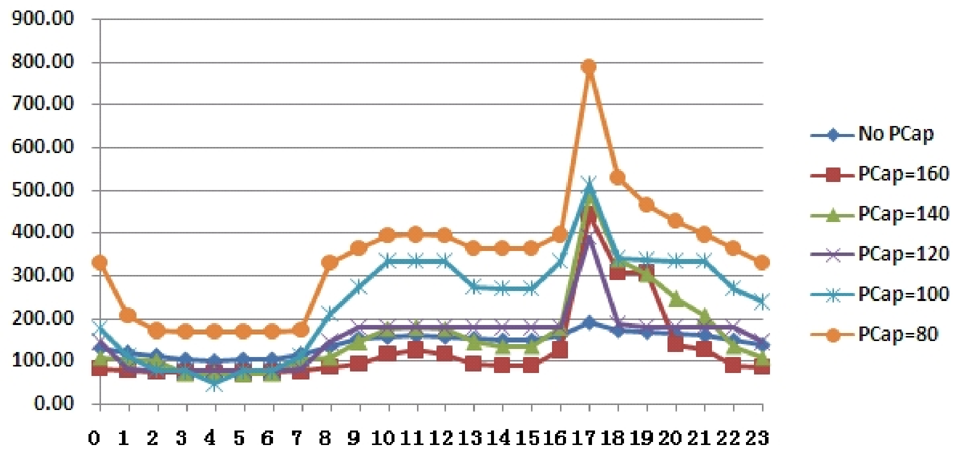

Figure 4.10 Mean outcomes for average hourly LMP levels on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and a supply-offer price cap varying from infinitely high (none) to low (80). Inadequacy Event (IE) LMP reserve prices are included in this figure.

Figure 4.10 provides a more disaggregated 24-hour depiction of the mean Avg LMP results reported in Table 4.9 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 4.11 presents still another way to visualize the mean Avg LMP outcomes reported in Table 4.9 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 relative to the no-PCap treatment. As explained more carefully in Chapter 3 Section 3.1.4, "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

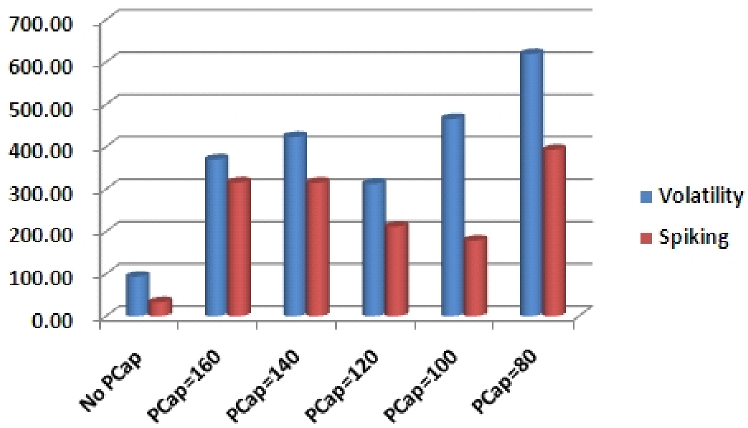

Figure 4.11 Mean outcomes for LMP spiking and LMP volatility range on day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and a supply-offer price cap varying from infinitely high (none) to low (80). Inadequacy Event (IE) LMP reserve prices are included in this figure.

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.

4.9 LMP Spatial Cross-Correlations

4.9.1 Correlation Experiment Preliminaries

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 4.3:

- 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.

4.9.2 GenCo Cross-Correlations

Table 4.10 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.

Table 4.10 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

4.9.3 GenCo-LMP Cross Correlations

Table 4.11 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 4.9.2. 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 4.11 is identical to the final column of values in Table 4.10. What explains these correlation patterns?

Table 4.11 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

One important explanatory factor is branch congestion and direction of branch power flows during hour 17. Recall from subsection 4.5 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 4.5. 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 4.6.

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

Table 4.12 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 4.11 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 4.10 and 4.11.

Table 4.12 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. 4.12 for the peak-demand hour 17 provides a visualization of the GenCo-LMP cross-correlation findings in Table 4.11. 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. 4.12 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. 4.12 for hour 17, note that GenCo 3's reported marginal cost is now perfectly positively correlated with the LMP at bus 3 and is 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 4.12.

Table 4.12 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%

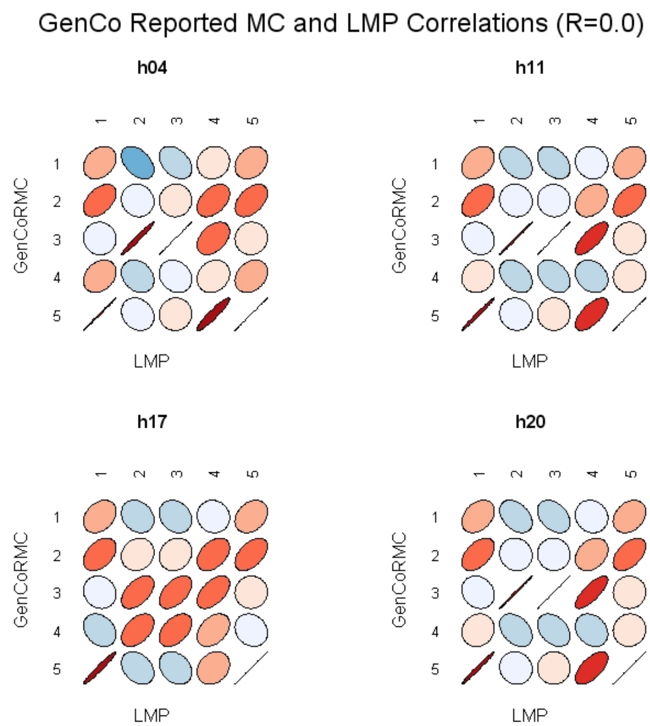

Figure 4.12 Pairwise cross-correlations between GenCo reported marginal costs and bus LMPs for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning. Demand for this case is 100% fixed (R=0.0).

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

Figures 4.13 and 4.14 report the effects on GenCo-LMP cross-correlations when the R ratio measuring the relative price-sensitivity of demand is systematically 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,

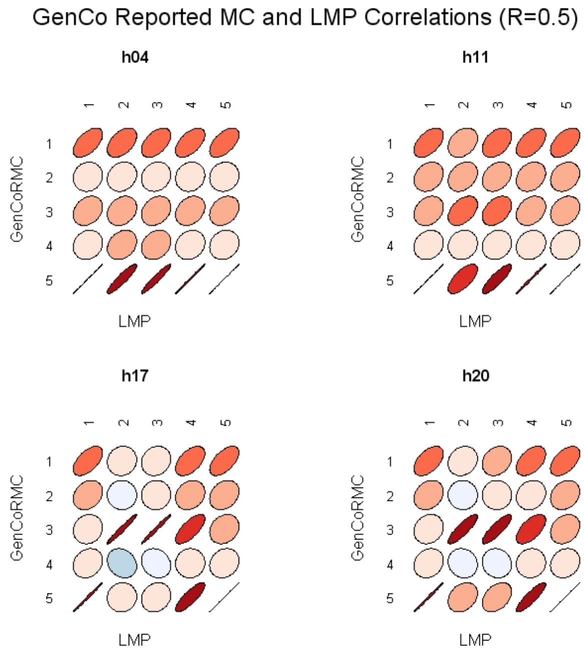

Figure 4.13 Pairwise cross-correlations between GenCo reported marginal costs and bus LMP for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and 50% potential price-sensitive demand (R=0.5).

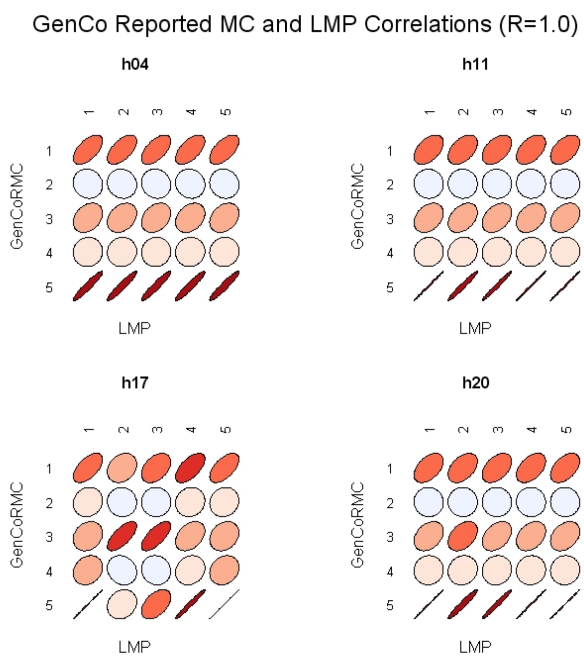

Figure 4.14 Pairwise cross-correlations between GenCo reported marginal costs and bus LMPs for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and 100% price-sensitive demand (R=1.0).

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 4.12 through 4.14. 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.

4.9.4 LMP-LMP Cross Correlations

Table 4.13 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.

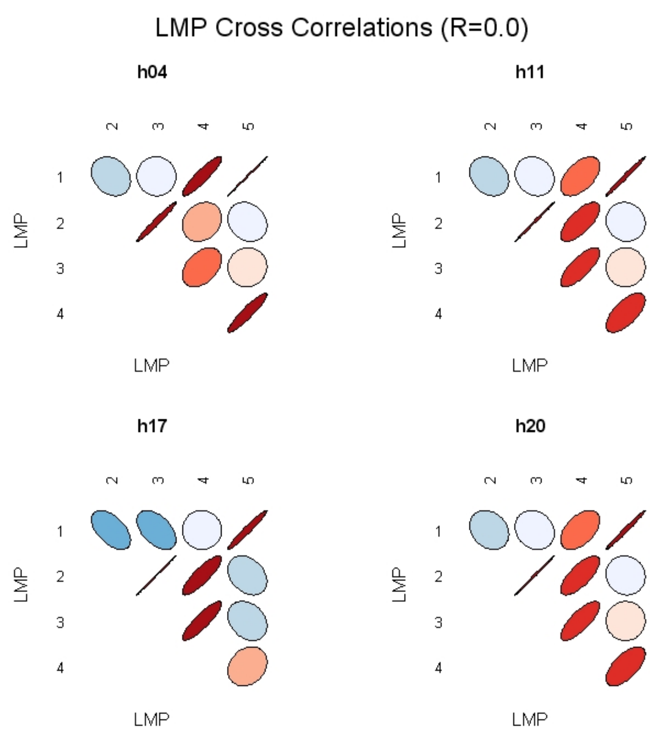

Figure 4.15 Pairwise LMP cross-correlations for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning. Demand for this case is 100% fixed (R=0.0).

Figs. 4.15 through 4.17 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 sen-

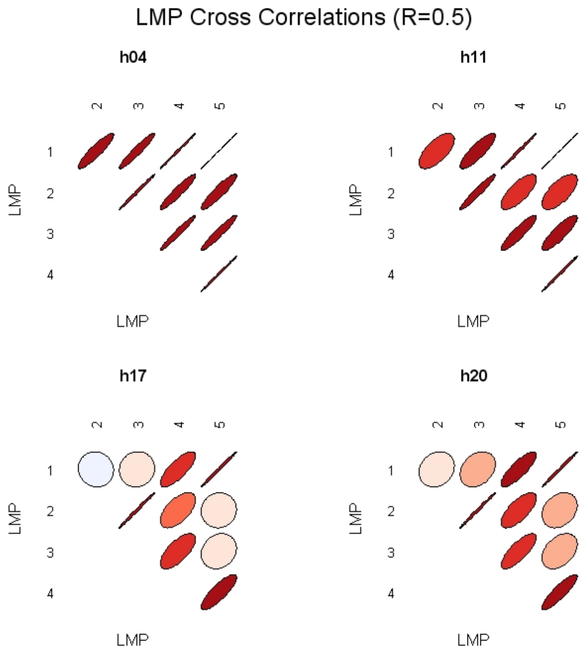

Figure 4.16 Pairwise LMP cross-correlations for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and 50% potential price-sensitive demand (R=0.5).

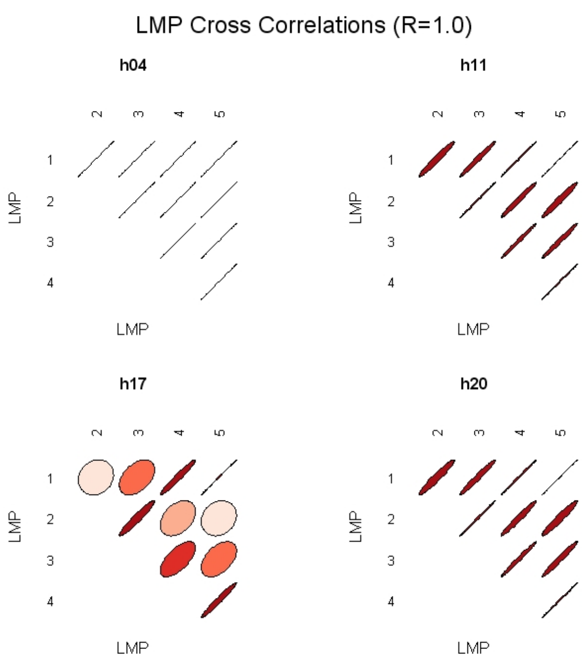

Figure 4.17 Pairwise LMP cross-correlations for hours 04, 11, 17, and 20 during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and 100% price-sensitive demand (R=1.0).

sitive).

Table 4.13 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

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.[7] 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 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 4.13 and Figs. 4.15-4.17 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 4.6. 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.

[7] See Liu, H. et al. (2009) for a rigorous presentation of this LMP derivation.

4.9.5 Empirical Evidence on LMP Correlations

In this subsection to calculate LMP cross-correlations using real-world price data. In particular, 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.[8] A BA is responsible for maintaining load-interchange-generation balance and the support of the Interconnection frequency.

From the geographical map depicted in Fig. 4.18, four neighboring BAs of MEC were picked 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). 24-hour historical data were obtained 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 Li, H. and Tesfatsion, L. (2009e). In particular, for ALTW the LMP for the load zone ALTW.MECB was used , and for the remaining four BAs interface LMPs were used. Then these data were used to calculate pairwise cross-correlations between the LMP reported for MEC and the LMPs reported at its four neighboring BAs.

Table 4.14 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

Table 4.14 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

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

are load-weighted prices determined for a load zone and interfaces and not for a single bus. The strong positive LMP cross-correlations in Table 4.14 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.

To differentiate between these various potential explanations for the strong positive correlations in Table 4.14—GenCo spillover effects, statistical artifact, and lack of congestion—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.

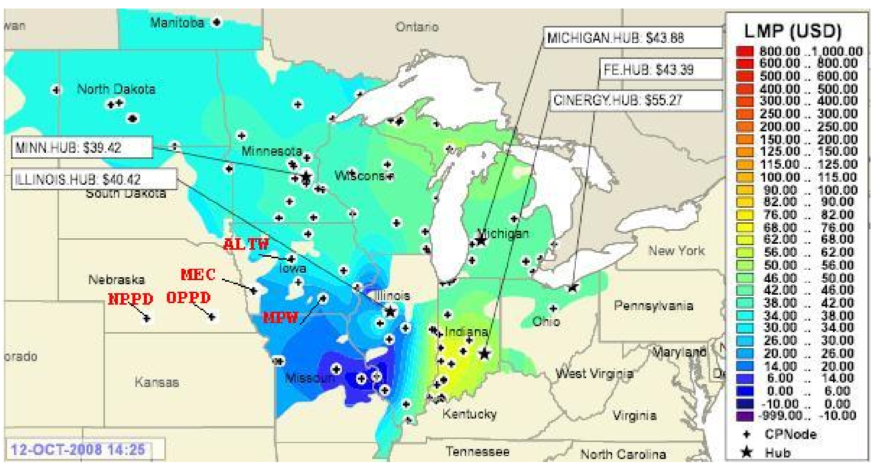

Figure 4.18 MidAmerican Energy Company (MEC) Balancing Authority and four neighboring Balancing Authorities in relation to MISO.

CHAPTER 5. ISO NET SURPLUS STUDY

This chapter uses dynamic 5-bus and 30-bus test cases to explore the social efficiency implications of the net surplus (congestion rents) collected and redistributed by ISOs in restructured wholesale power markets with grid congestion managed by locational marginal pricing (LMP). Demand price sensitivity and generator learning capabilities are taken as treatment factors. A key finding is that ISO net surplus substantially increases as the price-sensitivity of demand is reduced and the learning capabilities of generators are increased, conditions resulting in greater economic capacity withholding and a possible wastage of resources. A practical implication is that a more transparent public oversight of all net surplus collections and uses in wholesale power markets operating under LMP would be publicly prudent because these collections are not structurally well-aligned with social efficiency objectives.

5.1 Introduction

As elaborated in Joskow, P. (2006), over 50% of electric power generation in the U.S. is now traded at wholesale using locational marginal prices (LMPs). More precisely, an LMP at a particular grid location is the least cost to the system of providing an additional increment of power at that location. Congestion arising on any grid branch necessarily results in separation between the LMPS at two or more pricing locations. Ideally, as explained in Hausman, E. et al. (2006), this LMP separation should encourage transmission enhancements that relieve grid congestion, encourage new generation to locate where it has the greatest value, and encourage load to locate where it can be serviced most cheaply.

As is well-known, however, given branch and generation capacity limits, the physical laws regulating the flow of power on non-radial transmission grids can result in counter-intuitive LMP separation outcomes, see Gross, G. and Bompard, E. (2004), and Kirschen, D. S. and Strbac, G. (2004) Chp. 6. For example, LMPs at the two ends of a branch can separate without the branch being congested, power can flow from higher to lower price locations, and the price at a load-only location can be strictly higher or strictly lower than the marginal cost of all marginal (non-capacity-constrained) generation.

Another important LMP separation outcome is the creation of a net earnings stream whose use is discretionary to the independent system operator (ISO). When LMPs separate across the grid, the prices paid by load-serving entities (LSEs) can diverge from the prices paid to generation companies (GenCos). The difference between total LSE payments and total GenCo receipts, referred to below as *ISO net surplus*,is collected and redistributed by the ISO.

Previous research has shown that ISO net surplus is necessarily non-negative under standard DC OPF formulations; see Aldete, G. B. (2005) Prop. 2.1. Simulation findings in Gross, G. and Bompard, E. (2004) Fig. 11 demonstrate the strong sensitivity of ISO net surplus to variations in line flow limit specifications. To date, however, the social efficiency implications of ISO net surplus collections and uses do not appear to have been systematically examined.

Standard market efficiency analysis, entailing the maximization of the total net surplus extracted in a single market, is increasingly being applied to the study of power markets e.g., Liu, H. et al. (2008), Walawalkar, R. et al. (2008). However, this standard analysis focuses on seller and buyer net surplus collections. It does not consider the possibility that an agency tasked with clearing the market, here the ISO, is able to extract net surplus along with sellers and buyers, a feature that appears to raise conflict of interest issues. Moreover, it does not consider the more comprehensive issue of *social* efficiency.

This chapter uses the AMES Wholesale Power Market Test Bed to investigate ISO net surplus collections in relation to social efficiency for dynamic 5-bus and 30-bus test cases under systematically varied settings for demand-bid price sensitivity and the learning capabilities of GenCos.

5.2 Experimental Design

The benchmark dynamic 5-bus test case is detailed in Chapter 4 Section 4.3.

The benchmark dynamic 30-bus test case used in this study is a modified version of the IEEE 30-bus system presented in Shahidehpour, M. , Yamin, H. and Li, Z. (2002) App. D.4, 477-478. This 30-bus system has 9 GenCos, 21 LSEs, and 41 transmission grid branches. A complete input data file for this 30-bus system is included as Appendix C.

Two learning treatments are investigated in the experiments reported below: (a) As in the benchmark case, GenCos are non-learners; and (b) each GenCo i is a learning entity that makes daily use of a stochastic reinforcement learning algorithm to adjust the ordinate and slope parameters $\{a_i^R, b_i^R\}$ of its reported marginal cost functions (3.6) in pursuit of increased net earnings.

To control for random effects when GenCos are learners, thirty pseudo-random number seed values are used to initialize thirty distinct runs, each 1000 (5-bus) or 500 (30-bus) simulated days in length. See Table B.2 for a listing of these seed values.

To investigate LSE demand-bid price sensitivity, the ratio R of maximum potential price-sensitive demand to maximum potential total demand is systematically varied in each hour H, starting from R=0.0 (100% fixed demand) and ending with R=1.0 (100% price-sensitive demand). A positive R value indicates that the LSEs are able to exercise at least some degree of price resistance. Figure 3.6 illustrates the construction of R for the special cases R=0.0, R=0.5, and R=1.0.

5.3 5-Bus Benchmark Case

During a typical day D for the benchmark dynamic 5-bus test case the branch 1-2 connecting bus 1 to bus 2 is persistently congested. As a result, in each hour there is complete LMP separation across the grid.

As depicted in Table 5.1, GenCos 1 and 2 have relatively small net earnings in all hours and particularly in the peak-demand hour 17. This occurs for two reasons. First, as depicted in Fig. 4.1, these two GenCos are located at bus 1, hence they are semi-islanded away from the "load pocket" at buses 2 through 4 due to the persistent congestion on branch 1-2. Second, as seen in Fig. 4.2, these two GenCos have relatively small operating capacities.

In contrast, GenCo 3 located at the load-pocket bus 3 has relatively large net earnings in every hour, particularly in the peak-demand hour 17. This occurs because GenCo 3 is a pivotal supplier in most hours, meaning its relatively large capacity is needed to meet fixed demand. Moreover, during hour 17, GenCo 3 is dispatched at its maximum capacity and GenCo 5 is semi-islanded from bus 3 due to the congestion on branch 1-2. Consequently, to meet demand at bus 3 during hour 17, the ISO needs to call upon the expensive peaker unit, GenCo 4. This substantially spikes the LMP at bus 3 in hour 17,

Table 5.1 Hourly GenCo net earnings during a typical 24-hour day D for the benchmark dynamic 5-bus test case.

Hour	GenCo 1	GenCo 2	GenCo 3	GenCo 4	GenCo 5
00	67.81	1.15	1,105.79	0.00	1,377.42
01	67.24	1.08	725.83	0.00	1,340.07
02	66.87	1.04	518.48	0.00	1,315.68
03	66.68	1.02	427.08	0.00	1,303.45
04	66.49	0.99	345.93	0.00	1,291.50
05	66.59	1.01	385.44	0.00	1,297.48
06	66.68	1.02	427.08	0.00	1,303.45
07	67.06	1.06	618.74	0.00	1,327.95
08	68.00	1.18	1,247.51	0.00	1,389.76
09	68.75	1.28	1,909.70	0.00	1,440.36
10	68.94	1.30	2,097.94	0.00	1,453.20
11	69.03	1.31	2,193.68	0.00	1,459.54
12	68.94	1.30	2,097.94	0.00	1,453.20
13	68.75	1.28	1,909.70	0.00	1,440.36
14	68.66	1.26	1,820.44	0.00	1,434.06
15	68.66	1.26	1,820.44	0.00	1,434.06
16	69.03	1.31	2,193.68	0.00	1,459.54
17	0.02	0.00	18,654.46	142.27	1,912.03
18	57.62	0.22	4,980.40	0.00	1,573.60
19	69.41	1.37	2,601.82	0.00	1,485.24
20	69.31	1.35	2,497.56	0.00	1,478.84
21	69.13	1.33	2,291.68	0.00	1,465.89
22	68.66	1.26	1,820.44	0.00	1,434.06
23	68.09	1.19	1,324.32	0.00	1,396.18
Total	1,556.41	26.58	56,016.09	142.27	34,266.94

and hence the net earnings of GenCo 3.

GenCo 5 is essentially a base-load generator with large capacity and low marginal cost that is never dispatched at its maximum capacity. Consequently, although it is a pivotal supplier in most hours, its net earnings remain relatively flat.

Fig. 5.1 presents benchmark-case hourly financial flows during a typical day D. Note that LSE payments are persistently higher than GenCo revenues, particularly during the peak-demand hour 17. Consequently, ISO net surplus is persistently positive with a spike during hour 17.

Indeed, as will be seen below in Table 5.2, for a typical day D for the benchmark case (R=0.0), LSE payments are $754,919.61 and GenCo revenues are $545,508.54. Consequently, ISO net surplus is $209,411.07, which is about 2.3 times the amount $92,008.30 of GenCo net earnings.

5.4 5-Bus Case with Learning and Price-Sensitive Demand

For each R treatment, both with and without GenCo learning, congestion persistently occurs on branch 1-2. As seen in Fig. 4.9, however, the extension of the benchmark 5-bus test case to include

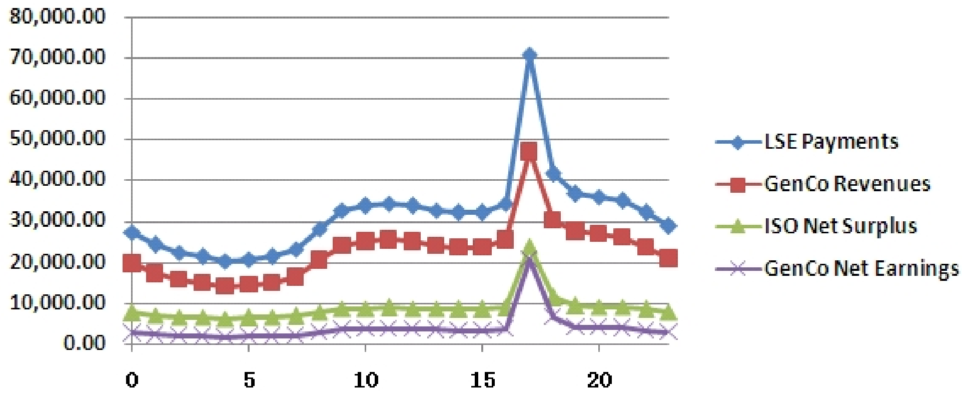

Figure 5.1 LSE payments, GenCo revenues, ISO net surplus, and GenCo net earnings during a typical 24-hour day D for the benchmark 5-bus test case.

GenCo learning and price-sensitive demand results in a substantial increase in mean LMP outcomes, particularly for small values of R. As more carefully explained in Li, H. , Sun, J. and Tesfatsion, L. (2009), this substantial LMP increase arises because each GenCo i learns over time to exercise *economic capacity withholding*, i.e., to submit to the ISO reported marginal cost functions (3.6) that lie strictly above its true marginal cost function (3.8).

This economic capacity withholding by the learning GenCos also has dramatic effects on ISO net surplus collection. These dramatic effects are graphically depicted in Figs. 5.2 and 5.3 and numerically reported in Tables 5.2 and 5.3.

Specifically, Fig. 5.2 and Table 5.2 present financial flows on a typical day D for the benchmark dynamic 5-bus test case extended to permit demand to vary from R=0.0 (100% fixed) to R=1.0 (100% price sensitive). As in the benchmark case, the GenCos submit supply offers to the ISO that reflect their true cost and capacity attributes. In contrast, Fig. 5.3 and Table 5.3 present corresponding financial flows on day 1000 for the case in which all five GenCos have learning capabilities. In particular, each GenCo applies stochastic reinforcement learning to its past net earnings outcomes in an attempt to determine which supply offer it should report to achieve the highest daily net earnings.

Consider, for example, the R=0.0 (100% fixed demand) daily data presented for the benchmark no-learning case in Fig. 5.2 and Table 5.2 and for the learning case in Fig. 5.3 and Table 5.3. Mean LSE payments on day 1000 for the learning case are $5,040,530.89, an approximately 6.7-fold increase

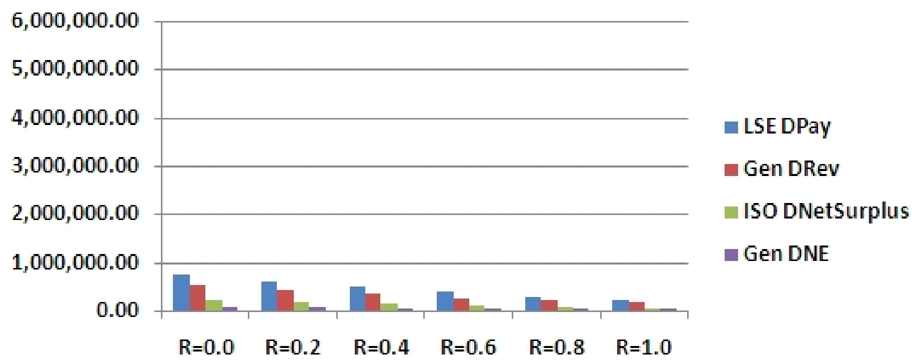

Figure 5.2 LSE payments, GenCo revenues, ISO net surplus, and GenCo net earnings during a typical day D for the benchmark dynamic 5-bus test case extended to permit demand to vary from R=0.0 (100% fixed) to R=1.0 (100% price sensitive)

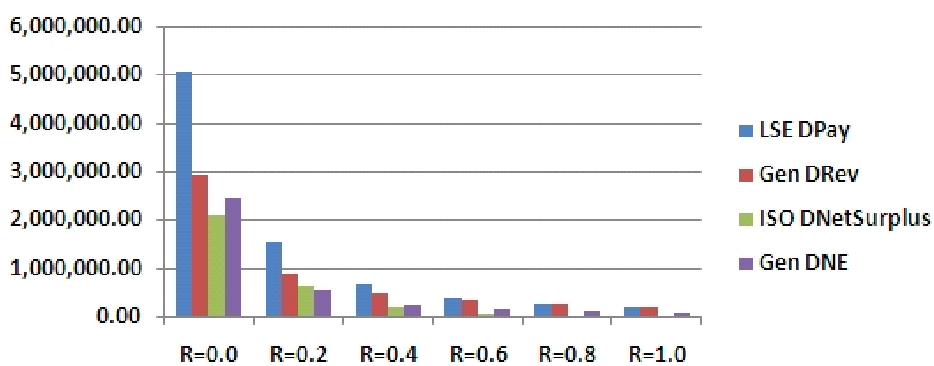

Figure 5.3 Mean outcomes for LSE payments, GenCo revenues, ISO net surplus, and GenCo net earnings during day 1000 for the benchmark dynamic 5-bus test case extended to include GenCo learning and demand varying from R=0.0 (100% fixed) to R=1.0 (100% price sensitive)

Table 5.2 GenCo net earnings, GenCo revenues, LSE payments, and ISO net surplus on a typical day D for the benchmark dynamic 5-bus test case extended to permit demand to vary 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 GenCo 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 GenCo 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

relative to the benchmark no-learning case. Also, mean GenCo revenues on day 1000 for the learning case are $2,942,909.93, an approximately 5.4-fold increase relative to the benchmark no-learning case. Note, however, that mean ISO net surplus on day 1000 for the learning case is then $2,097,620.96, an almost ten-fold increase relative to the benchmark no-learning case. Indeed, ISO net surplus under learning is similar in magnitude to GenCo net earnings ($2,441.646.71).

Since total demand for R=0.0 is the same under learning and no learning, the ten-fold increase in mean ISO net surplus collection under learning implies that the mean LMP paid by the LSEs is substantially higher than the mean LMP received by the GenCos. As suggested by Fig. 4.9, and more carefully detailed in Li, H. , Sun, J. and Tesfatsion, L. (2009), 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 buses 1 and 5, which have generation but no load.

Another regularity observed in Figs. 5.2 and 5.3, as well as in Tables 5.2 and 5.3, is that GenCo net earnings, GenCo revenues, LSE payments, and ISO net surplus all undergo marked monotonic

Table 5.3 Mean outcomes (with standard deviations) for GenCo net earnings, GenCo revenues, LSE payments, and ISO 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 GenCo $\overline{\text{DNE}}$	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 (257,39.64)	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 GenCo $\overline{\text{DRev}}$	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)

declines as R increases from R=0.0 (100% fixed demand) to R=1.0 (100% price-sensitive demand). The explanation for these monotonic declines is as follows.

Consider, first, the benchmark no-learning case. 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 carefully reported in Li, H. , Sun, J. and Tesfatsion, L. (2009), this results in lower LMPs, lower total demand, and lower avoidable costs of production. GenCo revenues and LSE payments are thus lower, and GenCo net earnings are also lower because the decline in GenCo avoidable costs is more than offset by the decline in GenCo revenues. Similarly, ISO net surplus is lower because the decline in GenCo revenues is more than offset by the decline in LSE payments.

Next consider the day-1000 data for the learning case. For R=0.0 the mean outcomes for LSE payments, GenCo revenues, and ISO daily net surplus under learning are all substantially higher than their corresponding values under no learning. As R increases, however, the mean outcomes for LSE payments, GenCo revenues, and ISO net surplus under learning eventually drop below their corresponding values under no learning. For mean GenCo revenues the switch point is at R=1.0, whereas for mean LSE payments and mean ISO net surplus the switch point is at R=0.6.

The explanation for these switch points can be deduced from the detailed LMP and total demand findings for the no-learning and learning cases presented in Li, H. , Sun, J. and Tesfatsion, L. (2009). 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 LSE payments, GenCo revenues, and ISO 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, 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 end result is that mean LSE payments, mean GenCo revenues, and mean ISO 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 GenCo net earnings data provided in Table 5.3, the most expensive GenCo 4 is at the greatest disadvantage in this competitive process while the least expensive GenCo 5 is most advantaged.

5.5 30-Bus Benchmark Case

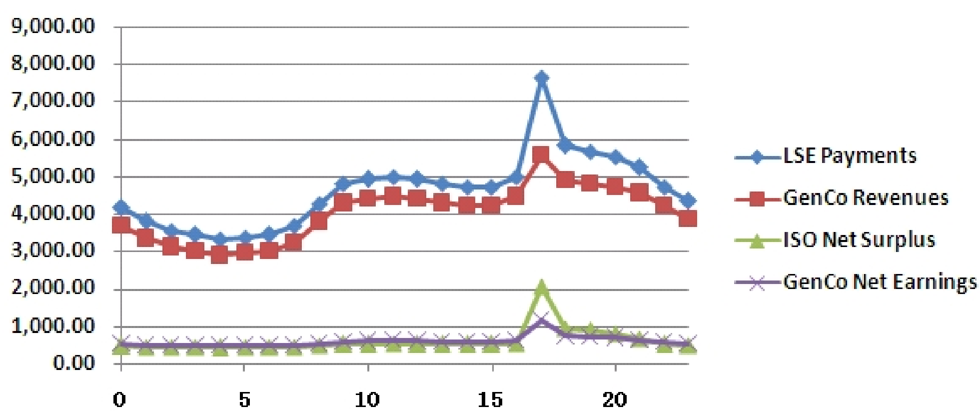

Figure 5.4 LSE payments, GenCo revenues, ISO net surplus, and GenCo net earnings during a typical 24-hour day D for the benchmark dynamic 30-bus test case.

Fig. 5.4 presents benchmark-case hourly financial flows during a typical day D for 30-bus test case. Note that LSE payments are persistently higher than GenCo revenues, particularly during the peak-demand hour 17. Consequently, ISO net surplus is persistently positive with a spike during hour 17.

Table 5.4 shows hourly GenCo net earnings during a typical 24-hour day D for the benchmark dynamic 30-bus test case.

Table 5.4 Hourly GenCo net earnings during a typical 24-hour day D for the benchmark dynamic 30-bus test case.

Hour	GenCo 1	GenCo 2	GenCo 3	GenCo 4	GenCo 5	GenCo 6	GenCo 7	GenCo 8	GenCo 9
00	16.56	3.27	240.85	248.88	0.00	0.00	10.24	0.00	0.00
01	16.56	2.91	235.56	239.10	0.00	0.00	3.48	0.00	0.00
02	16.56	2.68	232.09	232.69	0.00	0.00	0.99	0.00	0.00
03	16.56	2.57	230.31	229.40	0.00	0.00	0.30	0.00	0.00
04	16.56	2.46	228.59	226.23	0.00	0.00	0.01	0.00	0.00
05	16.56	2.51	229.46	227.84	0.00	0.00	0.11	0.00	0.00
06	16.56	2.57	230.31	229.40	0.00	0.00	0.30	0.00	0.00
07	16.56	2.79	233.83	235.91	0.00	0.00	2.05	0.00	0.00
08	16.56	3.39	242.58	252.08	0.00	0.00	13.21	0.00	0.00
09	16.56	3.95	253.41	272.61	0.00	0.16	39.59	0.00	0.00
10	16.56	4.12	255.81	277.07	0.00	0.85	45.74	0.00	0.00
11	16.56	4.21	257.02	279.32	0.00	1.39	48.84	0.00	0.00
12	16.56	4.12	255.81	277.07	0.00	0.85	45.74	0.00	0.00
13	16.56	3.95	253.41	272.61	0.00	0.16	39.59	0.00	0.00
14	16.56	3.86	252.22	270.41	0.00	0.02	36.55	0.00	0.00
15	16.56	3.86	252.22	270.41	0.00	0.02	36.55	0.00	0.00
16	16.56	4.21	257.02	279.32	0.00	1.39	48.84	0.00	0.00
17	16.08	5.89	442.43	647.01	0.00	43.52	5.65	0.00	0.00
18	16.56	4.81	306.29	376.69	0.00	17.66	12.86	0.00	0.00
19	16.56	4.68	298.41	361.16	0.00	12.82	15.20	0.00	0.00
20	16.56	4.56	290.28	345.15	0.00	8.64	17.83	0.00	0.00
21	16.56	4.32	274.04	313.12	0.00	2.74	23.70	0.00	0.00
22	16.56	3.86	252.19	270.35	0.00	0.02	36.48	0.00	0.00
23	16.56	3.46	243.46	253.70	0.00	0.00	14.87	0.00	0.00
Total	396.96	88.99	6247.60	6887.50	0.00	90.26	498.70	0.00	0.00

5.6 30-Bus Case with Learning

Fig. 5.5 presents Comparison of no GenCo learning and with GenCo learning case (day 500) for mean outcomes for LSE payments, GenCo revenues, ISO net surplus, and GenCo net earnings during day 500 for 30-bus test case when when demand is 100% fixed demand. Note that with GenCo learning, every item is much higher than no GenCo learning case.

Due to space limitations, we report only a sampling of results for the 30-bus test case with R=0.0 (100% fixed demand), both with and without GenCo learning.

As seen in Table 5.5, for the no-learning case the typical ISO daily net surplus collection is $28,588, and for the learning case the mean ISO net surplus collection on day 500 is $53,868.30, nearly double the amount for the no-learning case. This increase in ISO net surplus under learning is qualitatively similar to the findings for the 5-bus test case. However, the size of this increase under learning (an

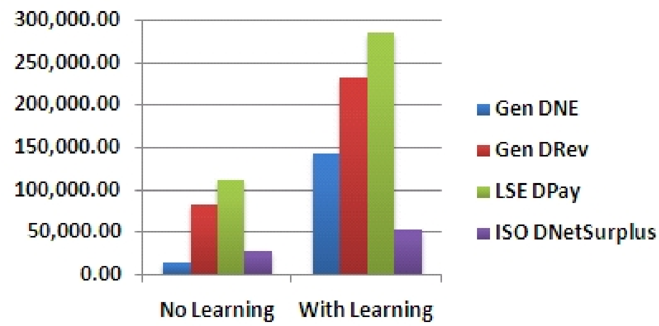

Figure 5.5 Comparison of no GenCo learning and with GenCo learning case (day 500) for mean outcomes for LSE payments, GenCo revenues, ISO net surplus, and GenCo net earnings for the benchmark dynamic 30-bus test case when demand is R=0.0 (100% fixed)

Table 5.5 Mean outcomes (with standard deviations) for GenCo net earnings, GenCo revenues, LSE payments, and ISO net surplus on day 500 for the benchmark dynamic 30-bus test with R=0.0 (100% fixed demand), both with and without GenCo learning.

	No GenCo Learning	With GenCo Learning
Total GenCo DNE	14,210.01	142,866.39 (67,530.19)
Total GenCo DRev	82,265.01	231,663.41 (68,839.94)
Total LSE DPay	110,853.01	285,531.70 (74,214.33)
ISO DNetSurplus	28,588.00	53,868.30 (32,322.68)

approximate doubling) is not as large as for the 5-bus test case, a reflection of the increased rivalry among the more numerous GenCos in the 30-bus test case that results in a more difficult learning environment and less economic capacity withholding.

5.7 Comparisons of Empirical ISO/RTO Day-Ahead Market Net Surplus

Comparisons of empirical ISO/RTO day-ahead market net surplus (congestion rent) collections with our simulation findings are not straightforward. For example, the “congestion costs,” “congestion charges,” and “net congestion revenues” data presented in the ISO/RTO state-of-the-market and market monitoring reports such as PJM (2009), CAISO (2009), ISO-NE (2009), MISO (2008) are in highly aggregated dollar form, and in some cases the explanation of calculation procedures is not fully given. In addition, these dollar amounts should ideally be normalized in some consistent fashion across ISO/RTOs to correct for the size differences among ISO/RTO market footprints. It does not appear to be possible to do this in a completely satisfactory manner based on current publicly released ISO/RTO data.

The most complete reporting appears to be provided by PJM. In PJM (2009) Table 2-47, p. 48, the 2008 average cleared fixed plus price-sensitive load in the PJM day-ahead market is given as 76,961 MWh whereas the average cleared price-sensitive load in the PJM day-ahead market is given as only 1,846 MWh. This implies an R-ratio equal to R=0.02, which is close to R=0.0 (100% fixed demand). In Section 7 (pp. 342) the "total congestion cost" is said to "represent the overall charge or credit to a zone," which we interpret to mean the difference between load payments to the ISO and generation credits (revenues) received from the ISO, i.e., ISO net surplus. On page 339 the 2008 day-ahead congestion costs for PJM are given as \$2.66 billion. This is approximately 7% of 2008 total PJM billings, listed as \$34.3 billion.

In comparison, consider the simulation findings reported in Table 5.5 for our 30-bus test case with R=0.0. For the no-learning case, the ratio of ISO daily net surplus to *total daily billings*, measured as [GenCo daily revenues + LSE daily payments], is about 15%. For the learning case, the ratio of ISO daily net surplus to total daily billings is about 10%. The latter learning-case findings are in line with the 7% empirical findings for PJM, particularly since total PJM billings include settlements for black start, ancillary services, reactive services, FTR payouts, ARR credits, and transmission charges in addition to settlements for load and generation day-ahead energy trades.

In CAISO (2009) Section 5, p. 5.3(103), the CAISO "inter-zonal congestion charges" for the day-ahead and hour-ahead markets in year 2008 are listed as \$176 million. However, these charges are calculated as the product of the "congestion price" (branch shadow price) and the power flow on the branches connecting variously specified zones rather than the difference between load payments and generator revenues as used in this study. In ISO-NE (2009) Section 3.4, p. 70, the combined 2008 Net Congestion Revenue for the ISO-NE real-time and day-ahead markets is listed as \$121 million, where "net congestion revenue" is calculated as the product of branch flows and branch shadow prices. In Figure 46 (p. 68) of MISO (2008) Section V, the 2007 congestion cost for the MISO day-ahead market is listed as approximately \$633 million and the 2007 total congestion cost combining both day-ahead and real-time markets is listed as approximately \$713 million. “Congestion cost” for the MISO is defined (p. 67) as “the difference in LMP prices between the locations multiplied by the amount of the transfer.”

5.8 Concluding Remarks

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

The findings presented in Section 5.3, Section 5.4, Section 5.5 and Section 5.6 demonstrate that ISO net surplus can be substantial in the presence of grid congestion. ISO net surplus was found to be particularly large when demand was predominately fixed (insensitive to price) and GenCos were able to learn over time to strategically report supply offers with higher-than-true marginal costs. Congestion and strategic reporting (encouraged by the presence of fixed demand) reduce market efficiency to the extent they result in the dispatch of more expensive generation in place of cheaper generation (out-of-merit-order dispatch) and/or a failure to meet serviceable price-sensitive demand; cf. Fig. 3.4.

In Litvinov, E. , Zhao, F. and Zheng, T. (2009) the authors note that “under the current ISO practice, the congestion revenue gathered by the ISO is largely returned to the load and transmission owners, resembling the government surplus as part of the social surplus in welfare economics.” However, welfare economists do not assert the *unqualified* desirability of assigning government tax revenue dollars the same weight as private trader net surplus dollars in market objective functions, as is done in the ISO day-ahead market objective function (3.2). An equal weighting is especially problematic if the government redistributes tax revenues to third parties with high entry barriers (e.g., transmission owners) and this redistribution effectively rewards these third parties for maintaining social costs (e.g., congestion) that the government hopes to alleviate.

A key social welfare issue for the ISO day-ahead market objective function (3.2) is whether a dollar flowing to the ISO is properly treated as having the same social benefit as a dollar flowing to a private energy buyer or seller. The answer surely depends on social opportunity costs, i.e., on the net social benefits of alternative uses to which such dollars could be put. This issue would seem to require further serious study.

For example, in some regions the ISO net surplus is used in part to encourage new transmission investment through the subsidization of financial transmission rights (FTRs) for those who invest in new transmission capacity. This practice could lead to a reduction of ISO net surplus to the extent that congestion is alleviated.[1] Nevertheless, transmission investment needs can arise for reasons other than congestion (e.g., the need to reach distributed energy resources), and congestion might better be alleviated by more local generation rather than by more transmission capacity. Consequently, the administrative subsidization of transmission investment through the distribution of ISO net surplus to FTR holders could inadvertently create additional sources of social inefficiency.

In addition, ISO net surplus is also used in part to compensate load. Here, however, it is important to keep in mind the intended market efficiency rationale for LMP pricing in relation to load: namely, to encourage load to locate where it can be serviced most cheaply. To the extent that redistribution of ISO net surplus to load dampens this incentive, it can be an unintended source of social inefficiency viewed from a dynamic vantage point.

Another troubling issue also arises. As seen in Section 5.3 and Section 5.4, ISO and GenCo net surplus collections dramatically increase when the price-sensitivity of demand is low and the GenCos have learning capabilities enabling them to exercise capacity withholding. On the other hand, LSE payments also dramatically increase. This would appear to give pure LSEs (those without generation ownership) an incentive to support congestion reduction measures, increased price-sensitivity of demand, and increased oversight to curtail GenCo capacity withholding.

However, if LSE payments in the wholesale power market are fully reimbursed through the receipt of regulated rates for sale of downstream (retail) electric power, and the LSEs are able to secure timely increases in these regulated rates in step with increases in wholesale power prices, then no direct participant in the wholesale power market suffers a loss of net surplus when LMPs increase due to con-

[1] The extent to which ISO net surplus payouts to FTR holders have actually resulted in new transmission investment is unclear. For example, in the CAISO report CASIO (2004) p. ES-3 reaches the following conclusion: "...the reality has been that the LMP differences have not provided enough incentives to upgrade key facilities even after many types of FTRs and ARRs are provided."

gestion, fixed demand, and/or GenCo capacity withholding. Rather, to the extent that retail consumers are protected against high wholesale prices by locally regulated retail prices, losses in net surplus are in effect borne by taxpayers. Moreover, barriers to entry into transmission, generation, and load servicing could then lead to the persistence over time of socially inefficient *rents*, i.e., net surplus collections in excess of the amounts needed to maintain resources in their current productive uses.

Power market researchers recognize that an important goal of market design is to ensure the structural alignment of participant objectives with socially desirable outcomes, thus reducing the need for oversight of participant behaviors, see Sauma, E. and Oren, S. (2009). The main conclusion drawn from the findings in this study is that net surplus collections are not structurally well-aligned with social efficiency objectives in ISO-managed wholesale power markets operating under LMP. The immediate practical import is the desirability of encouraging more transparent public reporting and oversight of net surplus collections and uses to maintain the confidence of both market participants and the public at large.

CHAPTER 6. GENCO CAPACITY WITHHOLDING STUDY

6.1 Introduction

This chapter uses a dynamic 5-bus test case implemented via the AMES Wholesale Power Market Test Bed to investigate strategic capacity withholding by generation companies (GenCos) in restructured wholesale power markets. The strategic behaviors of the GenCos are simulated by means of the VRE stochastic reinforcement learning algorithm introduced and explained in earlier chapters. The learning GenCos attempt to improve their net earnings over time by strategic selection of their reported supply offers. This strategic selection can involve *economic* capacity withholding (reporting of higher-than-true marginal costs), *physical* capacity withholding (reporting of lower-than-true maximum operating capacity), or combinations of the two.

Before undertaking these capacity withholding experiments, preliminary steps were first taken to help ensure that the GenCos' learning methods were calibrated to their decision environment. Specifically, the dynamic 5-bus test case with 100% fixed demand was used to conduct experiments involving extensive parameter sweeps for various key VRE learning parameters. These learning calibrations were done twice, once for economic capacity withholding and once for physical capacity withholding. For each case, the resulting average GenCo net earnings were determined and displayed using heat map visualizations. These visualizations were then used to determine the best ("sweet spot") learning parameter settings for the learning GenCos in all subsequent experiments.

Section 6.3 explains the experimental design used to explore GenCo capacity withholding under *economic* capacity withholding and *physical* capacity withholding separately when GenCos have sweet-spot VRE learning capabilities. Section 6.3.2 and Section 6.3.3 explain more carefully how these sweet-spot VRE learning capabilities were determined. Experimental findings are reported in Sections 6.2 through 6.6. Concluding remarks are given in Section 6.7.

6.2 5-Bus Benchmark Case: No Economic or Physical Capacity Withholding

For later reference below, recall that Table 5.1 in Chapter 5 (Section 5.3) shows hourly GenCo net earnings without GenCo learning for a typical benchmark-case day. GenCos 1 and 2 have relatively small net earnings in every hour, GenCo 3 located at the load-pocket bus 3 has relatively large net earnings in every hour, and GenCo 4 only has net earnings in the peak hour 17. GenCo 5, a base-load generator with large capacity and low marginal cost, has relatively flat net earnings across all hours.

6.3 Experimental Design and GenCo Learning Calibration

6.3.1 Experimental Design

All economic and physical capacity withholding experiments reported below are based on the benchmark 5-bus test case detailed in Chapter 4 (Section 4.2). GenCo net earnings are used to evaluate GenCo learning results. For simplicity, LSE demand bids are assumed to be in the form of 100% fixed demands (no price sensitivity).

The treatment factor for *economic* capacity withholding experiments is the extent to which each learning GenCo can exercise economic capacity withholding by reporting higher-than-true marginal costs in its supply offers. The treatment factor for *physical* capacity withholding experiments is the extent to which each learning GenCo can exercise physical capacity withholding by reporting lower-than-true maximum operating capacities in its supply offers.

When GenCos have learning capabilities, random effects are present in their supply offer selections. To control for these random effects, thirty seed values were generated using the standard Java class "random"; see Table B.2 for a listing of these seed values. For each learning treatment these thirty seed are used to implement thirty distinct runs, each 1000 simulated days in length.

6.3.2 GenCo Learning Calibration for Economic Capacity Withholding

In the economic capacity-withholding experiments reported below, each learning GenCo i 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 function (3.6) in pursuit of increased net earnings. The action domain AD_i for each

GenCo i is constructed as in Li, H. , Sun, J. and Tesfatsion, L. (2008) to include 100 candidate supply offer choices, each with a distinct setting for these ordinate and slope parameters.

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 Pentapalli, M. (2008). Finally, as explained and graphically depicted in Chapter 4 (Section 4.4), the particular sweet-spot settings (α, β) = (1,100) are used for each GenCo i's α and β learning parameters, which in turn imply sweet-spot settings for each GenCo i's VRE-RL initial propensity and temperature parameters $q(1)_i$ and T_i.

6.3.3 GenCo Learning Calibration for Physical Capacity Withholding

In the physical capacity-withholding experiments reported below, each learning GenCo i makes daily use of the VRE-RL algorithm to adjust the value of its reported maximum operating capacity Cap_i^{RU} in pursuit of increased net earnings. In particular, as clarified more carefully in Section 6.5, experiments are conducted for a range of settings for each GenCo's *Minimum Possible Reported Max Capacity (MPRMCap)*, as follows: 95% to 99% for GenCo 3; 75% to 95% for GenCo 1; and 70% to 95% for GenCo 5. For each different MPRMCap setting, the action domain AD_i for each learning GenCo i is constructed to include 30 candidate supply offer choices, each with a distinct setting for Cap_i^{RU} lying between MPRMCap and GenCo i's true maximum operating capacity Cap_i^U.

The VRE-RL recency and experimentation parameters r_i and e_i for each GenCo i are again fixed at 0.04 and 0.96, respectively. As indicated in Table 6.1, 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, or equivalently, for each GenCo i's α and β values.

Figure 6.1 depicts physical capacity withholding experimental findings for mean GenCo 3 net earnings on day 1000 under alternative (α,β) settings assumed to be set commonly across all GenCos. GenCo 3 is chosen for this calibration because in most hours this large supplier turns out to be pivotal (i.e., essential for meeting fixed demand).

Three interesting observations can be made. First, the settings for (α,β) substantially affect mean GenCo 3 net earnings in physical capacity learning experiments. Second, the sweet-spot (α,β) combi-

Table 6.1 GenCo action domain and learning parameter settings for physical capacity withholding experiments

Action Domain Parameters

GenCo i	$M1_i$	$M2_i$	$M3_i$	$RIMax_i^L$	$RIMax_i^U$	$RIMin_i^C$	SS_i
1	1	1	(1, 30)	0.75	0.75	(0.75, 0.80, 0.85, 0.90, 0.95)	0.001
2	1	1	1	0.75	0.75	1.00	0.001
3	1	1	(1, 30)	0.75	0.75	(0.95, 0.96, 0.97, 0.98, 0.99)	0.001
4	1	1	1	0.75	0.75	1.00	0.001
5	1	1	(1, 30)	0.75	0.75	(0.70, 0.75, 0.80, 0.85, 0.90, 0.95)	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	6,485.29	(1, 1/2, 1/4, 1/10, 1/24, 1/50, , 1/100)	(100, 50, 10, 2, 1, 1/2)
2	0.04	0.96	110.79	(1, 1/2, 1/4, 1/10, 1/24, 1/50, , 1/100)	(100, 50, 10, 2, 1, 1/2)
3	0.04	0.96	233,428.08	(1, 1/2, 1/4, 1/10, 1/24, 1/50, , 1/100)	(100, 50, 10, 2, 1, 1/2)
4	0.04	0.96	592.79	(1, 1/2, 1/4, 1/10, 1/24, 1/50, , 1/100)	(100, 50, 10, 2, 1, 1/2)
5	0.04	0.96	142,781.67	(1, 1/2, 1/4, 1/10, 1/24, 1/50, , 1/100)	(100, 50, 10, 2, 1, 1/2)

nations associated with the highest mean GenCo 3 net earnings lie within a small area spanning combinations from (1/24,100) to (1/24,50). The particular sweet-spot setting (α, β) = (1/24,100) is used for the learning GenCos in all *physical* capacity withholding experiments reported in the remainder of this chapter. Third, comparing the physical capacity withholding learning outcomes in Figure 6.1 with the economic capacity withholding learning outcomes in Figure 4.7, it is seen that the sweet-spot region for GenCo 3's (α, β) learning parameters strongly depends on the particular learning environment.

6.4 5-Bus Economic Capacity Withholding Experiments

In this section, two types of experiments are studied. The first type of experiment tests the extent to which a *single* learning GenCo can learn to achieve higher net earnings through economic capacity withholding when all other GenCos report their true cost and capacity attributes to the ISO. The second type of experiment tests the extent to which *two* learning GenCos can learn over time to achieve higher net earnings through economic capacity withholding when all other GenCos report their true cost and capacity attributes to the ISO. Of particular interest will be the extent to which the two learning GenCos collaborate with each other in determining their economic capacity withholding strategies.

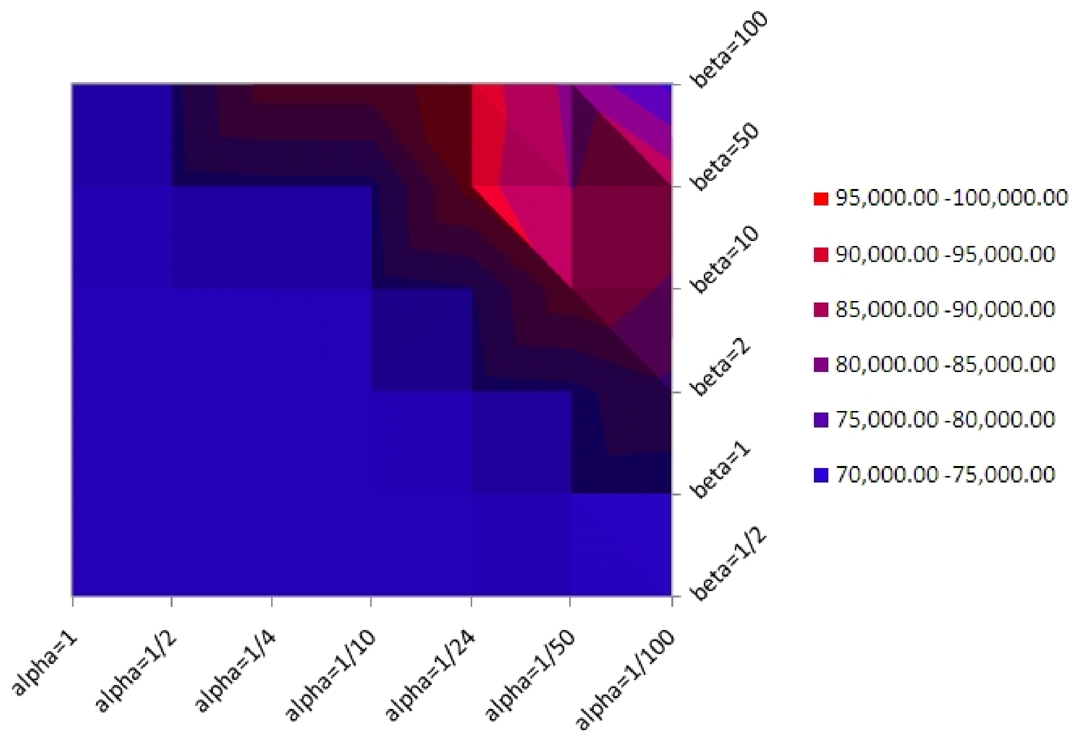

Figure 6.1 A 2D depiction of mean outcomes on day 1000 for GenCo 3's daily net earnings for the benchmark dynamic 5-bus test case extended to include physical capacity withholding learning capability (with MPRM-Cap=95%) for GenCo 3. Results are shown for a range of values for the learning parameters (α, β) commonly set across the GenCo 3.

6.4.1 Economic Capacity Withholding by One GenCo

GenCo 3 is selected for this first type of experiment because of the critical role it plays in the determination of LMPs. This critical role results for three reasons: (i) GenCo 3 has a large maximum operating capacity; (ii) GenCo 3 is a pivotal supplier during peak hours; and (iii) GenCo 3's true marginal costs of production are relatively high.

As discussed above, GenCo 3's learning parameters are set at the calibrated sweet-spot levels (α, β) = (1,100). Also, GenCo 3's action domain parameter values are set at $M1_3$=10, $M2_3$=10, $M3_3$=1, implying that GenCo 3's action domain consists of 100 possible marginal cost function choices with varying ordinate and slope values. All other GenCos are assumed to be non-learners that report their true cost and capacity attributes to the ISO. Consequently, each other GenCo i's action domain parameter values are set at $M1_i$=1, $M2_i$=1, and $M3_i$=1, implying that its action domain contains only one element: namely, GenCo i's true marginal cost function defined over its true operating capacity interval.

Table 6.2 Mean outcomes (with standard deviations) on day 1000 for GenCo daily net earnings and reported supply offers when GenCo 3 can learn to exercise economic capacity withholding.

	No GenCo Learning	With GenCo 3 Learning
GenCo 1 $\overline{\text{DNE}}$	1,556.41	0.00 (0.00)
GenCo 2 $\overline{\text{DNE}}$	26.58	0.00 (0.00)
GenCo 3 $\overline{\text{DNE}}$	56,016.09	1,699,368.20 (400,430.50)
GenCo 4 $\overline{\text{DNE}}$	142.27	253,468.03 (72,720.50)
GenCo 5 $\overline{\text{DNE}}$	34,266.94	33,097.68 (0.00)
GenCo 3 $\overline{a^R}$	25.0000 (true value)	91.0000 (14.5270)
GenCo 3 $\overline{b^R}$	0.0100 (true value)	0.2334 (0.0644)

Table 6.2 reports experimental findings both for the benchmark no-learning case and for the case in which GenCo 3 can learn to exercise economic capacity withholding. Clearly, under learning, GenCo 3 learns to report a much higher-than-true marginal cost function that results in a substantial increase in its net earnings. Interestingly, GenCo 4 substantially benefits along with GenCo 3, even though GenCo 4 has no learning capabilities and reports its true cost and capacity attributes. On the other hand, the mean net earnings of GenCo 1 and GenCo 2 are reduced to zero.

The reason for this is as follows. The branch from Bus 1 to Bus 2 is persistently congested whether or not GenCo 3 has learning capabilities. However, under learning, GenCo 3's high reported marginal costs during the peak hour 17 results in the higher dispatch of GenCo 4 (to max capacity) and also in the higher dispatch of GenCo 5 in order to meet demand in the load pocket surrounding GenCo 3 at Bus 3. GenCo 1 and GenCo 2 have to be backed down to 0 in order to permit GenCo 5 to be called up to service this demand without overloading the branch from Bus 1 to Bus 2.

6.4.2 Economic Capacity Withholding by Two GenCos

Two different pairs of learning GenCos are examined for this second type of experiment: Case (1) GenCo 1 and GenCo 3; and Case (2) GenCo 3 and GenCo 5. The reason for these choices is as follows.

For case (1), GenCo 1 is a small GenCo having low true marginal cost whereas GenCo 3 is a pivotal supplier during peak hours with relatively high true marginal costs. Can GenCo 1 learn to "free ride" on the market power exercised by GenCo 3 in order to improve its net earnings? For case (2), GenCo 3 and GenCo 5 both have large maximum operating capacities, but GenCo 5 has relatively lower marginal costs. Can GenCo 5 learn to undercut GenCo 3's supply offers when GenCo 3 reports aggressively high supply offers, thus raising its net earnings?

Table 6.3 reports mean outcomes for Case (1), in which GenCo 1 and GenCo 3 are the only learners. As indicated, GenCo 3 learns to report much higher-than-true marginal cost functions and attains much higher daily net earnings compared to the benchmark no-learning case. GenCo 1 also learns to report higher-than-true marginal cost functions, yet GenCo 1 does not manage to attain higher net earnings.

Interestingly, the net earnings and reported marginal cost results presented in Table 6.3 for the case in which GenCo 1 and GenCo 3 are both learners are similar to the corresponding results reported in Table 6.2 for the case in which only GenCo 3 is a learner. The reason for this is partly explained by the findings earlier discussed in Section 6.4.1 and Section 6.2. Due to the persistent congestion on the branch from Bus 1 to Bus 2, and to GenCo 3's relatively large operating capacity, GenCo 3 is a pivotal supplier in most hours, meaning that it's capacity is needed to meet fixed demand. On the other hand, GenCo 1 is a relatively small unit located on the "wrong" side of the congested branch and typically fails to be dispatched at any positive level. The result is that GenCo 3's reported supply offers have a much greater effect on dispatch results. GenCo 3 learns to take advantage of this situation by raising its reported marginal costs, resulting in an increase in the LMP at its Bus 3. In contrast, despite its learning capabilities, GenCo 1's supply offers are essentially irrelevant for the determination of price and dispatch levels, as well as for the determination of GenCo 3's reported supply offers.

Table 6.4 depicts mean outcomes for Case (2), in which GenCo 3 and GenCo 5 are the only learners. As indicated, both GenCo 3 and GenCo 5 learn to report much higher-than-true marginal costs and both attain substantially higher daily net earnings compared to the benchmark no-learning case. GenCo 3's

Table 6.3 Mean outcomes (with standard deviations) on day 1000 for GenCo net earnings and reported supply offers when both GenCo 1 and GenCo 3 can learn to exercise economic capacity withholding.

	No GenCo Learning	With GenCo 1, 3 Learning
GenCo 1 $\overline{\text{DNE}}$	1,556.41	0.00 (0.00)
GenCo 2 $\overline{\text{DNE}}$	26.58	0.00 (0.00)
GenCo 3 $\overline{\text{DNE}}$	56,016.09	1,699,368.20 (400,430.50)
GenCo 4 $\overline{\text{DNE}}$	142.27	253,468.03 (72,720.50)
GenCo 5 $\overline{\text{DNE}}$	34,266.94	33,097.68 (0.00)
GenCo 1 $\overline{a^R}$	14.0000	26.7006 (12.8204)
GenCo 1 $\overline{b^R}$	0.0050	0.1363 (0.2016)
GenCo 3 $\overline{a^R}$	25.0000	91.0000 (14.5270)
GenCo 3 $\overline{b^R}$	0.0100	0.2334 (0.0644)

net earnings, in particular, dramatically increase.

The explanation for these findings is as follows. GenCo 5 is essentially a base-load generator with large capacity and low true marginal cost. When both GenCo 3 and GenCo 5 report higher-than-true marginal costs, the branch connecting Bus 1 to Bus 2 becomes persistently congested, constraining the use of the relatively cheaper generation from GenCo 1 and GenCo 2 at Bus 1. GenCo 4, a relatively small unit, is then dispatched at its maximum capacity because its reported marginal costs are actually lower than the reported marginal costs of GenCo 3 and GenCo 5. This leaves GenCo 3 and GenCo 5 as pivotal suppliers. The dispatch of GenCo 5 is constrained by congestion considerations, which acts as a brake on its net earnings. GenCo 3, however, induces no such network constraint in terms of its pivotal status for the load at its own Bus 3. This permits GenCo 3 to raise its reported marginal costs to very high levels without concern for a cut-back in its dispatch, which in turn results in a very high LMP at its load-pocket Bus 3 and in correspondingly high daily net earnings for GenCo 3.

Comparing the Case (2) findings presented in Table 6.4 to the Case (1) findings presented in Table 6.3, it is seen that GenCo 3's daily net earnings are about the same. The implication is that GenCo 3 is not strategically interacting with GenCo 5 in Case (2); it behaves essentially the same way whether or

Table 6.4 Mean outcomes (with standard deviations) on day 1000 for GenCo daily net earnings and reported supply offers when both GenCo 3 and GenCo 5 can learn to exercise economic capacity withholding.

	No GenCo Learning	With GenCo 3, 5 Learning
GenCo 1 $\overline{\text{DNE}}$	1,556.41	20,552.07 (47,427.81)
GenCo 2 $\overline{\text{DNE}}$	26.58	17,516.78 (42,538.13)
GenCo 3 $\overline{\text{DNE}}$	56,016.09	1,689,877.00 (388,472.72)
GenCo 4 $\overline{\text{DNE}}$	142.27	299,156.86 (93,287.70)
GenCo 5 $\overline{\text{DNE}}$	34,266.94	129,744.81 (96,084.46)
GenCo 3 $\overline{a^R}$	25.0000	92.8333 (12.3654)
GenCo 3 $\overline{b^R}$	0.0100	0.2202 (0.0671)
GenCo 5 $\overline{a^R}$	10.0000	18.6560 (10.3939)
GenCo 5 $\overline{b^R}$	0.0070	0.0238 (0.0286)

not GenCo 5 has learning capabilities. On the other hand, in Case (2) GenCo 5 is able to take advantage of GenCo 3's economic capacity withholding to raise its own reported marginal costs without risking a cut-back in its dispatch, which substantially increases its daily net earnings.

Case (2) also differs from Case (1) in another interesting way. In Case (2), GenCo 5 ends up reporting marginal costs that are higher than the marginal costs of the non-learning GenCos 1 and 2. As a result, GenCo 1 and GenCo 2 located at Bus 1 are now dispatched at positive levels even though the branch connecting Bus 1 to Bus 2 is persistently congested. Consequently, these non-learning GenCos are better off in Case (2) than in Case (1).

6.5 5-Bus Physical Capacity Withholding Experiments

In parallel with subsection 6.4, this subsection considers two types of experiments. The first type of experiment tests the extent to which a *single* learning GenCo can learn to achieve higher net earnings through *physical* capacity withholding when all other GenCos report their true cost and capacity attributes to the ISO. The second type of experiment tests the extent to which *two* learning GenCos can learn over time to achieve higher net earnings through *physical* capacity withholding when all other GenCos report their true cost and capacity attributes to the ISO. As in subsection 6.4, of particular interest will be the extent to which the two learning GenCos collaborate with each other in determining their economic capacity withholding strategies.

The treatment factors for these experiments are the maximum permitted shrinkages for the GenCos' reported maximum capacities to the ISO. The tested ranges for these treatment factors are chosen to avoid supply inadequacy, i.e., the reporting of capacities that are insufficient to meet total fixed demand.

6.5.1 Physical Capacity Withholding by One GenCo

In the experiments presented in this subsection, as in subsection 6.4.1, only GenCo 3 has learning capabilities. However, here GenCo 3's learning is restricted to the ability to exercise *physical* capacity withholding. The only treatment factor is GenCo 3's action-domain parameter MPRMCap, i.e., the *minimum possible reported maximum capacity* that GenCo 3 is able to report to the ISO as a percentage of its true maximum capacity. For example, given MPRMCap = 99%, GenCo 3's reported maximum capacity must be at least 99% of its true maximum capacity. In the experiments presented below, MPRMCap is varied between 95% and 99%. All other GenCos are assumed to report their true costs and capacities to the ISO.

The findings presented in Table 6.5 show that GenCo 3 is able to substantially increase its mean daily net earnings through physical capacity withholding. Indeed, GenCo 3's daily net earnings steadily increase as it increases its physical capacity withholding from 1% to 5% of its true maximum capacity. These increases in daily net earnings are at the expense of GenCo 1 and GenCo 2, who do worse as GenCo 3's withholding increases. On the other hand, GenCo 4 and GenCo 5 experience modest gains in daily net earnings from GenCo 3's withholding.

Figure 6.2 depicts GenCo 3's actual reported maximum capacity versus its *optimal reported capacity (ORCap)* under a range of different MPRMCap settings for GenCo 3. More precisely, for each given MPRMCap setting, ORCap gives the best possible capacity value that GenCo 3 could report to the ISO in the sense that this reporting leads to the highest mean daily net earnings for GenCo 3. Figure 6.2 shows that the mean maximum capacity value that GenCo 3 *learns* to report to the ISO by day 500 is

Table 6.5 Mean outcomes (with standard deviations) on day 500 for GenCo daily net earnings and reported maximum capacity values when GenCo 3 can learn to exercise physical capacity withholding. Results are shown for a range of MPRMCap values for GenCo 3.

GenCo 3 MPRMCap	99%	98%	97%	96%	95%
GenCo 1 $\overline{\text{DNE}}$	1,519.60 (2.34)	1,515.29 (6.82)	1,461.74 (7.91)	1,451.66 (8.30)	1,397.42 (10.55)
GenCo 2 $\overline{\text{DNE}}$	26.36 (0.00)	26.21 (0.33)	24.99 (0.00)	24.72 (0.45)	23.69 (0.25)
GenCo 3 $\overline{\text{DNE}}$	65,532.62 (655.02)	66,389.84 (1,534.96)	78,855.13 (1,884.59)	80,724.79 (1,517.63)	92,830.96 (2,368.11)
GenCo 4 $\overline{\text{DNE}}$	176.51 (7.27)	209.16 (25.67)	300.41 (13.97)	353.49 (38.51)	471.42 (19.94)
GenCo 5 $\overline{\text{DNE}}$	34,576.94 (15.94)	34,530.62 (27.02)	34,839.52 (52.91)	34,815.19 (20.97)	35,121.16 (67.58)
GenCo 3 RepMaxCap (MW)	515.99 (0.77)	512.86 (2.44)	505.42 (1.01)	502.02 (2.40)	495.49 (1.04)
GenCo 3 RepMaxCap/TrueMaxCap (%)	99.23% (0.15%)	98.63% (0.47%)	97.20% (0.19%)	96.54% (0.46%)	95.29% (0.20%)

close to optimal for each MPRMCap setting.

6.5.2 Physical Capacity Withholding by Two GenCos

As in section 6.4.2, learning experiments are conducted for pairs of learning GenCos as follows: Case (3) GenCo 1 and GenCo 3; and Case (4) GenCo 3 and GenCo 5. Here, however, learning is restricted to physical capacity withholding. Choosing the same pairings as in section 6.4.2 permits meaningful comparisons between learning experiments for economic versus physical capacity withholding.

The treatment factors for the Case (3) experiments are the MPRMCap settings for GenCo 1 and GenCo 3. The MPRMCap setting for GenCo 1 is varied from 75% to 95% and the MPRMCap setting for GenCo 3 is varied from 95% to 99%. All non-learning GenCos report their true cost and capacity attributes to the ISO.

As a benchmark of comparison for Case (3), Figure 6.3 presents typical daily net earnings for GenCo 1 and GenCo 3 for the benchmark no-learning case under a range of settings for the maximum capacities for GenCo 1 and GenCo 3. From these findings it can be seen that GenCo 1, a relatively small

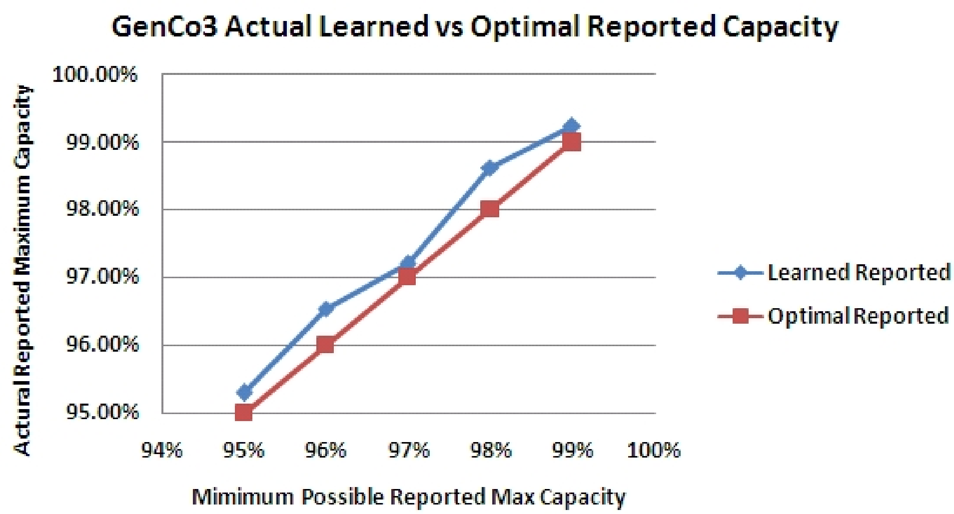

Figure 6.2 Mean outcomes on day 500 for the learned versus optimal values for GenCo 3's reported maximum capacity values when GenCo 3 can learn to exercise physical capacity withholding. Results are shown for a range of different minimum possible reported maximum capacity (MPRMCap) values for GenCo 3.

unit, does best when its maximum capacity is set at 75% of its benchmark true maximum capacity and the maximum capacity of the relatively large GenCo 3 is set at 99% of its benchmark true maximum capacity. In contrast, GenCo 3 does best when its maximum capacity is set at 95% of its benchmark true maximum capacity no matter what value is set for GenCo 1's maximum capacity.

Figure 6.4 presents mean daily net earnings for GenCo 1 and GenCo 3 on day 5000 when both GenCos can learn to exercise physical capacity withholding. From the left-hand side of this figure, it is seen that GenCo 1 attains its highest mean daily net earnings when its MPRMCap value is set at its lowest test level (75%) and the MPRMCap value for GenCo 3 is set at its highest tested level (99%). Comparing the left-hand side of Figure 6.4 to the left-hand side of Figure 6.3, it is also seen that the MPRMCap region over which GenCo 1 attains its highest mean daily net earnings under learning is smaller than the maximum capacity region over which it attains its highest daily net earnings in the benchmark no-learning case.

Interestingly, in parallel with the no-learning findings reported in Figure 6.3, it is seen in the right-hand side of Figure 6.4 that GenCo 3 attains its highest mean daily net earnings when its MPRMCap

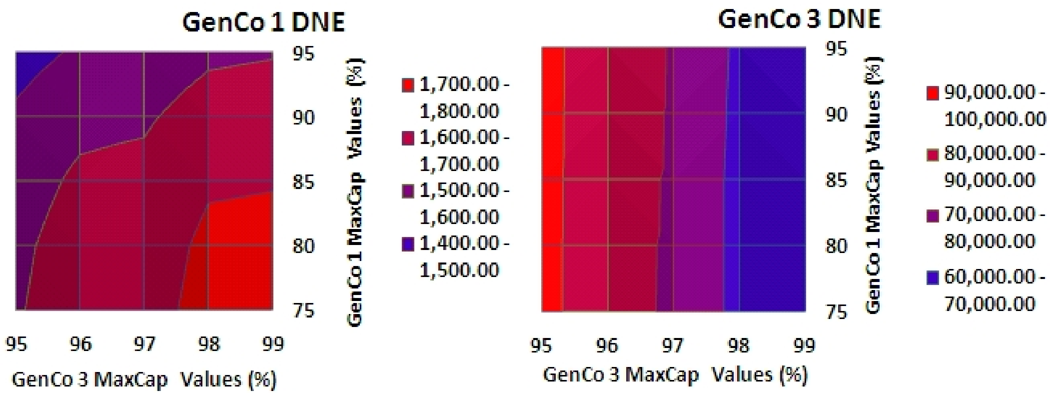

Figure 6.3 Typical daily net earnings for GenCo 1 and GenCo 3 for the benchmark no-learning case under a range of maximum capacity settings for each GenCo. The maximum capacity settings are depicted in percentage form (relative to benchmark true maximum capacities).

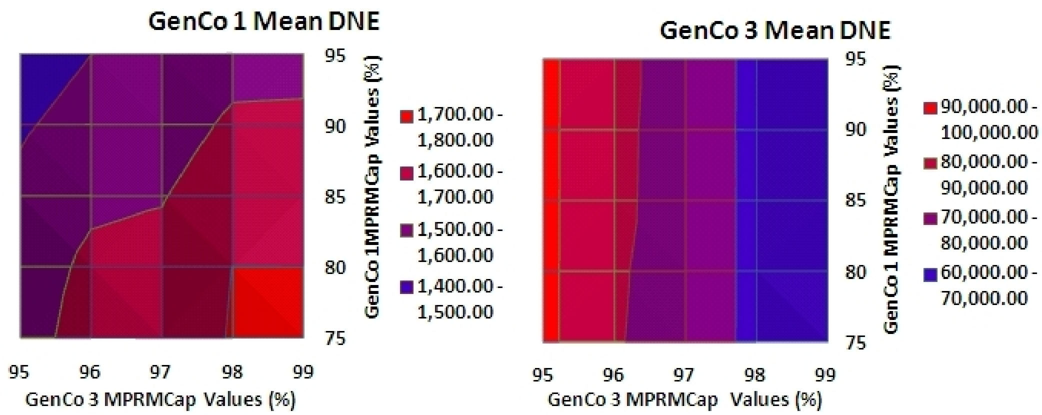

Figure 6.4 Mean daily net earnings on day 500 for GenCo 1 and GenCo 3 when both GenCos can learn to exercise physical capacity withholding. Results are shown for a range of MPRMCap values for each GenCo.

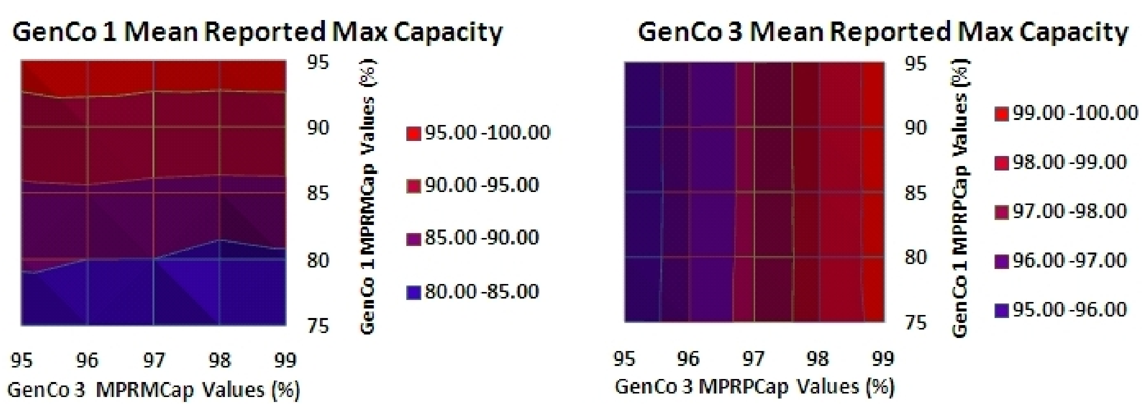

Figure 6.5 Mean reported maximum capacities (as a percentage of benchmark true maximum capacities) on day 500 for GenCo 1 and GenCo 3 when both GenCos can learn to exercise physical capacity withholding. Results are shown for a range of MPRMCap values for each GenCo.

value is set at its lowest tested level 95%. Also, the vertically-striped pattern for GenCo 3's mean daily net earnings indicates that GenCo 1's MPRMCap settings have essentially no effect on the daily net earnings attained by GenCo 3.

Figure 6.5 displays mean reported maximum capacities (as a percentage of benchmark true maximum capacities) for GenCo 1 and GenCo 3 on day 500 when both these GenCos have physical capacity learning capabilities. From the left-hand side of the figure, it is seen that GenCo 1's mean reported maximum capacity is somewhat higher than its MPRMCap setting for each tested pair of MPRMCap settings for GenCo 1 and GenCo 3. Moreover, as indicated by the horizontally-striped pattern in GenCo 1's reported maximum capacity results, GenCo 3's reported maximum capacity choices have essentially no effect on the reported maximum capacity choices made by GenCo 1. As indicated in the right-hand side of the figure, GenCo 3's mean reported maximum capacity is close to its MPRMCap setting for each tested pair of MPRMCap settings for GenCo 1 and 3. Moreover, as indicated by the vertically-striped pattern in GenCo 3's reported maximum capacity results, GenCo 1's reported maximum capacity choices have essentially no effect on the reported maximum capacity choices made by GenCo 3.

In summary, in the Case (3) physical capacity learning experiments involving GenCo 1 and GenCo 3, the smaller GenCo 1 is able to attain higher net earnings by essentially free riding on the strategic physical capacity reporting of the larger GenCo 3. In contrast, GenCo 3's reported maximum capacity choices are essentially uncorrelated with the reported maximum capacity choices of GenCo 1.

For the Case (4) physical capacity learning experiments involving GenCo 3 and GenCo 5, the treatment factors are the MPRMCap settings for GenCos 3 and GenCo 5. The MPRMCap setting for GenCo 3 is varied from 95% to 99% and the MPRMCap setting for GenCo 5 is varied from 70% to 95%. All non-learning GenCos report their true cost and capacity attributes to the ISO.

Figure 6.6 shows typical daily net earnings for GenCo 3 and GenCo 5 for the benchmark no-learning case under a range of different settings for their maximum capacities expressed as a percentage of their benchmark true maximum capacities. From the left-hand side of the figure it is seen that GenCo 3 attains its highest daily net earnings at the lowest maximum capacity setting (95%), regardless of the maximum capacity setting for GenCo 5. On the other hand, from the right-hand side of the figure it is seen that GenCo 5 attains its highest daily net earnings when its maximum capacity is set at the lowest tested level (70%) while at the same time the maximum capacity for GenCo 3 is set at its lowest tested level (95%). Thus, GenCo 5's daily net earnings are affected by the maximum capacity setting for GenCo 3.

Figure 6.7 shows GenCo 3 and GenCo 5 mean daily net earnings on day 500 when both GenCos have physical capacity learning capabilities. Results are shown for a range of MPRMCap values for each GenCo. From the left-hand side of the figure it is seen that GenCo 3 attains its highest daily net earnings when its MPRMCap value is set to the lowest tested level (95%), regardless of the MPRMCap setting for GenCo 5. In contrast, from the right-hand side of the figure it is seen that GenCo 5 attains its highest daily net earnings when its MPRMCap value is set to the lowest tested level (70%) while at the same time the MPRMCap value for GenCo 3 is set at its lowest tested level (95%). Consequently, in similarity to the no-learning case, GenCo 5's daily net earnings under learning are affected by the MPRMCap value set for GenCo 3.

Figure 6.8 shows GenCo 3 and GenCo 5 mean reported maximum capacities as a percentage of their benchmark true maximum capacities when both GenCos have physical capacity learning capabilities. Results are reported for a range of MPRMCap values for each GenCo. From the left-hand side of the figure it is seen that GenCo 3's mean reported maximum capacity is close to its MPRMCap value for each tested combination of MPRMCap settings for GenCo 3 and GenCo 5. Also, the horizontally-striped pattern of the results indicates that GenCo 5's reported maximum capacities have very little effect on GenCo 3's reported maximum capacities. The right-hand side of the figure shows that GenCo 5's mean reported maximum capacity is higher than its MPRMCap value for each tested combination of MPRMCap settings for GenCo 3 and GenCo 5. Also, GenCo 5's reported maximum capacities are weakly correlated with GenCo 3's reported maximum capacities.

6.6 5-Bus Combined Economic and Physical Capacity Withholding Experiments

In this section, two types of experiments are studied. The first type of experiment tests the extent to which a *single* learning GenCo can learn to achieve higher net earnings *through economic and/or physical capacity withholding* when all other GenCos report their true cost and capacity attributes to

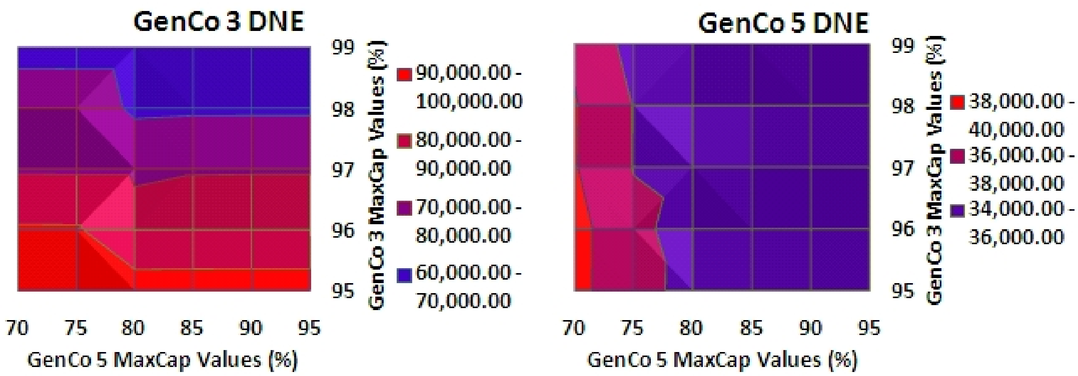

Figure 6.6 Typical daily net earnings for GenCo 3 and GenCo 5 for the benchmark no-learning case under a range of maximum capacity settings for each GenCo. The maximum capacity settings are depicted in percentage form (relative to benchmark true maximum capacities).

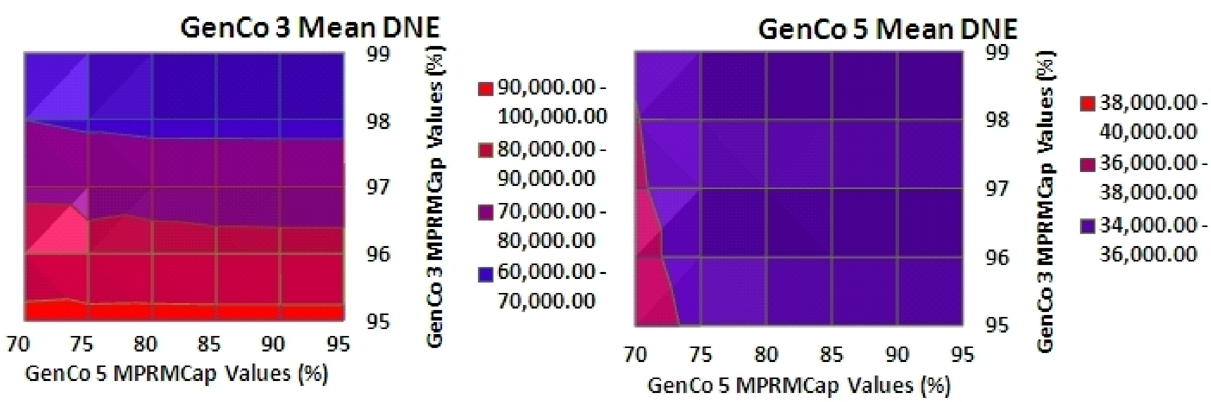

Figure 6.7 Mean daily net earnings on day 500 for GenCo 3 and GenCo 5 when both GenCos can learn to exercise physical capacity withholding. Results are shown for a range of MPRMCap values for each GenCo.

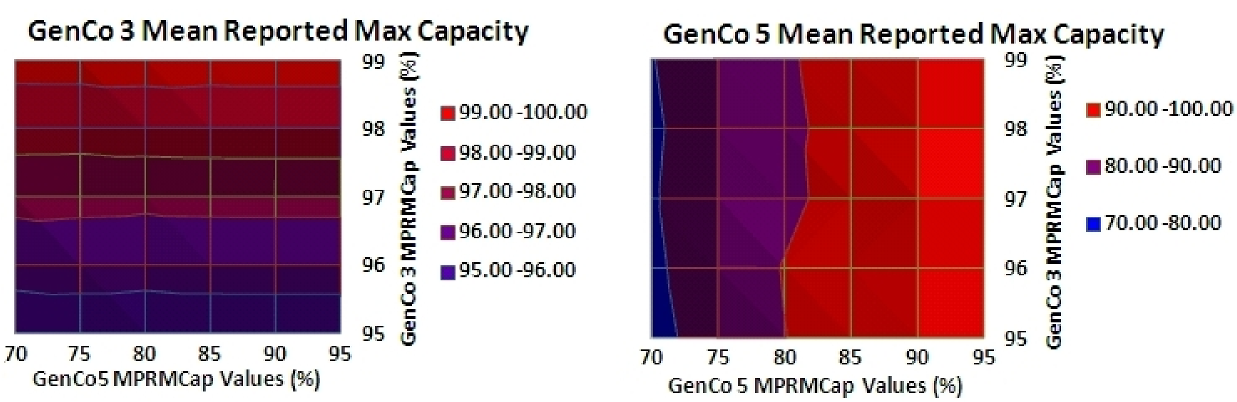

Figure 6.8 Mean reported maximum capacities (as a percentage of benchmark true maximum capacities) on day 500 for GenCo 3 and GenCo 5 when both GenCos can learn to exercise physical capacity withholding. Results are shown for a range of MPRMCap values for each GenCo.

the ISO. The second type of experiment tests the extent to which *two* learning GenCos can learn over time to achieve higher net earnings *through economic and/or physical capacity withholding* when all other GenCos report their true cost and capacity attributes to the ISO.

6.6.1 Combined Economic and Physical Capacity Withholding by One GenCo

For reasons elaborated in earlier subsections, GenCo 3 is selected as the one learning GenCo able to learn to exercise either economic or physical capacity withholding. Of particular interest will be whether one type of withholding dominates the other.

Table 6.6 presents mean outcomes (with standard deviations) for GenCo daily net earnings and GenCo 3 supply offers on Day 1000 under a range of MPRMCap settings for GenCo 3. It is seen that GenCo 3 attains much higher mean daily net earnings than in the benchmark no-learning case. Moreover, the mean daily net earnings for GenCo 3 monotonically increase with increases in the MPRMCap setting for GenCo 3.

However, comparing the findings in Table 6.6 with the benchmark no-learning findings in Table 6.2, it is seen that the increase in GenCo 3's mean net earnings through economic capacity withholding are substantially greater than the increases in its mean net earnings from successively higher physical capacity withholding. Thus, although both forms of capacity withholding add to GenCo 3's net earnings, economic capacity withholding is the primary channel through which it can increase its net earnings.

6.6.2 Combined Economic and Physical Capacity Withholding by Two GenCos

As in sections 6.4.2-6.5.2, learning experiments are conducted for a pair of learning GenCos: namely, GenCo 3 and GenCo 5. Here, however, the two learning GenCos are permitted to engage in both economic and physical capacity withholding. All other GenCos are assumed to report their true cost and capacity attributes to the ISO.

Table 6.7 shows mean outcomes (with standard deviations) on day 1000 for GenCo net earnings and reported supply offers when GenCo 3 and GenCo 5 having learning capabilities for both economic and physical capacity withholding. Comparing these results to the results presented in Table 6.2 for the benchmark no-learning case, it is seen that GenCo 3 and GenCo 5 both attain much higher mean

Table 6.6 Mean outcomes (with standard deviations) on day 1000 for GenCo net earnings and reported supply offers when GenCo 3 can learn to exercise both economic and physical capacity withholding. Results are reported for a range of different MPRMCap values for GenCo 3.

GenCo 3 MPRMCap	99%	98%	97%	96%	95%
GenCo 1 $\overline{\text{DNE}}$	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
GenCo 2 $\overline{\text{DNE}}$	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
GenCo 3 $\overline{\text{DNE}}$	1,973,577.19 (277,338.27)	1,980,160.94 (278,227.89)	1,986,854.87 (279,273.41)	1,993,661.93 (280,482.08)	2,000,585.19 (281,861.29)
GenCo 4 $\overline{\text{DNE}}$	303,449.92 (50,027.46)	304,616.65 (50,186.55)	305,802.88 (50,372.87)	307,009.16 (50,587.68)	308,236.04 (50,832.28)
GenCo 5 $\overline{\text{DNE}}$	33,097.68 (0.00)	33,097.68 (0.00)	33,097.68 (0.00)	33,097.68 (0.00)	33,097.68 (0.00)
GenCo 3 $\overline{\text{a}^R}$ (true a=25.00)	98.57 (7.82)	98.57 (7.82)	98.57 (7.82)	98.57 (7.82)	98.57 (7.82)
GenCo 3 $\overline{\text{b}^R}$ (true b=0.01)	0.27 (0.05)	0.27 (0.05)	0.27 (0.05)	0.27 (0.05)	0.27 (0.05)
GenCo 3 Reported Max Cap (MW) (TrueMaxCap=520)	517.62 (2.13)	515.23 (4.27)	512.85 (6.40)	510.47 (8.54)	508.08 (10.67)
GenCo 3 RepMaxCap/TrueMaxCap (%)	99.54% (0.41%)	99.08% (0.82%)	98.63% (1.23%)	98.17% (1.64%)	97.71% (2.05%)

net earnings under learning. These higher mean net earnings are primarily due to economic capacity withholding, in the form of substantially higher reported ordinate and slope values $\{a^R, b^R\}$ for the GenCos' reported marginal cost functions (3.6). For example, GenCo 3 raises its reported ordinate value a^R dramatically, to almost four times its true value, while GenCo 5 raises its reported ordinate value a^R to over double its true value.

Both GenCos also attain successively higher mean net earnings as the MPRMCap value for GenCo 3 is decreased, permitting GenCo 3 to exercise greater physical capacity withholding. However, these increases in mean net earnings are much smaller in percentage terms than the sharp increases resulting from economic capacity withholding.

In summary, while both economic and physical capacity holding add to the mean net earnings of GenCo 3 and GenCo 5, economic capacity withholding is the primary means through which they attain higher mean net earnings. From Table 6.7, it is also seen that GenCo 3 attains much higher mean net

Table 6.7 Mean outcomes (with standard deviations) on day 1000 for GenCo daily net earnings and reported supply offers when GenCo 3 and GenCo 5 can learn to exercise both economic and physical capacity withholding. Results are reported for a fixed MPRMCap setting of 70% for GenCo 5 and for a range of possible MPRMCap settings for GenCo 3.

GenCo 3 MPRMCap	99%	98%	97%	96%	95%
GenCo 1 $\overline{\text{DNE}}$	52,462.61 (75,157.13)	51,744.44 (75,285.93)	52,380.00 (74,686.83)	62,356.15 (78,612.73)	74,969.74 (80,573.60)
GenCo 2 $\overline{\text{DNE}}$	46,174.07 (67,610.86)	45,521.31 (67,739.12)	46,099.53 (67,179.75)	55,078.56 (70,738.39)	66,413.40 (72,545.32)
GenCo 3 $\overline{\text{DNE}}$	2,030,943.15 (345,177.83)	2,068,246.24 (315,194.58)	2,083,177.71 (339,489.78)	2,139,915.52 (452,137.70)	2,244,762.61 (440,142.52)
GenCo 4 $\overline{\text{DNE}}$	402,351.76 (133,209.50)	408,758.31 (130,103.54)	411,039.92 (130,289.40)	429,387.00 (155,201.83)	461,214.44 (148,768.56)
GenCo 5 $\overline{\text{DNE}}$	192,264.38 (153,958.15)	192,291.10 (153,948.98)	192,642.57 (153,918.35)	215,248.40 (165,392.36)	246,805.08 (169,829.01)
GenCo 3 $\overline{a^R}$ (true a=25.00)	98.57 (7.82)	98.57 (7.82)	100.00 (0.00)	97.14 (10.87)	98.57 (7.82)
GenCo 3 $\overline{b^R}$ (true b=0.01)	0.27 (0.05)	0.28 (0.05)	0.27 (0.06)	0.27 (0.06)	0.27 (0.06)
GenCo 5 $\overline{a^R}$ (true a=10.00)	22.90 (12.92)	22.90 (12.92)	22.90 (12.92)	24.47 (13.37)	26.07 (13.70)
GenCo 5 $\overline{b^R}$ (true b=0.007)	0.05 (0.05)	0.05 (0.05)	0.05 (0.05)	0.05 (0.05)	0.06 (0.05)
GenCo 3 Reported Max Cap (MW) (TrueMaxCap=520)	517.66 (2.03)	515.58 (4.22)	512.72 (6.70)	509.25 (8.07)	505.92 (10.67)
GenCo 5 Reported Max Cap (MW) (TrueMaxCap=600)	507.00 (60.18)	508.50 (59.66)	508.50 (59.66)	511.50 (60.82)	508.50 (61.95)
GenCo 3 RepMaxCap/TrueMaxCap (%)	99.55% (0.39%)	99.15% (0.81%)	98.60% (1.29%)	97.93% (1.55%)	97.29% (2.05%)
GenCo 5 RepMaxCap/TrueMaxCap (%)	84.50% (10.03%)	84.75% (9.94%)	84.75% (9.94%)	85.25% (10.14%)	84.75% (10.33%)

earnings than GenCo 5. The reasons for this are similar to the reasons discussed in section 6.4.2.

6.7 Comparisons of Results

This subsection summarizes the detailed experimental findings for economic and physical capacity withholding reported in previous subsections.

6.7.1 One GenCo Case Comparison

From Table 6.2, Table 6.5 and Table 6.6, it is seen that the relatively large and expensive GenCo 3 (located at the load-pocket Bus 3) attains much higher mean net earnings relative to the benchmark no-learning case when it is able to exercise economic capacity withholding, whether or not it engages in physical capacity withholding. Conversely, although GenCo 3's mean net earnings increase when it exercises only physical capacity withholding, increasingly so for successively smaller settings for its MPRMCap value, these gains are substantially smaller.

6.7.2 Two GenCos Case Comparison

Joint learning experimental findings for GenCo 3 and the relatively small and inexpensive GenCo 1 are reported in Table 6.3, Figure 6.3, Figure 6.4, and Figure 6.5. It is seen that GenCo 3 attains much higher mean net earnings relative to the benchmark no-learning case when using economic capacity withholding, whether or not it engages in physical capacity withholding. Moreover, GenCo 3 has more market power than GenCo 1 because if its pivotal suppler status. GenCo 1 suffers a loss in mean net earnings relative to the benchmark no-learning case when GenCo 3 exercises physical capacity withholding, due to network effects; and GenCo 1 loses out completely (zero dispatch level) when GenCo 3 engages in economic capacity withholding. Conversely, GenCo 3 is largely unaffected by the capacity withholding choices of GenCo 1.

Joint learning experimental findings for GenCo 3 and the relatively large but inexpensive base load GenCo 5 are presented in Table 6.4, Figure 6.6, Figure 6.7, Figure 6.8, and Table 6.7. It is seen that both GenCos attain much higher mean net earnings relative to the benchmark no-learning case when they exercise economic capacity withholding alone or a combination of economic and physical capacity withholding. Conversely, GenCo 3 and GenCo 5 achieve much smaller gains in mean net earnings when they only exercise physical capacity withholding. Also, GenCo 3's favorable load-pocket location gives it more market power than GenCo 5 because it is a pivotal supplier in almost every hour of every run. GenCo 5's best capacity withholding choices are affected by GenCo 3's choices, but GenCo 3's best capacity withholding choices are largely unaffected by the choices of GenCo 5.

Comparing these two joint learning experiments with each other, it is seen that GenCo 5 is in a much better position than GenCo 1 to take advantage of the capacity withholding choices of GenCo 3 to increase its own mean net earnings. Three factors work against GenCo 1 here: (i) the persistent congestion on the branch connecting Bus 1 to Bus 2; (ii) the location of GenCo 1 at Bus 1, semi-islanded away from the load pocket Bus 3; and (iii) the relatively small capacity of GenCo 1.

6.7.3 Conclusion

The experiments reported in this chapter indicate that economic capacity withholding is much more advantageous for GenCos than physical capacity withholding in terms of raising their mean net earnings. However, in these experiments the ISO does not mitigate the exercise of market power by the GenCos in any way. Effective market power mitigation requires monitoring of GenCo reported costs and capacities relative to true. It could be the case that economic capacity withholding is more easily monitored and controlled than physical capacity withholding, because true operating costs can be estimated rather well from publicly available information such as fuel type and fuel prices. Conversely, it could be more difficult to check whether forced outages of generation units are accurately being reported.

Finally, in the present experiments the inexpensive but small GenCo 1 located at Bus 1 is persistently non-marginal, hence its capacity withholding actions have little effect on the mean net earnings of other GenCos. Conversely, when the relatively big GenCos 3 and 5 exercise capacity withholding, GenCo 1 can either win or lose. Specifically, relative to the benchmark no-learning case, GenCo 1: (a) loses big (zero dispatch) when either GenCo 3 alone engages in economic capacity withholding, GenCo 3 and GenCo 1 both engage in economic capacity withholding, or GenCo 3 engages in combined economic and physical capacity withholding; (b) loses modestly when GenCo 3 alone engages in physical capacity withholding; and (c) gains big when GenCo 3 and GenCo 5 both engage in economic capacity withholding, with or without physical capacity withholding.

The key for GenCo 1 is the economic capacity withholding activity of GenCo 5 located at the neighboring Bus 5. If GenCo 5 rather aggressively reports higher-than-true marginal costs, GenCo 1 can appear to be the cheaper GenCo. In this case GenCo 1 is dispatched to full capacity in advance of GenCo 5 and thus manages to sell its generation at the very high price determined by the reported supply offers of the marginal GenCos 3 and/or 5.

CHAPTER 7. CONCLUSIONS AND FUTURE WORK

7.1 Conclusions

U.S. restructured wholesale power markets are large-scale systems encompassing physical constraints, administered rules of operation, and strategic human participants. The complexity of these systems makes it difficult to model and study them using standard analytical and statistical tools.

This dissertation develops and uses an agent-based simulation platform, the AMES Wholesale Power Market Test Bed, to systematically investigate the performance of these markets through intensive computational experiments. AMES includes three types of agents: ISO, GenCos, and LSEs. The ISO manages two types of markets - a day-ahead market and a real-time market. The interaction and activities of these agents are modeled and simulated based on information provided in business practices manuals from real-world restructured wholesale power markets.

The primary objective of this dissertation is to gain a better understanding of the basic performance capabilities of U.S. restructured wholesale power markets with respect to both market efficiency and supply adequacy, pointing out some issues and potential improvements in market design along the way.

The main contributions of the dissertation can be summarized as follows:

1. Findings from GenCo learning calibration experiments show that learning calibration strongly affects GenCo net earnings. GenCos are able to obtain substantially greater net earnings when their learning parameters have first been adjusted to ensure good average net earnings over the range of possible outcomes for the particular learning environment at hand. This suggests that future agent-based electricity market research should permit GenCos and other cognitive market participants to engage in learning-to-learn processes to endogenously adjust their learning algorithms to the learning problems they face.

2. Findings from GenCo learning experiments in which the GenCos use calibrated VRE stochastic reinforcement learning to decide which supply offers to report to the ISO demonstrate that this simple learning algorithm works well for the GenCos. It permits the GenCos to obtain high net earnings relative to the benchmark no-learning case in which they report their true cost and capacity attributes to the ISO. The VRE learning method is extremely easy to calibrate and implement, even in complicated multi-agent learning environments. It involves only simple updating rules, and the only required information is the decision maker's own past net earnings outcomes.

3. Findings from learning experiments in which GenCos use calibrated VRE stochastic reinforcement learning to choose their supply offers demonstrate that the GenCos are quickly able to learn that their most profitable strategy is to implicitly collude on higher-than-true reported marginal cost functions, i.e., to jointly engage in economic capacity withholding. The GenCos are especially successful at raising prices (and their net earning) when LSE demand is 100% fixed (no price sensitivity) and the ISO is forced to meet this demand in every hour no matter how expensive the required generation might be. These findings suggest the importance of encouraging a greater sensitivity of LSE demand to price. However, even with 100% price-sensitive LSE demand, the GenCos are still able to exercise some degree of market power, resulting in higher prices. Consequently, what appears to be needed is more active demand-side bidding on the part of the LSEs to counter the potential market power of the GenCos. Absent this, the only way that ISOs can hope to prevent the exercise of GenCo market power is through the imposition of strong market power mitigation rules that constrain GenCo supply-offer behaviors.

4. Findings from price-cap experiments, which investigate the market performance effects of changes in GenCo supply-offer price caps, show that improperly imposed price caps can lead to increased LMP spiking and volatility as well as increased reliability issues through inducement of supply inadequacy, particularly around peak-demand hours. This is the case even if LMP values are indeed lowered during other hours.

5. Findings from ISO net surplus (congestion rent) experiments show that ISO net surplus substantially increases as the price-sensitivity of demand is reduced and the learning capabilities

of generators are increased, conditions resulting in greater economic capacity withholding and possible market inefficiency (i.e., wastage of resources) due to out-of-merit-order dispatch. A practical implication is that a more transparent public oversight of all net surplus collections and uses in wholesale power markets operating under LMP would be publicly prudent because these collections are not structurally well-aligned with market efficiency objectives.

6. Findings from GenCo capacity withholding experiments demonstrate that economic capacity withholding is much more advantageous for GenCos than physical capacity withholding in terms of raising their average net earnings, even when both forms of withholding are simultaneously possible for the GenCos. These findings suggest that ISOs should monitor GenCo reported supply offers particularly for economic capacity withholding.

7. Findings from preliminary SCUC/SCED experiments (not reported on in this dissertation) show that, for a 5-bus test case, the LMPs resulting under SCUC/SCED are more volatile than the LMPs resulting under SCED alone. Also GenCo dispatch levels are different even though the branch power flows do not change much. The reason is that the SCUC/SCED problem formulations considers more constraints, such as ramping rates and start/stop costs.

8. In order to carry out the above research, several major extensions of the AMES Wholesale Power Market Test Bed have been developed and tested. To date, AMES appears to be the only software for wholesale power market research that has been made publicly available as an open source package. It has already provided a useful foundation for the work of other electricity market researchers (e.g., in Australia, Germany, and China), and it could facilitate the communication and accumulation of electricity market research results in the future.

7.2 Future Work

In the past few years, agent-based modeling has become increasingly popular, and it is now widely used in a number of different research areas. Specifically, as reported more carefully in Chapter 2, this is the case for electricity market research. One important reason for the growing interest in agent-based modeling among electricity market researchers is that it permits them to investigate the impacts of market designs on market participant behaviors as well as the impacts in turn of these behaviors on market performance, where these market processes operate over realistically rendered transmission and distribution grids.

The ultimate goal for agent-based electricity market test beds is to include more types of agents, permitting the modeling of market operations at very detailed levels. Agent-based modeling permits the incorporation of real-world market rules together with important physical aspects of real-world power system networks. The resulting agent-based models can exploit the powerful computational capability of multi-core cpu or super computers, thus enabling the useful simulation of real-world electricity markets for different purposes.

Based on the experience I have gained in performing this dissertation research, including the development of the AMES test bed, I am considering the following areas as possible directions for future work:

1. The study of financial and operational risk management for restructured wholesale power markets at three levels of concern: (a) regulatory level (FERC), (b) market operation level (ISO/RTO), (c) and market trader level (e.g. GenCos, LSEs, and other market traders). The aim would be to analysis financial and operational risks under current restructured wholesale power markets designs, and to propose practical financial and operational risk management methods to decrease or avoid these risks.

2. The study of impacts on market performance when GenCo agents use alternative learning behaviors and strategies. Different GenCo learning methods could be used to examine the effects of learning in given electricity market environments, and to determine which learning methods result in the best outcomes for the GenCo agents.

3. The open-source release of AMES (V3.0), which is currently in the final development stage. This version, an extension of AMES (V2.05) released in September 2009, will include a fully operational two-settlement system. In addition, for the day-ahead market it will include an option to run a full SCUC/SCED process in place of the SCED-only process incorporated in AMES(V2.05).

4. Fuller evaluation of the test bed capabilities of AMES through the development and study of additional empirically-motivated test cases. The requirements for these test cases are: (a) data input validity, including reasonable data for the physical and technical side of power systems such as transmission line parameters, generation unit costs and capacities, LSE fixed loads, and price-sensitive demand bids; and (b) data output validity, meaning reasonable matchings between simulated data and key real-world data such as LMPs, GenCo net earnings, LSE payments, and ISO net surplus collections.

5. Extension of AMES to permit the study of impacts on wholesale power market performance when LSE agents also have learning capabilities. Real-world LSEs do not typically engage in active strategic demand bidding in current wholesale power markets because their wholesale power market payments are reimbursed through regulated retail rates. However, the issue of LSE learning could become more important in the future if retail systems are restructured so LSE reimbursements from retail sales are more fully determined by competitive market forces.

6. Extension of AMES to include load, transmission operating conditions and generator operating uncertainties in real-time market implementations. This would provide a fully operational two-settlement system for AMES whose performance could be tested through intensive experiments.

7. Extension of Ames to permit the experimental study and evaluation of emission constraints and other mandated environmental protection measures. If carbon emission reductions are mandated by law, this will affect wholesale power markets dramatically. Preventive measures could by pre-tested using computational experiments.

8. Extension of AMES to include the existence of multiple ISO-managed wholesale power markets. This would permit the study and evaluation of seaming issues for the seven current ISO-managed energy regions in the U.S., including their actual and potential interactions and connections at the national level. This extended version of AMES could facilitate a high-level understanding of current U.S. wholesale power market operations in the short term, as well as inform the process of investment in generation and transmission network expansion in the longer term under alternative scenarios for future energy policies.

APPENDIX A. AMES WHOLESALE POWER MARKET TEST BED

A.1 Introduction

In an April 2003 white paper FERC (2003), the U.S. Federal Energy Regulatory Commission (FERC) proposed a new market design for U.S. wholesale power markets. Over 50% of U.S. generating capacity is now operating within the footprint of a wholesale power market restructured in compliance with the basic provisions of FERC's design.

These restructured wholesale power markets are complex, involving physical constraints, complicated market protocols, and behavioral dispositions of human participants. Moreover, time series are short due to the relative recency of the restructuring efforts, and the data that are available are often released only with a delay and only in partially masked form. Consequently, it is difficult to model and study these markets using standard analytical and statistical tools.

An additional complicating factor is that many economists are not familiar with transmission grid aspects of power systems, so they often focus on highly simplified two-bus or three-bus systems. Conversely, many power engineers are not familiar with basic economic market concepts, let alone the complicated design of restructured wholesale power markets. Modeling efforts by interdisciplinary teams capable of addressing both engineering and economic concerns would therefore be highly desirable.

In response to these concerns, an interdisciplinary group of researchers at Iowa State University has undertaken an OSS development of a wholesale power market test bed, referred to as AMES (*A*gent-based *M*odeling of *E*lectricity *S*ystems). The AMES test bed permits the systematic experimental study of strategic trading behaviors within restructured wholesale power markets operating over realistically rendered AC transmission grids. In addition, AMES facilitates augmentation of empirical input data with simulated input data to permit the study of a broader array of scenarios.

From the beginning, AMES was designed for research and teaching purposes rather than for commercial-grade application. AMES is entirely developed in the widely used Java programming language in order to facilitate readability and use. AMES is entirely OSS, combining together a collection of basic OSS modules for learning representation, optimal power flow solution, graphic display, and other functions. The modular and extensible OSS architecture of AMES permits users to modify and extend the code with relative ease to suit their special needs.

The first version of AMES was released as OSS at the IEEE Power and Energy Society General Meeting (PES GM) in 2007, and a substantially expanded second version was released as OSS at the IEEE PES GM in 2008. Downloads, manuals, and tutorial information for all AMES version releases to date are accessible at the AMES homepage Tesfatsion, L. (2009d). AMES is also available for downloading at the software site of the IEEE Task force on Open Source Software for Power Systems; see IEEE OSS (2009).

The release of AMES as OSS is intended to encourage the cumulative development of this test bed by multiple researchers in directions appropriate for their specific needs. It is also intended to encourage continual dialog with market stakeholders and regulators leading to successive refinements and improvements of the test bed.

A.2 Key Features

The latest version release AMES(V2.05) of the AMES Market Package incorporates, in simplified form, core features of FERC's proposed wholesale power market design, a detailed description can be found in Chapter 3.

As seen in Fig. A.1, 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 parameters for GenCos; (e) specification of simulation controls (e.g., the simulation stopping rule); and (f) customization of table and chart output displays.

The user can control the length of each simulation run by choosing to set (or not) any combination of the following five stopping rules:

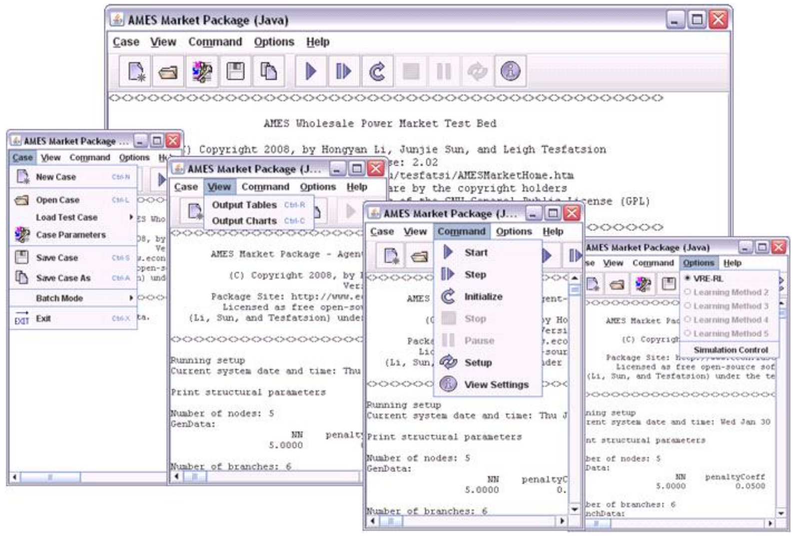

Figure A.1 AMES Graphical User Interface (GUI)

- Stop when a specified maximum day is reached.
- Stop when each GenCo is choosing a single supply offer with a probability that exceeds a user-specified threshold probability.
- Stop when the probability distribution used by each GenCo to select its supply offers has stabilized to within a user-specified threshold for a user-specified number of days.
- Stop when the supply offer selected by each GenCo has stabilized to within a user-specified threshold for a user-specified number of days.
- Stop when the net earnings of each GenCo have stabilized to within a user-specified threshold for a user-specified number of days.

When multiple stopping rules are flagged, the simulation run terminates as soon as any one of the flagged stopping rules is satisfied.

A.3 Running AMES Simulation Experiments

Detailed instructions for developing and running general AMES simulations in either single-run or batch-run mode can be found in the set-up information file included with the AMES software download; see Tesfatsion, L. (2009d). Here we briefly outline the general sequence of actions for a single simulation run, as follows:

1: To load one of the pre-set test cases (e.g., the 5-Bus Test Case), use the “Case → Load Test Case → 5-Bus Test Case” command sequence on the GUI menu. Alternatively, to create a new case, use the “Case → New Case” command sequence on the GUI menu.

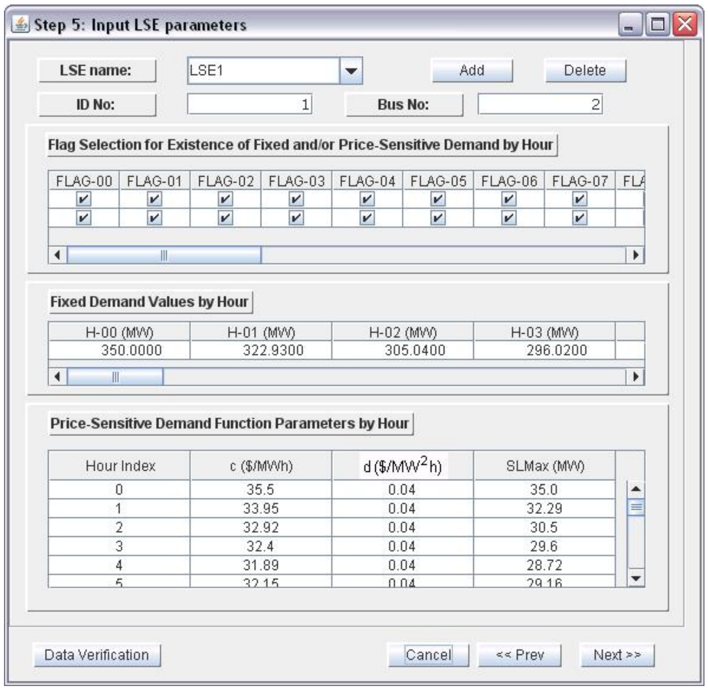

Figure A.2 AMES GUI: Setting screen for LSE fixed demand bids and price-sensitive demand function parameters for each hour

2: For a pre-set test case, either use the default parameter settings (including the default random seed value) or change some or all of these default settings to other admissible values. To change the default parameter settings, or to set parameters for a new case, use the “Case → Case Parameters” command sequence on the GUI menu to access a sequence of setting screens for grid, LSE, GenCo, and simulation control parameters; see, e.g., fig. A.2. Inadmissible parameter settings trigger explanatory error messages.

3: To run the case, click the “Start” button on the GUI toolbar or use the “Command → Start” command sequence from the GUI menu.

4: View a customizable output file in the AMES DATA directory using various programs (e.g., Microsoft Excel, Wordpad).

5: Alternatively, view output data in either table or chart form using either the "View → Output Tables" or the "View → Output Charts" command sequence from the GUI menu.

The AMES GUI tables and charts display six types of output for each run: GenCo commitments; GenCo profits and net earnings; cleared LSE price-sensitive demand; LSE net earnings corresponding to cleared price-sensitive demand; LMPs; and total supply and demand curves. This output is further subdivided into "benchmark" and "learning" portions as follows: (i) initially generated output for a no-learning *benchmark case* in which the supply offers that the GenCos report to the ISO reflect their true cost and capacity attributes; and (ii) subsequently generated output for a *learning case* in which the GenCos attempt to learn over time which supply offers to report to the ISO to increase their net earnings. The benchmark-case output provides a benchmark of comparison for the learning-case output.

A.4 Development Tools Used

As seen in Fig. A.3, various OSS tools were used in the development of AMES(V2.05). A more careful description of these tools is given below.

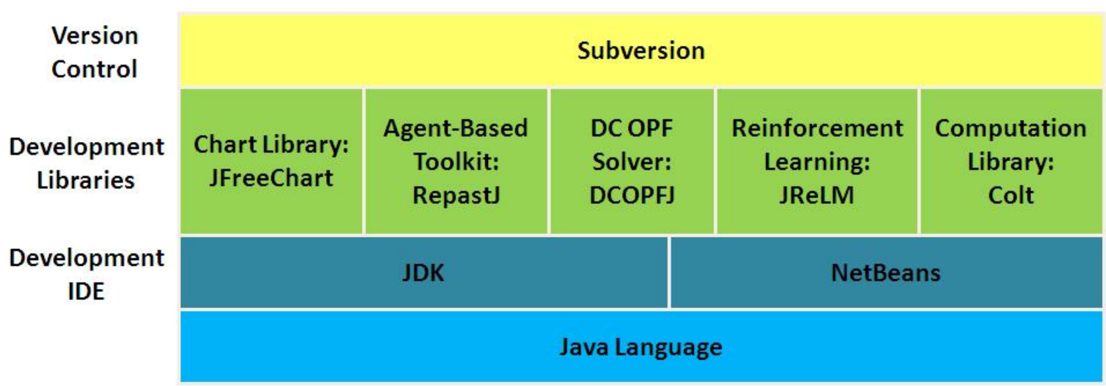

Figure A.3 OSS tools used in the development of AMES(V2.05)

- *(1) Java Development Kit (JDK):* Version 6 update 1 (6u1) of the Java SE Development Kit Java (2009) was used to develop the basic AMES code.

- *(2) Java Integrated Development Environment (IDE):* The NetBeans IDE 6.0 NetBeans (2009) was used in combination with (1) to develop AMES as a standard Java package. The NetBeans IDE is a powerful open-source cross-platform tool for Java programming.

- *(3) Java Chart Library:* JFreeChart JFreeChart (2009) is an OSS chart library that is used to create different kinds of charts in AMES.

- *(4) Java Agent-Based Toolkit:* The OSS agent-based toolkit Repast J RepastJ (2009) is used to control agent activities in AMES.

- *(5) Java DC-OPF Solver (DCOPFJ):* DCOPFJ, a free open-source solver for DC optimal power flow problems developed by Sun and Tesfatsion (Sun, J. and Tesfatsion, L. (2007b)), is used by the AMES ISO to solve hourly bid/offer-based DC-OPF problems.

- *(6) Java Reinforcement Learning Module (JReLM):* JReLM, an open-source Java learning module developed by Charles Gieseler Gieseler, C. (2005), is used to control agent learning behaviors in AMES.

- *(7) Colt Libraries for Java High-Performance Computing:* Colt Colt (2009) provides a set of open-source libraries for high-performance scientific and technical computing in Java.

- *(8) Subversion for Version Management:* Subversion Subversion (2009) is an open source version control system facilitating simultaneous software development by multiple project members.

A.5 Licensing and Release

The AMES Wholesale Power Market Test Bed is licensed by the copyright holders (Hongyan Li, Junjie Sun, and Leigh Tesfatsion) as free open-source software under the terms of the GNU General Public License (GPL) GNU (2009). Anyone who is interested is allowed to view, modify, and/or improve upon the code for AMES, but any software generated using all or part of this code must be released as free open-source software in turn.

The latest version of AMES can be downloaded in a zip file from the AMES homepage Li, H. and Tesfatsion, L. (2009a). This zip file includes three help files and three file directories, as follows:

- *(1) AMESMarketReadMe.htm:* This help file is a slightly modified version of the AMES homepage Li, H. and Tesfatsion, L. (2009a). It provides useful information about AMES, including an overview of capabilities, pointers to current and previous versions of AMES for download, and annotated pointers to project papers and publications making use of AMES.

- *(2) AMESMarketProjectSetupInfo.pdf:* This help file is a basic manual for AMES. Included are instructions for setting up AMES as a Java project, loading and viewing AMES test cases, developing new AMES test cases, modifying the AMES source code, and running AMES experiments in either individual or batch mode.

- *(3) AMESVersionReleaseHistory.htm:* This help file provides annotated pointers to all released AMES versions.

- *(4) src directory:* This directory includes all of the AMES source code.

- *(5) lib directory:* This directory includes all special OSS libraries used to implement AMES, e.g., an executable Java archive ("jar") file Colt.jar for Colt Colt (2009).

- *(6) DATA directory:* This directory provides sample data input files for 5-bus and 30-bus test cases that can be used as templates for the development of new AMES test cases. It also provides a sample batch-mode file that can be used as a template for setting up AMES experiments so that multiple runs can be implemented automatically in one batch.

A.6 Applications to Date

The AMES test bed has been used to conduct systematic experimental studies focusing on various performance aspects of restructured wholesale power markets operating over AC transmission grids when profit-seeking GenCos have learning capabilities allowing them to strategically evolve their supply offers over time (Sun, J. and Tesfatsion, L. (2007a)-Li, H. , Sun, J. and Tesfatsion, L. (2009)).

Basically, AMES is used by two categories:

(1) E3 group members in Iowa State University. Currently, a few PhD students build their research based on AMES, for example, two PhD students are studying short time LMP forecasting in day-ahead

markets; one PhD student is studying Financial Transmission Rights (FTR), one PhD student is seaming AMES with retail power markets and one PhD student is studying financial and operational risk management for restructured wholesale power markets, etc. For these PhD students, AMES provide platform to carry their research work.

(2) People who are interested in AMES from all over the world. The AMES software package has been downloaded from the AMES homepage Li, H. and Tesfatsion, L. (2009a) by numerous researchers worldwide, and several research groups have indicated through communications with AMES team members that the AMES OSS has helped them in the design of their own project software. For example, research groups in Germany and Australia have used aspects of the AMES OSS to build agent-based test beds for the study of C02 emission control systems for power markets, and a research group in China has developed an AC OPF module for possible inclusion in a future version of AMES.

Table A.1 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^U_{km}	Thermal limit (MW) for real power flow on km	$P^U_{km} > 0$, $km \in$ 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^F_{Lj}(H)	Benchmark-case hour-H fixed demand (MW) for LSE j	BP^F_{Lj}(H) > 0
p^F_{Lj}(H)	Actual hour-H fixed demand (MW) for LSE j	p^F_{Lj}(H) $=$ [1-R_j]*BP^F_{Lj}(H)
$SLMax_j$(H)	Hour-H upper limit for LSE j's price-sensitive demand (MW)	$SLMax_j$(H) $=$ R_j*BP^F_{Lj}(H)
$MPTD_j$(H)	Hour-H maximum potential total demand (MW) for LSE j	$MPTD_j$(H)= [p^F_{Lj}(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)) ≥ 0
$SCost_i$	Hourly pro-rated sunk cost ($/h) for GenCo i	$SCost_i \geq 0$
Cap^L_i	Lower real power operating capacity limit (MW) for GenCo i	$Cap^L_i \geq 0$
Cap^U_i	Upper real power operating capacity limit (MW) for GenCo i	$Cap^U_i > 0$
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	$MC_i(Cap^L_i) > 0$
$InitMoney_i$	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
$RIMax^L_i$	Ordinate range-index parameter for AD_i construction	$RIMax^L_i \in [0, 1)$
$RIMax^U_i$	Slope range-index parameter for AD_i construction	$RIMax^U_i \in [0,1)$
$RIMin^C_i$	Capacity-withholding range-index parameter for AD_i construction	$RIMin^C_i \in (0, 1]$
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)	$T_i > 0$
ρ_i	Recency parameter for GenCo i (learning)	$0 \leq \rho_i \leq 1$
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 A.2 Endogenous Variables

Variable	Description
p_{Lj}^S	Real-power price-sensitive demand (MW) by LSE j=1,...,J
a_i^R,b_i^R	Cost coefficients (\$/MWh,\$/MW^{2}h) reported by GenCo i=1,...,I
Cap_i^{RU}	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

APPENDIX B. SYSTEM DATA FOR 5-BUS TEST SYSTEM

Table B.1 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_i^L	Cap_i^U	$InitMoney_i$
1	1	0.00	14.0	0.005	0.0	110.0	\$1M
2	1	0.00	15.0	0.006	0.0	100.0	\$1M
3	3	0.00	25.0	0.010	0.0	520.0	\$1M
4	4	0.00	30.0	0.012	0.0	200.0	\$1M
5	5	0.00	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	$BP^F(16)$	$BP^F(17)$	$BP^F(18)$	$BP^F(19)$	$BP^F(20)$	$BP^F(21)$	$BP^F(22)$	$BP^F(23)$
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_{km}^U (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 B.2 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 B.3 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)

APPENDIX C. SYSTEM DATA FOR 30-BUS TEST SYSTEM

Table C.1 Numerical input specifications for the benchmark dynamic 30-bus test case (branch and GenCo cost data)

From	To	MaxCap	x				
2	6	30	0.1763				
4	6	30	0.0414				
4	12	65	0.2560				
5	7	30	0.1160				
6	7	30	0.0820				
6	8	30	0.0420				
6	9	30	0.2080				
6	10	30	0.5560				
6	28	30	0.0599				
8	28	30	0.2000				
9	10	30	0.1100				
9	11	30	0.2080				
10	17	32	0.0845				
10	20	32	0.2090				
10	21	30	0.0749				
10	22	30	0.1499				
12	13	65	0.1400				
12	14	32	0.2559				
12	15	32	0.1304				
12	16	32	0.1987				
14	15	16	0.1997				
15	18	16	0.2185				
15	23	16	0.2020				
16	17	16	0.1923				
18	19	16	0.1292				
19	20	32	0.0680				
21	22	30	0.0236				
22	24	30	0.1790				
23	24	16	0.2700				
24	25	30	0.3292				
25	26	30	0.3800				
25	27	30	0.2087				
27	28	30	0.3960				
27	29	30	0.4153				
27	30	30	0.6027				
29	30	30	0.4533				
GenCo i	at bus	SCost_i	a_i	b_i	Cap_i^L	Cap_i^U	InitMoney_i
1	1	0.00	10.6940	0.0046	0.00	100.0	$1M
2	2	0.00	18.1000	0.0061	0.00	80.0	$1M
3	5	0.00	13.3270	0.0087	0.00	50.0	$1M
4	8	0.00	13.3530	0.0089	0.00	50.0	$1M
5	11	0.00	37.8890	0.0143	0.00	20.0	$1M
6	13	0.00	19.3270	0.0103	0.00	70.0	$1M
7	15	0.00	18.3000	0.0071	0.00	60.0	$1M
8	24	0.00	39.8890	0.0163	0.00	20.0	$1M
9	30	0.00	49.3270	0.0243	0.00	20.0	$1M

Table C.2 Numerical input specifications for the benchmark dynamic 30-bus test case (100% fixed demand data)

LSE j	at bus	$\mathrm{BP}^F(00)$	$\mathrm{BP}^F(01)$	$\mathrm{BP}^F(02)$	$\mathrm{BP}^F(03)$	$\mathrm{BP}^F(04)$	$\mathrm{BP}^F(05)$	$\mathrm{BP}^F(06)$	$\mathrm{BP}^F(07)$
1	2	16.93	15.62	14.76	14.32	13.89	14.11	14.32	15.19
2	3	1.87	1.73	1.63	1.58	1.54	1.56	1.58	1.68
3	4	52.74	48.66	45.97	44.60	43.27	43.94	44.60	47.32
4	5	26.68	24.62	23.26	22.57	21.89	22.23	22.57	23.94
5	7	17.79	16.41	15.50	15.04	14.59	14.82	15.04	15.96
6	8	23.41	21.59	20.40	19.79	19.20	19.50	19.79	21.00
7	10	4.53	4.17	3.94	3.83	3.71	3.77	3.83	4.06
8	12	8.74	8.06	7.62	7.39	7.17	7.28	7.39	7.84
9	14	4.84	4.46	4.22	4.09	3.97	4.03	4.09	4.34
10	15	6.40	5.90	5.58	5.41	5.25	5.33	5.41	5.74
11	16	2.73	2.52	2.38	2.31	2.24	2.28	2.31	2.45
12	17	7.02	6.48	6.12	5.94	5.76	5.85	5.94	6.30
13	18	2.50	2.30	2.18	2.11	2.05	2.08	2.11	2.24
14	19	7.41	6.84	6.46	6.27	6.08	6.18	6.27	6.65
15	20	1.72	1.58	1.50	1.45	1.41	1.43	1.45	1.54
16	21	13.65	12.60	11.90	11.55	11.20	11.38	11.55	12.25
17	23	2.50	2.30	2.18	2.11	2.05	2.08	2.11	2.24
18	24	6.79	6.26	5.92	5.74	5.57	5.66	5.74	6.09
19	26	2.73	2.52	2.38	2.31	2.24	2.28	2.31	2.45
20	29	1.87	1.73	1.63	1.58	1.54	1.56	1.58	1.68
21	30	8.27	7.63	7.21	6.99	6.79	6.89	6.99	7.42
LSE j	at bus	$\mathrm{BP}^F(08)$	$\mathrm{BP}^F(09)$	$\mathrm{BP}^F(10)$	$\mathrm{BP}^F(11)$	$\mathrm{BP}^F(12)$	$\mathrm{BP}^F(13)$	$\mathrm{BP}^F(14)$	$\mathrm{BP}^F(15)$
1	2	17.36	19.10	19.53	19.75	19.53	19.10	18.88	18.88
2	3	1.92	2.11	2.16	2.18	2.16	2.11	2.09	2.09
3	4	54.08	59.49	60.84	61.52	60.84	59.49	58.83	58.83
4	5	27.36	30.10	30.78	31.12	30.78	30.10	29.76	29.76
5	7	18.24	20.06	20.52	20.75	20.52	20.06	19.84	19.84
6	8	24.00	26.40	27.00	27.30	27.00	26.40	26.11	26.11
7	10	4.64	5.10	5.22	5.28	5.22	5.10	5.05	5.05
8	12	8.96	9.86	10.08	10.19	10.08	9.86	9.75	9.75
9	14	4.96	5.46	5.58	5.64	5.58	5.46	5.40	5.40
10	15	6.56	7.22	7.38	7.46	7.38	7.22	7.14	7.14
11	16	2.80	3.08	3.15	3.19	3.15	3.08	3.05	3.05
12	17	7.20	7.92	8.10	8.19	8.10	7.92	7.83	7.83
13	18	2.56	2.82	2.88	2.91	2.88	2.82	2.78	2.78
14	19	7.60	8.36	8.55	8.65	8.55	8.36	8.27	8.27
15	20	1.76	1.94	1.98	2.00	1.98	1.94	1.91	1.91
16	21	14.00	15.40	15.75	15.93	15.75	15.40	15.23	15.23
17	23	2.56	2.82	2.88	2.91	2.88	2.82	2.78	2.78
18	24	6.96	7.66	7.83	7.92	7.83	7.66	7.57	7.57
19	26	2.80	3.08	3.15	3.19	3.15	3.08	3.05	3.05
20	29	1.92	2.11	2.16	2.18	2.16	2.11	2.09	2.09
21	30	8.48	9.33	9.54	9.65	9.54	9.33	9.22	9.22
LSE j	at bus	$\mathrm{BP}^F(16)$	$\mathrm{BP}^F(17)$	$\mathrm{BP}^F(18)$	$\mathrm{BP}^F(19)$	$\mathrm{BP}^F(20)$	$\mathrm{BP}^F(21)$	$\mathrm{BP}^F(22)$	$\mathrm{BP}^F(23)$
1	2	19.75	21.70	20.83	20.62	20.40	19.96	18.88	17.58
2	3	2.18	2.40	2.30	2.28	2.26	2.21	2.09	1.94
3	4	61.52	67.60	64.90	64.22	63.54	62.19	58.81	54.76
4	5	31.12	34.20	32.83	32.49	32.15	31.46	29.75	27.70
5	7	20.75	22.80	21.89	21.66	21.43	20.98	19.84	18.47
6	8	27.30	30.00	28.80	28.50	28.20	27.60	26.10	24.30
7	10	5.28	5.80	5.57	5.51	5.45	5.34	5.05	4.70
8	12	10.19	11.20	10.75	10.64	10.53	10.30	9.74	9.07
9	14	5.64	6.20	5.95	5.89	5.83	5.70	5.39	5.02
10	15	7.46	8.20	7.87	7.79	7.71	7.54	7.13	6.64
11	16	3.19	3.50	3.36	3.33	3.29	3.22	3.05	2.84
12	17	8.19	9.00	8.64	8.55	8.46	8.28	7.83	7.29
13	18	2.91	3.20	3.07	3.04	3.01	2.94	2.78	2.59
14	19	8.65	9.50	9.12	9.03	8.93	8.74	8.27	7.70
15	20	2.00	2.20	2.11	2.09	2.07	2.02	1.91	1.78
16	21	15.93	17.50	16.80	16.63	16.45	16.10	15.23	14.18
17	23	2.91	3.20	3.07	3.04	3.01	2.94	2.78	2.59
18	24	7.92	8.70	8.35	8.27	8.18	8.00	7.57	7.05
19	26	3.19	3.50	3.36	3.33	3.29	3.22	3.05	2.84
20	29	2.18	2.40	2.30	2.28	2.26	2.21	2.09	1.94
21	30	9.65	10.60	10.18	10.07	9.96	9.75	9.22	8.59

Table C.3 Additional numerical input specifications for the benchmark dynamic 30-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
6	10	10	1	0.75	0.75	1.00	0.001
7	10	10	1	0.75	0.75	1.00	0.001
8	10	10	1	0.75	0.75	1.00	0.001
9	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	383889.60	1	100
2	0.04	0.96	520350.72	1	100
3	0.04	0.96	239368.80	1	100
4	0.04	0.96	239824.80	1	100
5	0.04	0.96	272665.44	1	100
6	0.04	0.96	485835.84	1	100
7	0.04	0.96	394672.32	1	100
8	0.04	0.96	287046.24	1	100
9	0.04	0.96	354923.04	1	100

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

BIBLIOGRAPHY

Aldete, G. B. (2005). Alternative models to analyze market power and financial transmission rights in electricity markets. *Ph.D. Thesis*, Department of Electrical and Comp. Eng. U. of Waterloo, Ontario.

Baldick, R. and Hogan, W. (2002). 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.

Benders, J. F. (1962). Partitioning procedures for solving mixed-variables programming problems. *Journal Numerische Mathematik*, *4*, (1) 238–252.

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

CAISO (2009). Market Issues & Performance, 2008 Annual Report. Department of Market Monitoring, California Independent System Operator Corporation. www.caiso.com/2390/239087966e450.pdf

CASIO (2004). Transmission Economic Assessment Methodology. www.caiso.com/docs/2004/06/03/2004060313241622985.pdf

Cheng, X. and Overbye, T. J. (2006). An Energy Reference Bus Independent LMP Decomposition Algorithm. *IEEE Transactions on Power Systems*, *21*,(3) 1041–1049.

Colt (2009). Colt Project: A Set of Open-Source Libraries for High Performance Scientific and Technical Computing in Java. acs.lbl.gov/~hoschek/colt/

Conejo, A. J. , Castillo, A. J. , Minguez, R. and Milano, F. (2005). Locational Marginal Price Sensitivities. *IEEE Transactions on Power Systems*, *20*(4), 2026–2033.

Conzelmann, G. , North, M. , Boyd, G. , Cirillo, R. , Koritarov, V. , Macal, C. , Thimmapuram, P. and Veselka, T. (2004). Simulating Strategic Market Behavior Using an Agent-Based Modeling Approach. *6th IAEE European Energy Conference on "Modeling in Energy Economics and Policy 28*(4), 305–342.

Deshpande, A. and Riehle, D. (2008). The Total Growth of Open Source. *Proceedings*, Fourth Conference on Open Source Systems (OSS 2008), Springer Verlag.

E3 (2009). Electric Energy Economics (E3) Group. Iowa State University, Ames, Iowa, www.econ.iastate.edu/tesfatsi/E3GroupISU.htm

EMCAS (2009). Electricity Market Complex Adaptive System (EMCAS). Available: http://www.dis.anl.gov/projects/emcas.html

Erev, I. and Roth, A. E. (1998). Predicting how people play games with unique mixed-strategy equilibria. *American Economic Review*, *88*, 848–881.

FERC (2003). *Notice of White Paper*, U.S. Federal Energy Regulatory Commission.

Geoffrion, A. M. (1972). Generalized benders decomposition. *Journal of Optimization Theory and Applications*, *10*, (4) 237–260.

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

GNU (2009). GNU General Public License (GPL). www.gnu.org/licenses/licenses.html

Gross, G. and Bompard, E. (2004). Optimal power flow application issues in the pool paradigm. *Electric Power and Energy System*, *26*, 787–796.

Hausman, E. , Fagan, R. , White, D. , Takahashi, K. and Napoleon, A. (2006). LMP Electricity Markets: Market Operations, Market Power, and Value for Consumers. *Synapse Energy Economics*.

Hobbs, B. , Rothkopf, M. , ,O'Neill, R. and Chao, H. (2001). The Next Generation of Electric Power Unit Commitment Models. Norwell, MA: Kluwer Academic Publishers.

IEEE OSS (2009). Task Force on Open Source Software for Power Systems, ewh.ieee.org/cmte/psace/CAMS_taskforce/index.htm

ISO-NE (2009). ISO New England Inc. Home page. [Online]. Available: www.iso-ne.com/

ISO-NE Report(2009). 2008 Annual Markets Report in ISO New England. Independent Market Monitoring Unit, ISO New England, Inc. http://www.iso-ne.com/markets/mktmonmit/rpts/other/amr08_final_061709.pdf

Java (2009). Java SE Development Kit (JDK), java.sun.com

JFreeChart (2009). JFreeChart, www.jfree.org/jfreechart

Joskow, P. (2006). Markets for power in the united states: An interim assessment. *The Energy Journal*, *27*(1), 1–36.

Kirschen, D. S. and Strbac, G. (2004). *Fundamentals of Power Economics*, Wiley, 2004.

Lally, J. (2002). Financial transmission rights: Auction example in Financial Transmission Rights Draft (01-10-02), ISO New England, Inc., section 6.

Li, H. , Sun, J. and Tesfatsion, L. (2010). Testing Institutional Arrangements via Agent-Based Modeling: A U.S. Electricity Market Application, in H. Dawid (Ed.). *Computational Methods in Economic Dynamics*, Springer-Verlag, to appear.

Li, H. , Sun, J. and Tesfatsion, L. (2009). Separation and volatility of locational marginal prices in restructured wholesale power markets. *ISU Economics Working Paper #09009*. www.econ.iastate.edu/research/publications/viewabstract.asp?pid=13075

Li, H. and Tesfatsion, L. (2009a). ISO Net Surplus Extraction in Restructured Wholesale Power Markets. *ISU Economics Working Paper No. 09015*, revise and resubmit requested by the *IEEE Transactions on Power Systems*.

Li, H. and Tesfatsion, L. (2009b). The AMES Wholesale Power Market Test Bed: A Computational Laboratory for Research, Teaching, and Training. *IEEE Proceedings*, Power and Energy Society General Meeting, Calgary, Canada.

Li, H. and Tesfatsion, L. (2009c). Development of Open Source Software for Power Market Research: The AMES Test Bed. *Journal of Energy Markets*, *2*,(2) 111–128.

Li, H. and Tesfatsion, L. (2009d). Capacity Withholding in Restructured Wholesale Power Markets: An Agent-Based Test Bed Study. *IEEE Proceedings*, Power Systems & Exposition Conference, Seattle, WA.

Li, H. and Tesfatsion, L. (2009e), MISO Energy Project Homepage, Data section http://www.econ.iastate.edu/tesfatsi/MISOenergygroup.htm#Data

Li, H. , Sun, J. and Tesfatsion, L. (2008). Dynamic LMP Reponse Under Alternative Price-Cap and Price-Sensitive Demand Scenarios. *IEEE Proceedings*, Power and Energy Society General Meeting, Pittsburgh, PA.

Liu, H. , Tesfatsion, L. and Chowdhury, A. A. (2009). Derivation of Locational Marginal Prices for Restructured Wholesale Power Markets. *Journal of Energy Markets*, *2*,(1) 3–27.

Liu, H. , Shen, Y. , Zabinsky, Z. , Liu, C.C. , Courts, A. and Joo, S.K. (2008). Social Welfare Maximization in Transmission Enhancement Considering Network Congestion. *IEEE Transactions on Power Systems*, *23*, 1105–1114.

Litvinov, E. , Zhao, F. and Zheng, T. (2009). Alternative auction objectives and pricing schemes in short-term electricity markets. *IEEE Proceedings*, Power and Energy Society General Meeting, Calgary, Canada.

MISO (2009). Midwest ISO Inc. Homepage. [Online]. Available: www.midwestiso.org/

MISO Report(2008). 2007 State of the Market Report for the Midwest ISO. *Potomac Economics, Independent Market Monitor for the Midwest ISO*. www.midwestiso.org/publish/Document/

NetBeans (2009). NetBeans Integrated Development Environment. www.netbeans.org

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

Orfanogianni, T. and Gross, G. (2007). A general formulation for LMP evaluation. *IEEE Transactions on Power Systems*, *22*(3), 1163–1173.

Pentapalli, M. (2008). 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.

Pinto, H. , Magnago, F. , Brignone, S. , Alsac, O. , Stott, B. (2006). Security Constrained Unit Commitment: Network Modeling and Solution Issues. *IEEE Proceedings*, Power Systems & Exposition Conference, *1*,1759–1766.

PJM Report(2009). 2008 State of the Market Report, Vol 2: Detailed Analysis. Monitoring Analytics LLC, Independent Market Monitor for PJM. www.monitoringanalytics.com/reports/PJM_State_of_the_Market/2008/

Rassenti, S. , Smith, V. L. and Wilson, B. (2003). Controlling market power and price spikes in electricity networks: Demand-side bidding. *Proceedings of the National Academy of Sciences*, *100*, (5), 2998–3003.

RepastJ (2009). Repast J version 3.1, repast.sourceforge.net/repast_3/download.html

Roth, A. E. and Ido, E. (1995). Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games and Econ. Behavior*, *8*, 164–212.

Salazar, H. (2008). A critical appraisal of economic-driven transmission enhancement. *M.S. Creative Component*, Economics Department, Iowa State University.

Sauma, E. and Oren, S. (2009). Do generation firms in restructured electricity markets have incentives to support social-welfare-improving transmission investments? *Energy Economics*, *31*, 676–689.

Sensfuß, F. , Ragwitz, M. , Genoese, M. and Möst, D. (2007). Agent-Based Simulation of Electricity Markets: A Literature Review. *Working Paper Sustainability and Innovation*, Fraunhofer Institute Systems and Innovation Research

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

Somani, A. and Tesfatsion, L. (2008). An Agent-Based Test Bed Study of Wholesale Power Market Performance Measures. *IEEE Computational Intelligence Magazine*, *3*,(4) 56–72. Working paper available: http://www.econ.iastate.edu/tesfatsi/AMESPerformanceMeasures.ASLT.IEEECIM2008.pdf

Sourceforge (2009). Sourceforge.Net: Open Source Software, a repository of OSS, http://sourceforge.net/

Streiffert, D. , Philbrick, R. and Ott, A. (2005). A Mixed Integer Programming Solution for Market Clearing and Reliability Analysis. *IEEE Proceedings*, Power and Energy Society General Meeting, *3*, 2724–2731.

Subversion (2009). Subversion: An open source version control system, http://subversion.tigris.org/

Sueyoshi, T. and Tadiparthi, G. R. (2008). Wholesale power price dynamics under transmission line limits: A use of agent-based simulator. *IEEE Transactions on Systems, Man, and Cybernetics*, *38*(2), 229–241.

Sun, J. and Tesfatsion, L. (2008). DC-OPF Formulation with Price-Sensitive Demand Bids. *Working Paper*, ISU Economics Department. Available: www.econ.iastate.edu/tesfatsi/DC-OPF.PriceSensitiveDemandBids.pdf

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

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

Sutton, R. S. and Barto, A. G. (2000). Reinforcement Learning: An Introduction, Third Printing, The MIT Press, Cambridge, MA.

Tan, C.S. , Liu, Y. and Lo, K.L. (2008). A More Transparent Way of Financial Settlement for Congestion Cost in Electricity Markets. *DRPT2008*, Nanjing, China.

Tellidou, A. and Bakirtzis, A. (2007). Agent-Based Analysis of Capacity Withholding and Tacit Collusion in Electricity Markets. *IEEE Transactions on Power Systems*, *22*,(4) 1735–1742.

Tesfatsion, L. (2009a). Auction basics for restructured wholesale power markets: Objectives and pricing rules. *IEEE Proceedings*, Power and Energy Society General Meeting, Calgary, Alberta, CA. www.econ.iastate.edu/tesfatsi/AuctionBasics.IEEEPES2009.LT.pdf

Tesfatsion, L. (2009b). Agent-Based Computational Economics, hosted by the Economics Department, Iowa State University. Available: www.econ.iastate.edu/tesfatsi/ace.htm

Tesfatsion, L. (2009c). Agent-Based Computational Economics Research on Restructured Electricity Markets, hosted by the Economics Department, Iowa State University. Available: www.econ.iastate.edu/tesfatsi/aelect.htm

Tesfatsion, L. (2009d). AMES Wholesale Power Market Test Bed Homepage, hosted by the Economics Department, Iowa State University, http://www.econ.iastate.edu/tesfatsi/AMESMarketHome.htm

Tesfatsion, L. and Judd, K. L., eds. (2006). Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics, Handbooks in Economics Series, Amsterdam, the Netherlands: North-Holland/ Elsevier.

Walawalkar, R. , Blumsack, S. , Apt, J. and Fernands, S. (2008). An economic welfare analysis of demand response in the PJM electricity market. *Energy Policy*, *26*, 3692–3702.

Watkins, C. (1989). Learning from Delayed Rewards. Ph.D. Thesis, Cambridge University, UK.

Weidlich, A. and Veit, D. (2008). A Critical Survey of Agent-Based Wholesale Electricity Market Models. *Energy Economics*

Zhou, Z. , Chan, W. and Chow, J. (2007). Agent-based simulation of electricity markets: a survey of tools. *Artificial Intelligence Review 28*(4), 305–342.