Co-Learning Patterns As Emergent Market Phenomena: An Electricity Market Illustration

Hongyan Li, Leigh Tesfatsion

Working Paper No. 10042
December 2010
Revised on June 2011

IOWA STATE UNIVERSITY
Department of Economics
Ames, Iowa, 50011-1070

Co-learning patterns as emergent market phenomena: An electricity market illustration

Hongyan Li and Leigh Tesfatsion

Hongyan Li, Consulting Engineer, ABB Inc., Raleigh, North Carolina, and Leigh Tesfatsion (corresponding author: tesfatsi@iastate.edu), Professor of Economics, Mathematics, and Electrical and Computer Engineering, Iowa State University, Ames, IA 50011-1070 USA.

Abstract

The definition of emergence remains problematic, particularly for systems with purposeful human interactions. This study explores the practical import of this concept within a specific market context: namely, a double-auction market for wholesale electric power that operates over a transmission grid with spatially located buyers and sellers. Each profit-seeking seller is a learning agent that attempts to adjust its daily supply offers to its best advantage. The sellers are co-learners in the sense that their supply offer adjustments are in response to past market outcomes that reflect the past supply offer choices of all sellers. Attention is focused on the emergence of co-learning patterns, that is, global market patterns that arise and persist over time as a result of seller co-learning. Examples of co-learning patterns include correlated seller supply offer behaviors and correlated seller net earnings outcomes. Heat maps are used to display and interpret co-learning pattern findings. One key finding is that co-learning strongly matters in this auction market environment. Sellers that behave as Gode-Sunder budget-constrained zero-intelligence agents, randomly selecting their supply offers subject only to a break-even constraint, tend to realize substantially lower net earnings than sellers that tacitly co-learn to correlate their supply offers for market power advantages.

Key words: Emergence, co-learning, double auction, wholesale electric power market, capacity withholding, market power, AMES wholesale power market testbed, heat maps

Preprint submitted to New Perspectives on Emergence in Economics *June 8, 2011*

1. Introduction

Emergence is an intriguing multi-faceted concept whose meaning remains controversial, particularly for systems involving purposeful human interactions. Consequently, it is of interest to study the practical import of this concept for economics by examining its role in specific realistically-rendered economic contexts.

This study examines emergence in an empirically-based model of a double-auction market for wholesale electric power. The market operates over a 5-bus transmission grid with spatially located buyers and sellers. Each profit-seeking seller is a learning agent that attempts to adjust its daily supply offers to its best advantage. The sellers are individual learners in the sense that the learning method of each seller is calibrated (pre-tuned) to the attributes of the seller's specific decision environment to capture learning-to-learn effects. However, the sellers are also *co-learners* in the sense that the adjustments of their daily supply offers are in response to past market outcomes that reflect the past supply offer choices of all sellers.

Each seller in our model can engage in two forms of strategic capacity withholding in an attempt to influence market prices to its own advantage, i.e., in an attempt to exercise *market power*. The seller can engage in *economic capacity withholding* (reporting supply offers with higher-than-true marginal costs), and/or it can engage in *physical capacity withholding* (reporting supply offers with lower-than-true maximum generation capacities). Economic and physical capacity withholding are the two main ways in which real-world energy sellers can exercise market power. Consequently, it is of interest to energy market operators, for market power mitigation purposes, to understand which form of market power affords greatest advantage to energy sellers. Economic capacity withholding is relatively easy to monitor, to the extent that a seller's fuel type gives a strong indication of its true marginal costs. Strategic physical capacity withholding can be difficult to distinguish from outages and other events that cause unintentional reductions in available generation capacity.

Systematic computational experiments are then conducted to explore the emergence of *co-learning patterns*, that is, global market patterns that arise and persist over time as a result of seller co-learning. The specific co-learning patterns of interest here are correlated seller supply offer behaviors and correlated seller net earnings outcomes.

One key finding is that learning strongly matters in our double-auction environment. Sellers that behave as Gode and Sunder (1993,1997) budget-constrained zero-intelligence agents, randomly selecting their supply offers subject only to a break-even constraint, tend to realize substantially lower net earnings than sellers that tacitly co-learn to correlate their supply offers for market power advantages. A second key finding is that learning-to-learn strongly matters. The co-learning sellers perform much better when the parameters of their learning methods are calibrated to sweet-spot values reflecting the attributes of their particular decision environment, including both own attributes (e.g., size, cost function, and location) and rival seller attributes. A third key finding is that the pure exercise of economic capacity withholding is typically much more profitable for sellers than any use of physical capacity withholding.

A number of previous electricity researchers have separately explored either economic capacity withholding or physical capacity withholding exercised by learning traders, including the current authors. For example, Li and Tesfatsion (2009a) conduct preliminary learning experiments focusing on seller physical capacity withholding. Li, Sun, and Tesfatsion (2008,2009) explore the emergence of spatially correlated price patterns supported by seller co-learning when sellers can learn to exercise economic capacity withholding. Li and Tesfatsion (2011) explore the effects of seller co-learning on total net surplus (efficiency) and the distribution of surplus among sellers, buyers, and the ISO when sellers can learn to exercise economic capacity withholding.

The only previous work we are aware of that permits learning traders to engage simultaneously in both economic and physical capacity withholding is Tellidou and Bakirtzis (2007). The latter authors analyze an electricity market operating over a 2-bus transmission grid in which seller supply offers take the form of an offered quantity and an offered price. The offered quantity can be less than or equal to the seller's true maximum generation capacity, and the offered price can be greater than or equal to the seller's true reservation price. However, the authors do not undertake any comparative analysis to determine the relative advantages to sellers of the two forms of market power exercise. Moreover, all sellers are assumed to use the same identically parameterized learning method.

With regard to the general economics literature, it is rare to see physical capacity withholding treated at all.[1] This could be due, in part, to the analytical complications that arise when physical capacity withholding leads to binding capacity constraints. It could also be due to the folk belief that, when it comes to the exercise of market power, economic and physical capacity withholding are essentially equivalent means. When physical capacity withholding is considered, it is typically within game contexts in which the focus is on the existence of Nash equilibria without consideration of learning capabilities [e.g., Dechenaux and Kovenock (2007)].

We begin our study in Section 2 with a summary discussion of emergence as it has previously been defined and used for economic systems. A key conclusion from this section is that the concept of weak emergence developed by Bedau (1997) is particularly relevant for the study of real-world economic systems – such as electric power markets – whose complex interweaving of physical constraints, institutional rules, and strategic human behaviors renders them analytically intractable. Roughly, Bedau defines a property P of a system to be *weakly emergent* if P can be systematically generated for the system through a finite simulation, but through no other means.

Section 3 presents our wholesale electric power market model. This model is implemented by means of the AMES Wholesale Power Market Testbed [Li and Tesfatsion (2009b,c), Tesfatsion (2010)], an agent-based computational laboratory that incorporates institutional and structural features characterizing actual U.S. wholesale electric power markets. In keeping with actual practice, AMES implements a two-settlement system consisting of a forward day-ahead market and a real-time balancing market that operate over a high-voltage alternating current (HVAC) transmission grid. The day-ahead market is organized as a double auction in which wholesale buyers submit daily demand bids to buy energy, wholesale sellers submit daily supply offers to sell energy, and "locational marginal prices" are determined locally (for each hour at each grid bus) to maximize total net surplus subject to transmission and generation constraints. Traders in AMES can be modeled as learning agents who adjust their demand bids and supply offers over time in an attempt to exercise market power.

In Section 4 we explain the experimental design used to test for the (weak)

[1]For example, firm behavior with potentially binding production capacity constraints is only considered within one relatively small section (pp. 211-234) of the well-known 479-page industrial organization textbook by Tirole (2003) used in graduate and advanced undergraduate economics courses.

emergence of two types of co-learning patterns in our market model: namely, correlated seller supply offer behaviors, and correlated seller net earnings outcomes. In particular, we develop a series of test cases for a 5-bus wholesale electric power market in which the exercise of seller market power takes one of three forms: economic capacity withholding; physical capacity withholding; or some combination of the two.

Section 5 explains steps taken prior to our test-case experimentation to calibrate each seller's learning parameters to sweet-spot values attuned to each seller's actual decision environment. For example, each seller's initial aspiration level for net earnings is calibrated to its actual net earnings opportunities as structurally determined by its feasible supply offers in relation to its true marginal cost function. Heat maps are used to display and interpret these sweet-spot patterns. A *heat map* is a two-dimensional graphical depiction of data in which groups of associated data values are distinguished from one another by distinct colorings.

Sections 6-8 report our test-case experimental findings regarding the emergence of two types of co-learning patterns: correlated seller supply offer behaviors, and correlated seller net earnings outcomes. Heat maps are used to display and interpret these correlations. These heat map depictions can be viewed as extensions of traditional industrial organization measures for (strategic) substitution and complementarity defined in terms of the signs of (cross) partial derivatives evaluated at a point in time [Bulow et al. (1985)].[2] In the present context, which involves repeated stochastic choice by learning profit-seeking traders embedded in an interaction network operating over a physical network, more comprehensive ways are needed to measure the effects of one trader's actions on the actions and outcomes of other traders.

The key findings from the test-case experiments reported in Sections 6-8 are summarized and compared in Section 9. Concluding remarks are given in Section 10. For improved expositional clarity, technical materials regarding seller cost and net earnings functions, seller learning methods, and sweet-spot learning parameter calibrations are gathered together in appendices to this study.

[2] More precisely, a firm A's product is said to be a substitute (or complement) for firm B if more "aggressive" action by firm A, measured by an increase in some variable S^A, results in a decrease (increase) in the profits π^B of firm B. *Strategic* substitutes (complements) are defined in terms of the effects of a change in S^A on the *marginal* profitability of firm B, that is, in terms of the sign of $\partial^2\pi^B/\partial S^B\partial S^A$.

2. Emergence in Economic Systems

This section briefly reviews how emergence has previously been defined and used for the study of economic systems.[3] Of special interest are the complications caused by the presence of one or more agents with learning capabilities.

Kuperberg (2006) provides an historical overview of emergence conceptions in the economics literature. He begins in 1976 with the "invisible hand" of Adam Smith (1937). He then proceeds to a discussion of the "micromotives of macrobehavior" ideas of Schelling (1978) and the work of modern-day economists such as Alan Kirman (1993) espousing "a new kind of economics" now known as agent-based computational economics [Tesfatsion and Judd (2006)].[4] Kuperberg concludes there is no universally agreed upon definition of emergence, yet three core characteristics can be identified. First, there must be at least two distinct levels of organization. Second, at the lower level of organization, a multitude of individual agents operate in accordance with rules. Third, aggregate outcomes occur at the higher level of organization that result from the interactions of these individual agents but that are not easily derivable from the rules followed by the individual agents.

Harper and Endres (2010) focus on conceptions of emergence in the modern economics literature. They conclude (page 3) that this literature reveals "an incomplete patchwork of fragmented and contradictory notions of emergence." Table 1 (page 6) of their study homes in on differences in usage between evolutionary-institutional economists and complexity economists who view economies as dynamic systems of interacting components (agents, units, entities, ...). Each specific usage captures a potentially important facet of messy economic reality: for example, novelty, non-reducibility of wholes to their parts, downward causation, self-organization (bottom-up growth), recurrence of regular patterns, and unpredictability.

[3]Detailed discussions examining historical and current controversies surrounding the concept of emergence as used in a variety of disciplines can be found in Anderson (1972), Auyang (1998), Stephan (1999), Corning (2002), O'Connor and Wong (2006), and Dessalles et al. (2008). Discussions of emergence for general social science systems can be found in Gilbert (1995), Hodgson (2000), and Squazzoni (2006).

[4]Somewhat surprisingly, Kuperberg makes no mention of Hayek (1948), whose recognition that markets can be "spontaneously ordered" through the coordination capabilities of price mechanisms surely represents a relatively early identification of an emergent property for economic systems [Lewis (2011)].

Referring to works by Stephan (1998) and Corning (2002), among others, Harper and Endres identify core features commonly seen in delineations of emergence for economic systems which they suggest should be viewed as minimally necessary for emergence. Briefly, in order for a pattern, property, or relation to be emergent for an economic system, it must arise from material components and depend upon the organization (connections and interactions) of these components, not simply upon individual component properties. Harper and Endres then identify additional features particularly relevant to evolutionary economics which, together with the core features, constitute a stronger form of emergence: namely, genuine novelty; unpredictability in principle; and irreducibility to system component properties either in isolation or in other simpler systems.

What are the implications for emergence when one of more "constitutent parts" of a system are agents with learning capabilities? The general discussions of potential downward causation (from macroscopic to microscopic levels) appearing in some articles [e.g., Section 5 of Lewis (2011)] clearly relate to this issue. However, the issue is more directly addressed in the literature surveyed by Dessalles et al. (2008) on emergence in multi-agent systems. This literature explicitly recognizes that cognitive agents can be constrained "by the whole" because perceived social phenomena can affect their behaviors and decisions. Dessalles et al. (2008) also provide their own thoughtful discussion of implications for emergence when agents are cognitive observers of the world within which they interact and understand to various extents their collective ability to affect social outcomes.[5]

As elaborated in Borrill and Tesfatsion (2011), real people combine constructive and non-constructive aspects; they can only acquire new data about the world constructively, through interactions, but they can have possibly uncomputable beliefs about the world that influence their interactions. These beliefs about the world can arise from inborn attributes, from communications received from other agents, or from the use of non-constructive methods

[5]There is a strand of research not covered by Dessalles et al. (2008) that focuses on the emergence of social conventions in multi-agent systems with co-learning agents; see, for example, Shoham and Tennenholtz (1994), Kittock (1995), Delgago (2002), and Urbano et al. (2009) However, the focus of this literature is on network topology: specifically, how do different interaction networks (e.g., scale-free versus small-world networks) affect the speed with which conventions emerge. The relationship between co-learning and emergence is not specifically addressed.

(e.g., proof by contradiction) to interpret acquired data. To the extent that these beliefs involve perceptions of systemic properties, they provide a channel through which systemic properties can act back upon microscopic states. Interestingly, agent-based modeling permits decision-making agents to be learners who form both constructive and non-constructive beliefs about their virtual computational world as guides for their interactions, and whose interactions in turn inform their beliefs. These powerful modeling capabilities should greatly facilitate the study of emergence in real-world social systems.

In all of the emergence studies mentioned above, the focus is on the emergence of some pattern, property, or relation within a structurally given dynamic system. As will be concretely demonstrated in the remaining sections of this study, emergence in dynamic systems can also usefully be studied at a higher level using structural perturbation methods.

Specifically, structural perturbation methods similar to those used to study chaotic processes and fractal systems [Devaney and Keen, 1989] can by used to generate an ensemble of possible dynamic system trajectories in trajectory space. By an appropriate coloring of these trajectories (or of their pre-images in a parameter domain), interesting patterns can sometimes be revealed that provide a unique and informative "fingerprint" for the dynamic system.[6]

As a simple economic illustration, consider a discrete-time dynamic system that generates a unique net earnings level for each of N market traders over time, starting from an exogenously given system state z_1 in time period 1. Suppose the system structure depends on a parameter α. The distribution of net earnings for the N traders at any given time $t \geq 1$ would typically be represented by some form of histogram that associates each possible net earnings level with a frequency of occurrence in the population. Alternatively, however, this distribution can be represented as a single point $p_t(z_1, \alpha)$ in an N-dimensional space, where each coordinate of $p_t(z_1, \alpha)$ gives the net earnings level of a particular trader in period t, conditional on z_1 and α. One

[6]A famous example of this is the Mandelbrot set; see Branner (1998). The Mandelbrot set M is the set of all points c in the complex plane $\mathcal{C}$ for which the sequence $(z_t(c))$ in $\mathcal{C}$ does *not* diverge to ∞ as t tends to ∞, where $z_1(c) = 0$ and $z_{t+1}(c) = [z_t(c)]^2 + c$ for $t \geq 1$. The intricately beautiful irregular regularity of M is revealed by (i) coloring black all c-points lying in M, (ii) roughly partitioning the c-points in $\mathcal{C}$ lying outside M into finitely many subsets in accordance with the differential divergence rates of $z_t(c)$, and (iii) using (non-black) colors to differentially visualize the c-points lying in these subsets.

can then consider the loci of points traced out by $p_t(z_1, \alpha)$ in N-space as α is systematically varied over a feasible range of α-values, or the trajectory (orbit) traced out by $p_t(z_1, \alpha)$ in N-space as t is varied from 1 to ∞, or various other tracings resulting from individual or combined changes in t, z_1, and α.

In what sense do the patterns revealed by differential colorings of these tracings represent emergent patterns? Here we make recourse to Bedau (1997) to argue that these patterns can be emergent in a weak sense for dynamic models attempting to capture the salient characteristics of complicated real-world systems.

The concept of weak emergence was introduced by Bedau (1997) in an attempt to obtain a well-defined, practical, and scientifically relevant definition of emergence that sidesteps difficult philosophical and practical issues associated with definitions involving stronger requirements. Specifically, Bedau defines a macrostate P of a system S with a microdynamic D, initial conditions C, and possibly additional external conditions E to be *weakly emergent* for S if and only if P can be derived from $\{D, C, E\}$, but only by means of a finite simulation.

Bedau further notes that this core concept of weak emergence, restricted to a given system macrostate P, can be extended in a natural way to characterize the weak emergence of *system patterns*, i.e., collections of suitably related system macrostates. As stressed by Laughlin et al. (2000) and Laughlin (2000), even when the underlying microscopic properties and relationships for a dynamic system are not fully understood, the hope would be to find higher-level organizational principles that reliably associate collections of related microscopic states to collections of related macroscopic states, thus allowing some form of quantitative analysis to proceed.

Relating this back to earlier discussion, one way to visually differentiate among Bedau's distinct system patterns would be through the use of distinct colorings for the collections of system macrostates constituting these patterns. Heat maps could then be obtained by projecting the colored macrostates into various two-dimensional subspaces of interest. These heat maps could help to elucidate complex relationships between structural, institutional, and behavioral conditions and the appearance and persistence of system patterns. In the remaining sections of this study we illustrate how heat-maps can be used to visualize the weak emergence of system patterns for the concrete case of a wholesale electric power market.

- **Traders**
  - LSEs (bulk-power buyers)
  - GenCos (bulk-power sellers with learning capabilities)
- **Independent System Operator (ISO)**
  - Day-ahead hourly scheduling via bid/offer-based DC optimal power flow (OPF)
  - System reliability assessments
- **Two-settlement process**
  - Day-ahead market (double auction, financial contracts)
  - Real-time market (settlement of differences)
- **AC transmission grid**
  - LSEs and GenCos located at user-specified buses across the transmission grid
  - Congestion managed via locational marginal pricing

Figure 1: Key Features of the AMES Wholesale Power Market Testbed.

3. AMES Wholesale Power Market Testbed: Overview

The wholesale electric power market model used in this study is implemented using the AMES software platform. AMES (*A*gent-based *M*odeling of *E*lectricity *S*ystems) is an open-source wholesale power market testbed developed entirely in Java by researchers at Iowa State University. The latest version of AMES can be freely downloaded either at the AMES homepage [Tesfatsion (2010)] or through the IEEE Task Force on Open Source Software [IEEE (2011)].

This section describes the key features of Version 2.05 of AMES, used in this study.[7] These key features reflect, in simplified form, actual U.S. wholesale electric power market operations in the Midwest (MISO), New England (ISO-NE), New York (NYISO), the mid-Atlantic states (PJM), California (CAISO), Texas (ERCOT), and the Southwest (SPP). A summary listing of these key features is provided in Fig. 1.

The AMES(V2.05) wholesale power market operates over a high-voltage alternating current (HVAC) transmission grid starting with hour 00 of day 1 and continuing through hour 23 of a user-specified maximum day. AMES includes an *Independent System Operator (ISO)* that manages market op-

[7] Technical details are relegated to appendices. Appendix A provides quantitative descriptions for seller cost and net earnings functions, and Appendix B presents the precise quantitative form of the learning method used by the sellers to update their daily supply offers.

erations and a collection of energy traders distributed across the grid who buy or sell power at wholesale. The wholesale buyers service the energy demands (load) of retail consumers and are referred to as *Load-Serving Entities (LSEs)*. The wholesale sellers are energy producers referred to as *Generation Companies (GenCos)*.

The objective of the not-for-profit ISO is the maximization of *Total Net Surplus (TNS)* subject to branch capacity limits, GenCo generation capacity limits, and balance constraints.[8] As detailed in Li and Tesfatsion (2011), TNS is the sum of GenCo net surplus, LSE net surplus, and ISO net surplus. In an attempt to attain this objective, the ISO operates a day-ahead energy market settled by means of *Locational Marginal Prices (LMPs)*, the pricing of electric power in accordance with the timing and location of its injection into, or withdrawal from, the transmission grid.[9]

The welfare of each LSE is measured by the net earnings it secures for itself through the purchase of power in the day-ahead market and the resale of this power to its retail customers. During the morning of each day D, each LSE reports a *demand bid* to the ISO for the day-ahead market for day D+1. Each demand bid consists of two parts: fixed demand (i.e., a 24-hour load profile) to be sold downstream at a regulated price r to its retail customers with fixed-price contracts; and 24 price-sensitive inverse demand functions, one for each hour, reflecting the price-sensitive demand (willingness to pay) of its retail customers with dynamic-price contracts.[10]

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 s_i^R 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.[11]

[8]For technical reasons, power injections into a grid (supply) must at all times be in balance with power withdrawals (demands plus losses) to maintain grid stability.

[9]The price LMP_{kt} at bus k for time t is determined as the shadow price of the balance constraint at bus k for time t.

[10]The LSEs in AMES(V2.05) have no learning capabilities; LSE demand bids are user-specified at the beginning of each simulation run. However, as explained more carefully in Li and Tesfatsion (2009b,c), AMES(V2.05) includes a learning module, JReLM, that can be used to implement a wide variety of stochastic reinforcement learning methods for decision-making agents. Extension to include LSE learning is planned for future AMES releases.

[11]Whether GenCos are permitted to report only one supply offer or 24 supply offers

Each supply offer s_i^R in AD_i takes the form of a reported linear marginal cost function (characterized by a reported ordinate a_i^R and a reported slope $2b_i^R$) defined over a production capacity interval spanning the range from 0 to a reported maximum generation capacity Cap_i^{RU}. GenCo i's ability to vary its choice of a supply offer s_i^R from AD_i permits it to adjust the ordinate/slope of its reported marginal cost function and/or the upper limit of its reported generation 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 posts hourly bus LMP levels as well as LSE cleared demands and GenCo dispatch levels for the day-ahead market for day D+1. These hourly outcomes are determined via *Security-Constrained Economic Dispatch (SCED)* formulated as bid/offer-based *DC Optimal Power Flow (DC-OPF)* problems with approximated TNS objective functions based on reported rather than true GenCo costs.[12]

At the end of each day D the ISO settles the day-ahead market for day D+1 by receiving all purchase payments from LSEs and making all sale payments to GenCos based on the LMPs for the day-ahead market for day D+1, collecting any difference as *ISO net surplus*. As explained and demonstrated in Li and Tesfatsion (2011), this ISO net surplus is guaranteed to be nonnegative and, under congested grid conditions, will typically be strictly positive due to the separation of bus LMPs.

As detailed in Appendix B, each GenCo i at the end of each day D uses a variant of a well-known "stochastic reinforcement learning" method due to Roth and Erev (1995) 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"). Roughly described, if GenCo i's supply offer on day D results in a relatively good reward, GenCo i increases

---

for use in the day-ahead energy market varies from one energy region to another. For example, the ISO-NE permits only one supply offer whereas MISO permits 24 separate supply offers. Baldick and Hogan (2002) suggest that imposing limits on the ability of GenCos to report distinct hourly supply offers could reduce their ability to exercise market power, a conjecture that would be interesting to put to a test.

[12]A technical presentation of the bid/offer-based DC-OPF problem formulation for the ISO in AMES(V2.05) is provided in Li and Tesfatsion (2011). When demand is 100% fixed (price insensitive), the objective of maximizing TNS is equivalent to the objective of minimizing total GenCo avoidable costs of operation; 100% fixed demand is the case treated in all experiments reported in this study.

the probability it will choose to report this same supply offer on day D+1, and conversely. Hereafter this learning method is referred to as the *Variant Roth-Erev (VRE) learning method*.

There are no system disturbances (e.g., weather changes) or shocks (e.g., line outages). Consequently, the dispatch levels determined on each day D for the day-ahead energy market for day D+1 are carried out as planned without need for settlement of differences in the real-time energy market for day D+1.

4. Experimental Design

All of the experiments reported in this study were conducted using a 5-bus wholesale electric power market model based on a transmission grid configuration developed by Lally [2002] that is now commonly used in many ISO-managed U.S. energy regions for training purposes. Our main goal is to implement an experimental design within this 5-bus framework that permits us to explore how systematic variations in the ability of the GenCos to exercise market power result in systematic variations in GenCo reported supply offers and GenCo net earnings outcomes that, when visualized through heat maps, reveal interesting correlation patterns.

Seller market power is exercised in wholesale electric power markets in two possible ways: *economic capacity withholding* (offering energy at higher-than-true marginal cost); and *physical capacity withholding* (offering lower-than-true maximum generation capacity). Consequently, our experimental design encompasses four types of test cases: (i) a benchmark test case in which the GenCos have no learning capabilities and always report their true cost and capacity attributes to the ISO; (ii) test cases in which the GenCos use VRE learning to strategically report higher-than-true marginal costs to the ISO but always report their true maximum generation capacities to the ISO; (iii) test cases in which the GenCos use VRE learning to report lower-than-true maximum generation capacities to the ISO but always report their true marginal costs to the ISO; and (iv) test cases in which the GenCos use VRE learning to report higher-than-true marginal costs and/or lower-then-true maximum generation capacities to the ISO.

A more detailed description of these test cases is provided below.

4.1. Benchmark Five-Bus Test Case: No Learning

Complete input data for our benchmark 5-bus test case are provided in the input data file for the 5-bus test case included in the data directory of the AMES(V2.05) download available at the AMES homepage [Tesfatsion (2010)]. Briefly summarized, this benchmark case has the the following structural, institutional, and behavioral features.

The wholesale electric power market for our benchmark case operates over a 5-bus transmission grid as depicted in Fig. 2, with branch attributes (e.g., thermal limits), locations of LSEs and GenCos, and initial hour-0 LSE demands adopted from Lally (2002). The Lally grid configuration has proved to be highly useful in practice for ISO training purposes because it is small enough to be manageable while still retaining many important real-world features. For example, the Lally grid is connected yet not completely connected (i.e., not every pair of grid busses is connected by a branch), which has important ramifications for the physical flow of power across the grid. The branch 1-2 connecting Bus 1 to Bus 2 has a thermal limit and hence is vulnerable to overload (congestion). The LSE demands (loads) are concentrated in a *load pocket* at Busses 2, 3, and 4, which gives GenCo 3 located at Bus 3 market power advantages for the servicing of this load whenever branch 1-2 is congested. Finally, LSE demands are 100% fixed (no price sensitivity), an empirically accurate reflection of the highly inelastic demand currently characterizing U.S. wholesale electric power markets.

The Lally configuration was primarily designed for point-in-time reliability studies, not for dynamic economic studies. Hence, GenCo marginal costs (assumed constant) and LSE demands (loads) are given for one point in time, and strategic learning behaviors are not considered.

Our benchmark case extends the Lally configuration by specifying GenCo marginal cost functions as depicted in Fig. 3. The GenCos range from GenCo 5, a relatively large coal-fired baseload unit with low marginal operating costs, to GenCo 4, a relatively small gas-fired peaking unit with relatively high marginal operating costs. Moreover, our daily LSE 24-hour fixed demand profiles are adopted from a case study presented on pages 296–297 in Shahidehpour et al. (2002). Hourly fixed demand varies between low (hour 4:00) and peak (hour 17:00) each day.

However, for our benchmark case we retain the Lally (2002) presumption that GenCos are non-learners. Specifically, we assume the GenCos report supply offers to the ISO for the day-ahead energy market that convey their true marginal cost functions and true maximum generation capacities.

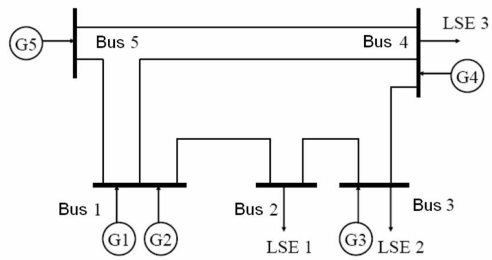

Figure 2: Transmission grid for the benchmark 5-bus test case.

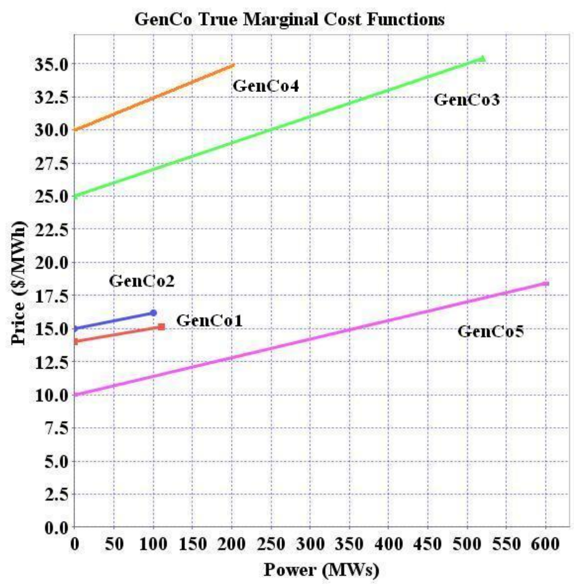

Figure 3: GenCo true marginal cost functions and true capacity attributes for the benchmark 5-bus test case.

As it turns out, during a typical day D for our benchmark 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 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. 2, 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. 3,

Table 1: Hourly GenCo net earnings during a typical 24-hour day D for the benchmark 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

these two GenCos have relatively small generation 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, and hence the net earnings of GenCo 3.

4.2. Five-Bus Test Cases with GenCo Learning

Each of the test cases with GenCo learning extends the benchmark 5-bus test case described in Section 4.1 by permitting one or more GenCos to use VRE learning to exercise either economic capacity withholding, or physical capacity withholding, or combinations of the two. Complete input data for these test cases (including initial random seed values) are provided in the input data file for the 5-bus test case included in the data directory of the AMES(V2.05) download available at the AMES homepage [Tesfatsion (2010)].

All other aspects of these learning test cases are the same as for the benchmark 5-bus test case. In particular, GenCo net earnings are used to evaluate GenCo welfare, and the hourly energy demand bids of the LSEs are fixed quantities that are insensitive to price.

The treatment factor for *economic* capacity withholding experiments is whether or not each GenCo can learn to exercise economic capacity withholding by reporting higher-than-true marginal costs in its supply offers. The treatment factor for *physical* capacity withholding experiments is whether or not each GenCo can learn to exercise physical capacity withholding by reporting lower-than-true maximum generating capacities in its supply offers. For *combined* economic and physical capacity withholding experiments, the treatment factor is whether or not each GenCo can learn to report higher-than-true marginal costs and/or lower-than-true maximum generating 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". For each learning treatment these thirty seed values are used to implement thirty distinct runs, each 1000 simulated days in length.

5. Sweet-Spot Calibration of GenCo Learning Parameters

5.1. Motivation for VRE Learning Calibration

In actual U.S. wholesale electric power markets, relatively small numbers of profit-seeking GenCos repeatedly make supply offers in an attempt to secure good net earnings. It seems reasonable to assume that these GenCos are able to adjust their learning methods over time to their particular decision environments, that is, that they are able to *learn-to-learn*.

Consequently, prior to conducting our capacity withholding experiments with learning GenCos, we first undertook intensive parameter sensitivity studies in an attempt to determine *sweet-spot* values for each GenCo's VRE learning parameters. Specifically, we attempted to determine values for these learning parameters capable of yielding relatively high GenCo net earnings in our 5-bus test case experiments by selecting values that in fact yielded relatively high GenCo net earnings in our calibration experiments.

In this section we briefly report on these calibration experiments, relegating technical details to Appendix C. In particular, we highlight several interesting implications regarding the importance of learning and learning-to-learn in complicated decision environments as exemplified by our current wholesale electric power market setting.

5.2. Calibration of VRE Learning Parameters

As detailed in Appendix B, the VRE learning method used to implement learning for each GenCo i depends on four key parameters: ($\mathrm{q}_i(1)$, T_i, r_i, e_i). We first briefly summarize the role of each parameter in the learning process.

The *initial propensity level* $\mathrm{q}_i(1)$ is a measure of GenCo i's net earnings aspirations at the beginning of the initial day 1, which the VRE learning method then successively updates in an action-conditioned manner as GenCo i successively selects new actions (supply offers) and new rewards (own-net earnings outcomes) are realized. After each updating, these action-conditioned propensities are transformed into action-conditioned probability assessments which GenCo i uses to select its next supply offer to report to the ISO. The *temperature* parameter T_i (which has nothing to do with actual weather) enters into the mapping from propensities to probabilities in a manner that affects the extent to which GenCo i experiments with new actions, particularly in early learning stages. Larger ("hotter") values of T_i encourage increased experimentation.

The *recency parameter* r_i enters into the propensity updating relationship; higher values of r_i dampen the rate at which GenCo i's action-conditioned propensities change over time. The *experimentation parameter* e_i also enters into the propensity updating relationship; higher values of e_i permit the reinforcement effects of rewards to spill over in greater proportion from chosen to non-chosen actions, thus encouraging GenCo i to experiment across a broader range of actions.

For the calibration of the recency and experimentation parameters (r_i,e_i), we relied on the sensitivity findings reported by Pentapalli (2008) for 3-bus and 5-bus wholesale electric power market experiments. For the calibration of the initial propensity and temperature parameters ($\mathrm{q}_i(1)$,T_i), we used 5-bus experiments to determine GenCo net earnings for a range of positive values for α_i and β_i, defined as follows:

- $\alpha_i = \mathrm{q}_i(1)/\mathrm{MaxDNE}_i$, where MaxDNE_i denotes GenCo i's (positive) estimated maximum possible daily net earnings as determined structurally from its action domain of feasible supply offers and its true marginal cost function;

- $\beta_i = q_i(1)/T_i$.

5.3. Illustrative Calibration Findings

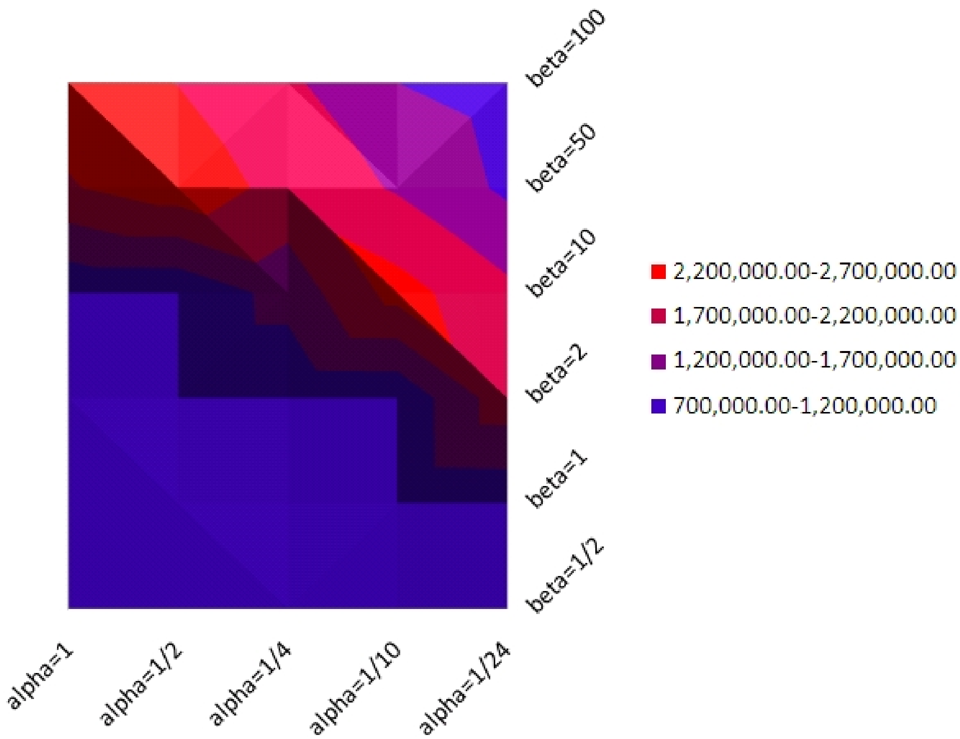

Figure 4: Economic capacity withholding calibrations: A 2D heat-map depiction of mean outcomes for total GenCo daily net earnings on day 1000 for the 5-bus test case extended to permit all five GenCos to use VRE learning to exercise economic capacity withholding. Outcomes are depicted for a range of settings for the two key VRE learning parameters α and β set commonly across all five GenCos.

The specific sweet-spot settings for ($q_i(1)$, T_i, r_i, e_i) selected for each GenCo i for use in each of our 5-bus test case experiments are explained in Appendix C. In this section we report and interpret illustrative findings for some of our calibration experiments, using heat maps to visualize the mapping between learning parameter settings and GenCo net earnings outcomes.

Figure 4 depicts calibration-experiment findings for mean total GenCo daily net earnings attained for a 5-bus case in which each GenCo is a VRE learner able to exercise pure economic capacity withholding. The findings are generated for a range of values for α and β, commonly set across all GenCos. An interesting pattern is immediately evident: namely, the (α,β) combinations associated with the highest mean net earnings outcomes lie along a nonlinear ridge line that traverses from (1, 100) in the northwest corner to (1/24, 2) in the south-central region. What causes this nonlinear coupled dependence of mean net earnings outcomes 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., with different supply offers (reported marginal cost functions). 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. In this case the behavior of the GenCos corresponds to Gode-Sunder (1993,1997) budget-constrained zero-intelligence (ZI-B) market sellers who randomly select supply offers subject only to a budget constraint. That is, each GenCo randomly chooses supply offers from its action domain, which by construction only includes possible reported marginal cost functions that lie on or above the GenCo's true marginal cost function (thus enforcing a break-even constraint). The deleterious effect of this ZI-B GenCo behavior is seen in the uniformly low mean net earnings outcomes achieved in Fig. 4 for the lowest tested β levels 1/2 and 1.

The bottom line following from the calibration-experiment findings reported in Fig. 4 is that the learning representation for the GenCos definitely matters. GenCos achieve their highest mean net earnings along a nonlinear ridge-line of sweet-spot (α, β) values traversing from (1, 100) to (1/24, 2), and their lowest mean net earnings for extremely low β values that induce the GenCos to behave like Gode-Sunder ZI-B market sellers.

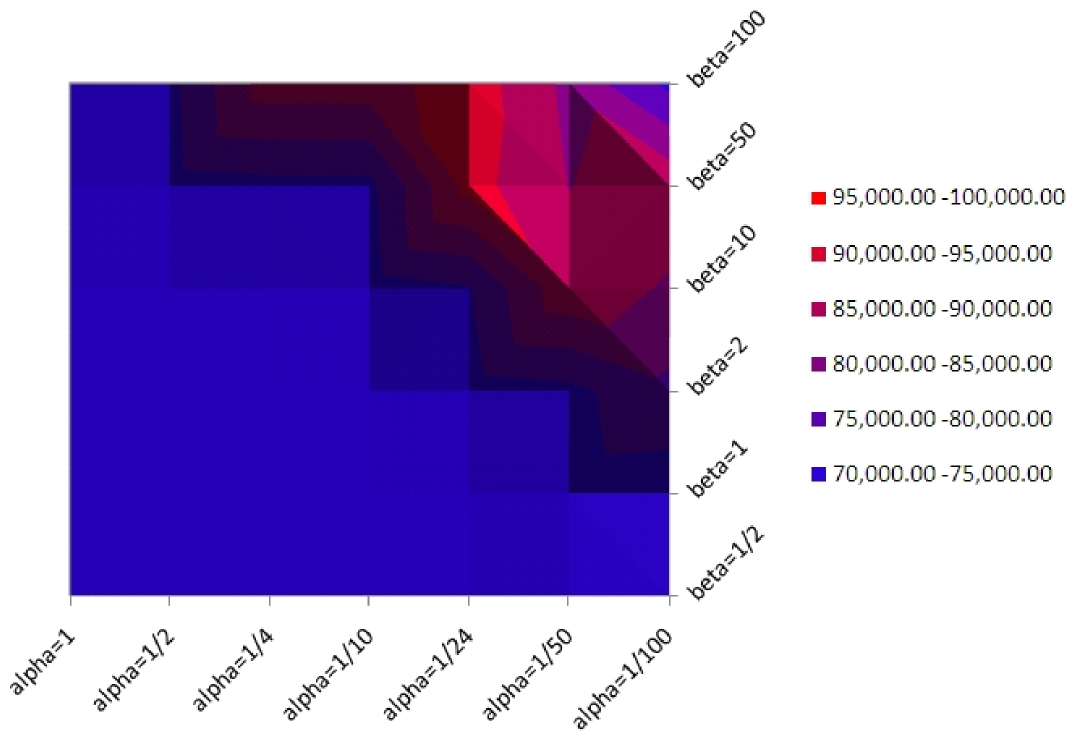

Figure 5: Physical capacity withholding calibration results for GenCo 3: A 2D heat-map depiction of mean outcomes on day 1000 for GenCo 3's daily net earnings for the benchmark 5-bus test case extended to permit GenCos 1, 3, and 5 to use VRE learning to exercise physical capacity withholding. Results are shown for a range of values for the VRE learning parameters (α, β), commonly set for GenCos 1, 3, and 5.

Figure 5 depicts calibration-experiment findings for mean GenCo 3 net earnings on day 1000 attained for a 5-bus case in which GenCos 1, 3, and 5 are VRE learners able to exercise pure physical capacity withholding. The findings are generated for a range of values for α and β, commonly set for GenCos 1, 3, and 5. The net earnings of GenCo 3 are chosen as the criterion for the calibration of (α, β) for our pure physical capacity withholding experiments because in most hours this large supplier turns out to be a *pivotal* supplier, i.e., its generation capacity is essential for meeting fixed demand.

Three interesting observations can be made about the results reported in Fig. 5. First, learning matters: the setting for (α, β) substantially affects mean GenCo 3 net earnings. Second, the sweet-spot (α, β) combinations associated with the highest mean GenCo 3 net earnings roughly lie along a vertical ridge line ranging from (1/24, 50) to (1/24, 100). Third, comparing the pure physical capacity withholding outcomes in Figure 5 with the pure economic capacity withholding outcomes in Figure 4, it is seen that the sweet-spot region for GenCo 3's (α, β) learning parameters strongly depends on the particular learning environment.

This third finding indicates the importance of learning-to-learn. No one setting for the parameters of a learning method can be expected to do well across all possible decision environments in which the learning method might be applied.

6. Test-Case Findings for Pure Economic Capacity Withholding

This section reports findings for two types of 5-bus test case experiments. The first type of experiment tests the extent to which a *single* 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* GenCos can co-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 is the extent to which the second type of experiment results in correlated supply offer selections and correlated net earnings outcomes for the two learning GenCos.

6.1. Economic Capacity Withholding by One Learning GenCo

GenCo 3 is selected as the sole learner for this first type of experiment because of the critical role it plays in the determination of locational marginal prices (LMPs). This critical role results for three reasons: (a) GenCo 3 has a relatively large generation capacity; (b) GenCo 3 is a pivotal supplier during peak (high) demand hours, meaning that its capacity is needed to meet LSE fixed demand; and (c) GenCo 3's true marginal costs of production are relatively high.

As carefully explained in Appendix A, each supply offer reported to the ISO by a learning GenCo takes the form of a reported linear marginal cost function with ordinate a^R and slope $2b^R$. This function is defined over a reported generation capacity interval ranging from 0 to a reported maximum generation capacity Cap^{RU}.

For the economic capacity withholding experiment at hand, GenCo 3 always reports its true maximum generation capacity. Consequently, GenCo 3's action domain AD_3 is specified to consist of 100 possible marginal cost functions corresponding to 100 different specifications for the ordinate and slope values (a^R,b^R), where each of these functions is defined over GenCo 3's true generation capacity interval. All other GenCos are assumed to be non-learners that report their true cost and capacity attributes to the ISO. The action domain for each of these other GenCos thus contains only one element: namely, this GenCo's true marginal cost function defined over its true generation capacity interval.

Table 2: Mean outcomes (with standard deviations) on day 1000 for GenCo daily net earnings (DNE) and GenCo 3's reported supply offers (a^R,b^R) when GenCo 3 uses VRE learning to exercise economic capacity withholding.

	No GenCo Learning	With GenCo 3 Learning
GenCo 1 $\overline{\mathrm{DNE}}$	1,556.41	0.00 (0.00)
GenCo 2 $\overline{\mathrm{DNE}}$	26.58	0.00 (0.00)
GenCo 3 $\overline{\mathrm{DNE}}$	56,016.09	1,699,368.20 (400,430.50)
GenCo 4 $\overline{\mathrm{DNE}}$	142.27	253,468.03 (72,720.50)
GenCo 5 $\overline{\mathrm{DNE}}$	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 2 reports experimental findings both for the benchmark no-learning case and for the case in which GenCo 3 uses VRE learning 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, the mean net earnings for GenCo 4 also substantially increase (complementarity effect) whereas the mean net earnings of GenCo 1 and GenCo 2 are reduced to zero (substitution effect), even though GenCos 4, 1, and 2 are not learning agents and hence always report their true cost and capacity attributes to the ISO.

The reason for these findings 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-demand 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.

As noted in the introduction, these complicated substitution and complementarity effects arising through network interactions are not well captured using traditional derivative measures [Bulow et al. (1985)].

6.2. Economic Capacity Withholding by Two Co-Learning GenCos

Two different pairs of co-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 with relatively low true marginal cost whereas GenCo 3 is a pivotal supplier during peak-demand 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 relatively large maximum generation 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 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 the net earnings of GenCo 1 decline to zero (substitution effect).

Interestingly, the net earnings and reported marginal cost results presented in Table 3 for the case in which GenCo 1 and GenCo 3 are co-learners are similar to the corresponding results reported in Table 2 for the case in which GenCo 3 is the sole learner. The reason for this is partly explained by the findings earlier discussed in Section 4.1 and Section 6.1. Due to the persistent congestion on the branch from Bus 1 to Bus 2, and to GenCo 3's relatively large generation capacity, GenCo 3 is a pivotal supplier in most hours, meaning that its 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 1-2 and it 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 levels, as well as for the determination of GenCo 3's reported supply offers.

Table 3: Mean outcomes (with standard deviations) on day 1000 for GenCo daily net earnings (DNE), and for reported supply offers (a^R,b^R) for GenCo 1 and GenCo 3, when both GenCo 1 and GenCo 3 use VRE learning 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 (true value)	26.7006 (12.8204)
GenCo 1 $\overline{b^R}$	0.0050 (true value)	0.1363 (0.2016)
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 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 (complementarity effect) compared to the benchmark no-learning case. GenCo 3's net earnings, in particular, dramatically increase.

The explanation for these findings is as follows. GenCo 5 is essentially a

Table 4: Mean outcomes (with standard deviations) on day 1000 for GenCo daily net earnings (DNE), and for reported supply offers (a^R,b^R) for GenCo 3 and GenCo 5, when both GenCo 3 and GenCo 5 use VRE learning 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 (true value)	92.8333 (12.3654)
GenCo 3 $\overline{b^R}$	0.0100 (true value)	0.2202 (0.0671)
GenCo 5 $\overline{a^R}$	10.0000 (true value)	18.6560 (10.3939)
GenCo 5 $\overline{b^R}$	0.0070 (true value)	0.0238 (0.0286)

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 4 to the Case (1) findings presented in Table 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 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 cutback 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).

7. Test-Case Findings for Pure Physical Capacity Withholding

In parallel with Section 6, this section considers two types of experiments. The first type of experiment tests the extent to which a *single* 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* GenCos can co-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. Of particular interest is the extent to which the second type of experiment results in correlated supply offer selections and correlated net earnings outcomes for the two learning GenCos.

The treatment factor for these experiments is the maximum possible shrinkage for reported maximum generation capacity that a learning GenCo can submit to the ISO, the interpretation being that any larger shrinkage would risk detection by the ISO. The tested ranges for this treatment factor are chosen to avoid *supply inadequacy*, i.e., the reporting of capacities that are insufficient to meet total fixed demand.

7.1. Physical Capacity Withholding by One Learning GenCo

In the experiments presented in this section, as in Section 6.1, only GenCo 3 has learning capabilities. Here, however, GenCo 3's learning is restricted to the ability to exercise *physical* capacity withholding. The only treatment factor is GenCo 3's *shrinkage value* $\mathrm{MPRMCap}_3$, i.e., the setting for GenCo 3's *minimum possible reported maximum capacity* that determines the lowest possible maximum generation capacity that GenCo 3 is able to report to the ISO. For example, given $\text{MPRMCap}_3 = 0.99$, GenCo 3's reported maximum generation capacity Cap_3^{RU} must be at least 99% of its true maximum generation capacity Cap_3^U In the experiments presented below, MPRMCap_3 is varied between 0.95 and 0.99. All other GenCos are assumed to report their true costs and capacities to the ISO.

Table 5: Mean outcomes (with standard deviations) on day 500 for GenCo daily net earnings (DNE), and for GenCo 3's reported maximum generation capacity values Cap^{RU}, when GenCo 3 uses VRE learning to exercise physical capacity withholding. Results are shown for a range of MPRMCap_3 shrinkage values for GenCo 3.

MPRMCap for GenCo 3	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 $\overline{\text{Cap}^{RU}}$ (TrueMaxCap=520)	515.99 (0.77)	512.86 (2.44)	505.42 (1.01)	502.02 (2.40)	495.49 (1.04)
GenCo 3 $\overline{\text{Cap}^{RU}/\text{TrueMaxCap}}$	99.23% (0.15%)	98.63% (0.47%)	97.20% (0.19%)	96.54% (0.46%)	95.29% (0.20%)

The findings presented in Table 5 show that GenCo 3 is able to subtantially 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 generation capacity. These increases in daily net earnings are at the expense of GenCo 1 and GenCo 2, whose daily net earnings decline as GenCo 3's withholding increases (substitution effect). On the other hand, GenCo 4 and GenCo 5 experience modest gains in daily net earnings from GenCo 3's withholding (complementarity effect).

Figure 6 depicts GenCo 3's *actual* reported maximum generation capacity Cap_3^{RU} versus the *optimal* value for its reported maximum generation capacity under a range of different MPRMCap_3 shrinkage values. More precisely, for each given MPRMCap_3 setting, the optimal reported value gives the best possible maximum generation 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; these optimal settings were determined by direct off-line search. Figure 6 shows that the mean maximum generation capacity value that GenCo 3 *learns* to report to the ISO by day 500 is close to optimal for each $MPRMCap_3$ setting.

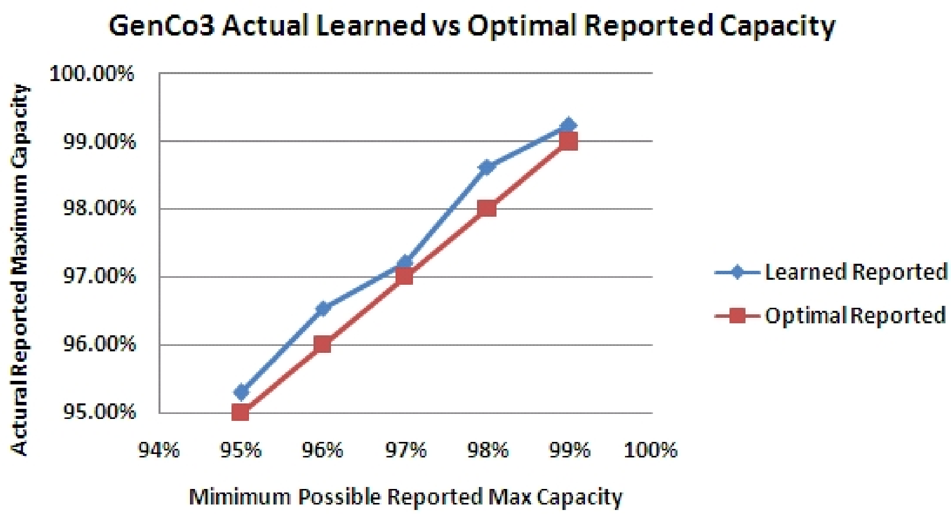

Figure 6: Mean outcomes on day 500 for learned versus optimal values for GenCo 3's reported maximum generation capacity values when GenCo 3 uses VRE learning to exercise physical capacity withholding. Results are shown for a range of $MPRMCap_3$ shrinkage values for GenCo 3.

7.2. Physical Capacity Withholding by Two Co-Learning GenCos

As in Section 6.2, 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 6.2 permits meaningful comparisons between learning experiments for economic versus physical capacity withholding.

The treatment factors for the Case (3) experiments are the MPRMCap shrinkage values for GenCo 1 and GenCo 3. The $MPRMCap_1$ setting for GenCo 1 is varied from 0.75 to 0.95, and the $MPRMCap_3$ setting for GenCo 3 is varied from 0.95 to 0.99. All non-learning GenCos report their true cost and capacity attributes to the ISO.

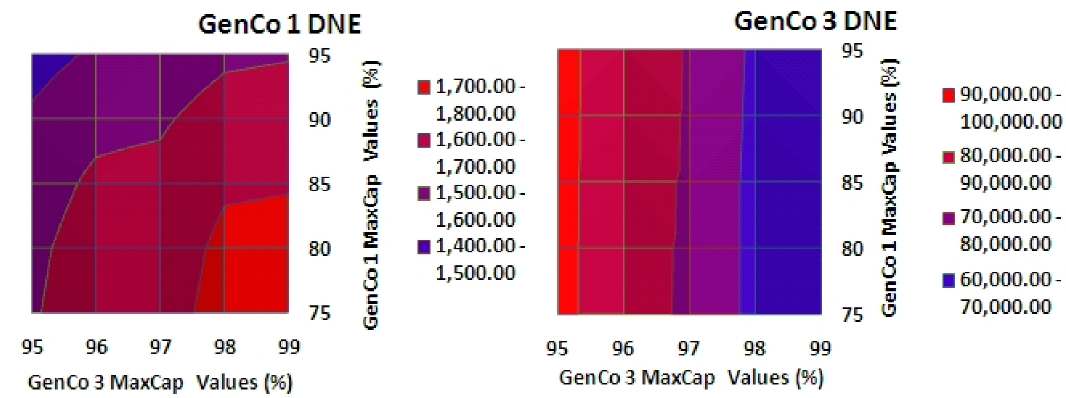

Figure 7: Daily net earnings (DNE) for GenCo 1 and GenCo 3 for the benchmark 5-bus test case extended to permit varied true maximum generation capacity values for GenCo 1 and GenCo 3. The latter values are depicted in percentage form (relative to original benchmark values).

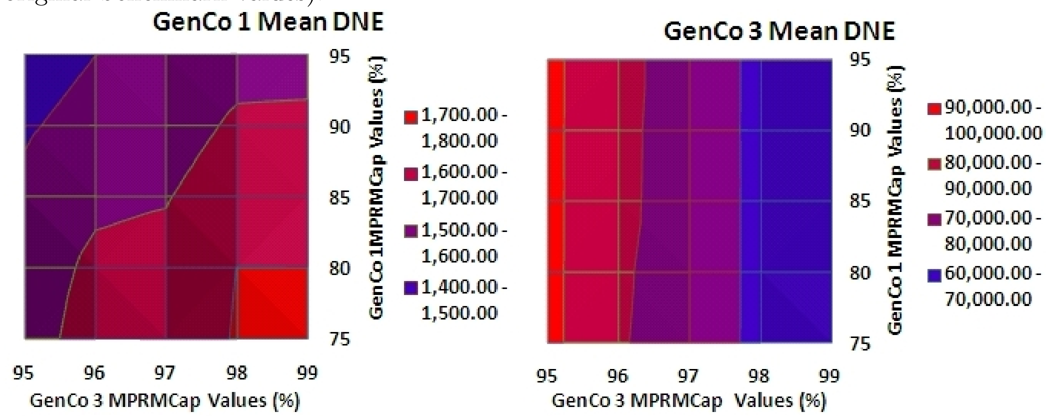

Figure 8: Mean daily net earnings (DNE) on day 500 for GenCo 1 and GenCo 3 when both GenCos use VRE learning to exercise physical capacity withholding. Mean DNE results for each GenCo are shown for various combinations of MPRMCap shrinkage values for the two GenCos.

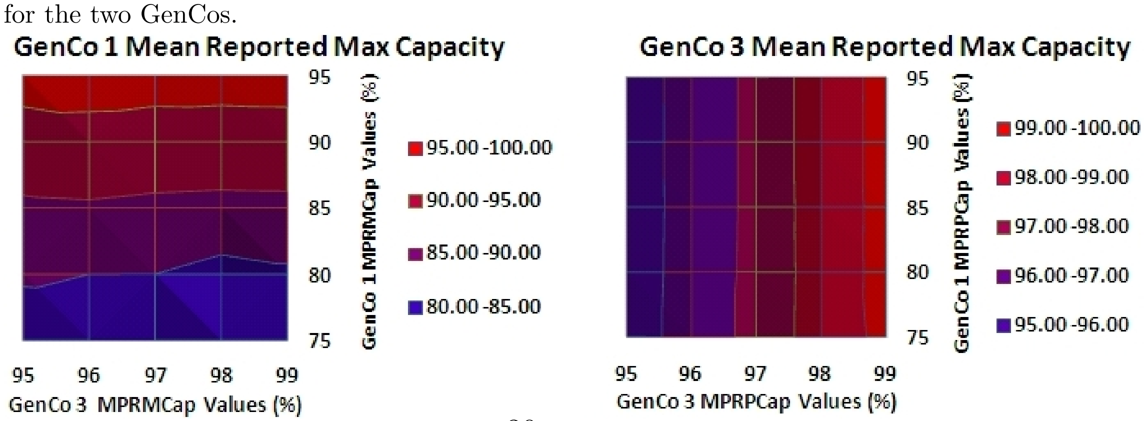

Figure 9: Mean reported maximum generation capacities (as a percentage of true maximum generation capacities) on day 500 for GenCo 1 and GenCo 3 when both GenCos use VRE learning to exercise physical capacity withholding. Mean reported maximum capacity results for each GenCo are shown for various combinations of MPRMCap shrinkage values for the two GenCos.

As a benchmark of comparison for Case (3), Figure 7 presents typical daily net earnings for GenCo 1 and GenCo 3 for the benchmark 5-bus test case (no learning) extended to permit a range of settings for the true maximum generation capacities of GenCo 1 and GenCo 3. From these findings it can be seen that GenCo 1, a relatively small unit, does best when its maximum generation capacity is 75% of its benchmark maximum generation capacity and the maximum generation capacity of the relatively large GenCo 3 is 99% of its benchmark maximum generation capacity. In contrast, GenCo 3 does best when its maximum generation capacity is 95% of its benchmark maximum generation capacity no matter what value is set for GenCo 1's maximum generation capacity.

Figure 8 presents mean daily net earnings for GenCo 1 and GenCo 3 on day 5000 when both GenCos use VRE learning 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_1 shrinkage value is set at its lowest tested level (0.75) and the MPRMCap_3 shrinkage value for GenCo 3 is set at its highest tested level (0.99). Comparing the left-hand side of Figure 8 to the left-hand side of Figure 7, it is also seen that the MPRMCap_1 region over which GenCo 1 attains its highest mean daily net earnings under learning is smaller than the maximum generation 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 7, it is seen in the right-hand side of Figure 8 that GenCo 3 attains its highest mean daily net earnings when its MPRMCap_3 shrinkage value is set at its lowest tested level (0.95). Also, the vertically-striped pattern for GenCo 3's mean daily net earnings indicates that GenCo 1's MPRMCap_1 settings have essentially no effect on the daily net earnings attained by GenCo 3.

Figure 9 displays mean reported maximum generation capacities (as a percentage of benchmark true maximum generation capacities) for GenCo 1 and GenCo 3 on day 500 when both these GenCos use VRE learning to exercise physical capacity withholding. From the left-hand side of the figure, it is seen that GenCo 1's mean reported maximum generation capacity is somewhat higher than its MPRMCap_1 shrinkage value 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 generation capacity choices have essentially no effect on the reported maximum generation capacity choices made by

GenCo 1. As indicated in the right-hand side of the figure, GenCo 3's mean reported maximum generation capacity is close to its $\mathrm{MPRMCap}_3$ shrinkage value 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 generation capacity results, GenCo 1's reported maximum generation capacity choices have essentially no effect on the reported maximum generation capacity choices made by GenCo 3.

In summary, in the Case (3) physical capacity withholding experiments in which only GenCo 1 and GenCo 3 are learners, 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 (complementarity effect). In contrast, GenCo 3's reported maximum generation capacity choices are essentially uncorrelated with the reported maximum generation capacity choices of GenCo 1.

For the Case (4) physical capacity withholding experiments in which only GenCo 3 and GenCo 5 are learners, the treatment factors are the MPRMCap shrinkage values for GenCo 3 and GenCo 5. The $\mathrm{MPRMCap}_3$ shrinkage value for GenCo 3 is varied from 0.95 to 0.99, and the $\mathrm{MPRMCap}_5$ shrinkage value for GenCo 5 is varied from 0.70 to 0.95. All non-learning GenCos report their true cost and capacity attributes to the ISO.

Figure 10 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 generation capacities expressed as a percentage of their benchmark true maximum generation capacities. From the left-hand side of the figure it is seen that GenCo 3 attains its highest daily net earnings at its lowest tested maximum generation capacity setting (0.95), regardless of the maximum generation 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 generation capacity is set at its lowest tested level (0.70) while at the same time the maximum generation capacity for GenCo 3 is set at its lowest tested level (0.95). Thus, GenCo 5's daily net earnings are affected by the maximum generation capacity setting for GenCo 3.

Figure 11 shows GenCo 3 and GenCo 5 mean daily net earnings on day 500 when both GenCos use VRE learning to exercise physical capacity withholding. Results for each GenCo are shown for various combinations of MPRMCap shinkage values for the two GenCos. From the left-hand side of the figure it is seen that GenCo 3 attains its highest daily net earnings

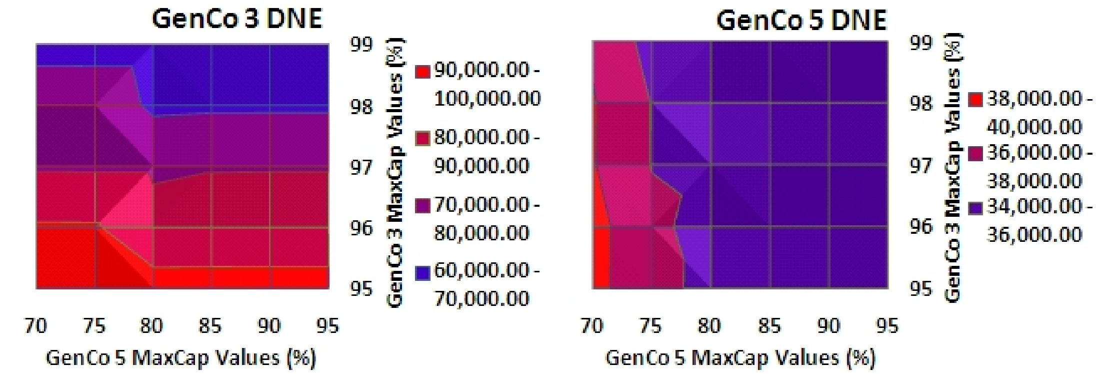

Figure 10: Daily net earnings (DNE) for GenCo 3 and GenCo 5 for the benchmark 5-bus test case modified to permit varied true maximum generation capacity values for GenCo 3 and GenCo 5. The latter values are depicted in percentage form (relative to original benchmark values).

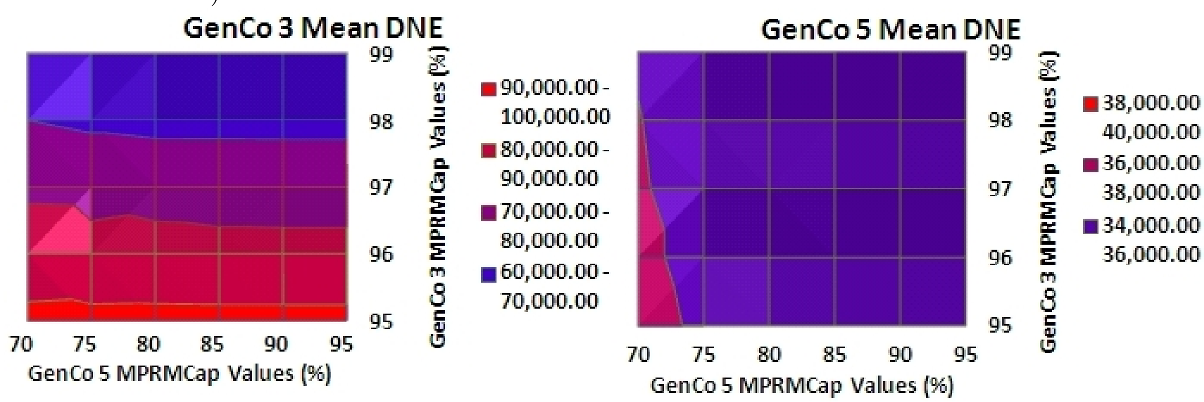

Figure 11: Mean daily net earnings (DNE) on day 500 for GenCo 3 and GenCo 5 when both GenCos use VRE learning to exercise physical capacity withholding. Mean DNE results for each GenCo are shown for various combinations of MPRMCap shrinkage values for the two GenCos.

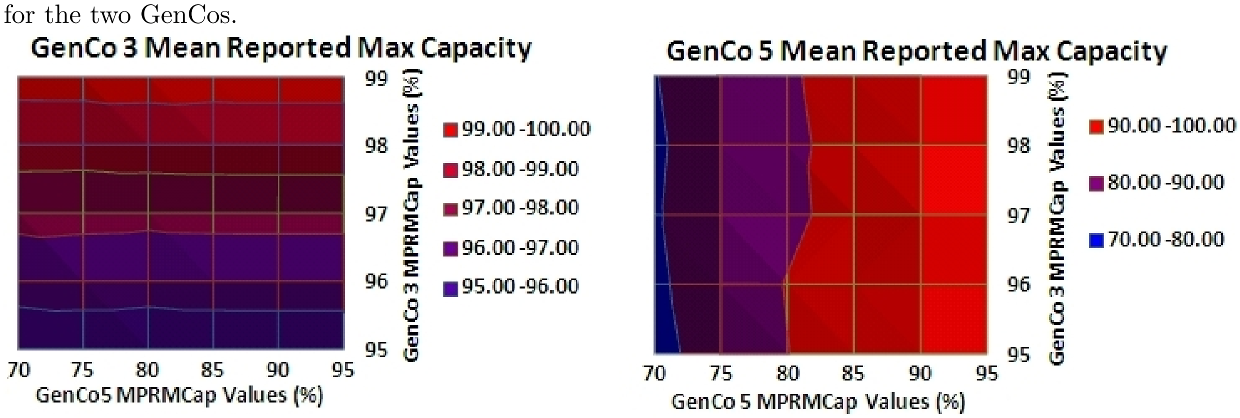

Figure 12: Mean reported maximum generation capacities (as a percentage of true maximum generation capacities) on day 500 for GenCo 3 and GenCo 5 when both GenCos use VRE learning to exercise physical capacity withholding. Mean reported maximum capacity results for each GenCo are shown for various combinations of MPRMCap shrinkage values for the two GenCos.

when its MPRMCap_3 value is set to its lowest tested level (0.95), regardless of the MPRMCap_5 shrinkage value 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_5 shrinkage value is set to its lowest tested level (0.70) while at the same time the MPRMCap_3 shrinkage value for GenCo 3 is set at its lowest tested level (0.95). Consequently, in similarity to the no-learning case, GenCo 5's daily net earnings under learning are affected by the MPRMCap_3 shrinkage value set for GenCo 3.

Figure 12 shows GenCo 3 and GenCo 5 mean reported maximum generation capacities as a percentage of their true maximum generation capacities when both GenCos use VRE learning to exercise physical capacity withholding. Results for each GenCo are shown for various combinations of MPRMCap shinkage values for the two GenCos. From the left-hand side of the figure it is seen that GenCo 3's mean reported maximum generation capacity is close to its MPRMCap_3 shrinkage value for each tested combination of MPRMCap shrinkage values for GenCo 3 and GenCo 5. Also, the horizonally-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 generation capacity is higher than its MPRMCap_5 shrinkage value for each tested combination of MPRMCap shrinkage values for GenCo 3 and GenCo 5. Also, GenCo 5's reported maximum generation capacities are weakly positively correlated with GenCo 3's reported maximum generation capacities (weak complementarity effect).

8. Test-Case Findings for Experiments with Combined Economic and Physical Capacity Withholding

In this section, two types of experiments are studied. The first type of experiment tests the extent to which a *single* 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 the ISO. The second type of experiment tests the extent to which *two* GenCos can co-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.

8.1. Combined Economic and Physical Capacity Withholding by One GenCo

For reasons elaborated in earlier sections, 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 for GenCo 3 in the sense that it yields consistently higher mean daily net earnings for GenCo 3.

Table 6 presents mean outcomes (with standard deviations) on day 1000 for GenCo net earnings, and for GenCo 3's reported supply offers, for a range of $\mathrm{MPRMCap}_3$ shrinkage value 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 $\mathrm{MPRMCap}_3$ setting for GenCo 3.

However, comparing the findings in Table 6 with the benchmark (no learning) and pure economic capacity withholding findings in Table 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 GenCo 3 increases its net earnings.

8.2. Combined Economic and Physical Capacity Withholding by Two GenCos

As in Section 6.2 and Section 7.2, learning experiments are conducted for two co-learning GenCos: namely, GenCo 3 and GenCo 5. Here, however, the two co-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 7 shows mean outcomes (with standard deviations) on day 1000 for GenCo net earnings, and for GenCo 3 and GenCo 5 reported supply offers, when GenCo 3 and GenCo 5 use VRE learning to exercise both economic and physical capacity withholding. Comparing these results to the results presented in Table 2 for the benchmark (no-learning) and pure economic capacity withholding cases, it is seen that GenCo 3 and GenCo 5 both attain much higher mean net earnings under learning. However, 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 (1). For example, GenCo 3 raises its reported ordinate value a^R dramatically, to almost four times its true value,

Table 6: Mean outcomes (with standard deviations) on day 1000 for GenCo net earnings, and for GenCo 3's reported supply offers (a^R,b^R,Cap^{RU}), when GenCo 3 uses VRE learning to exercise both economic and physical capacity withholding. Results are reported for a range of different $MPRMCap_3$ shrinkage values for GenCo 3.

MPRMCap for GenCo 3	99%	98%	97%	96%	95%
GenCo 1 $\overline{DNE}$	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
GenCo 2 $\overline{DNE}$	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
GenCo 3 $\overline{DNE}$	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 $\overline{DNE}$	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 $\overline{DNE}$	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 $\overline{a^R}$ (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 $\overline{b^R}$ (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 $\overline{Cap^{RU}}$ (TrueMaxCap=520)	517.62 (2.13)	515.23 (4.27)	512.85 (6.40)	510.47 (8.54)	508.08 (10.67)
GenCo 3 $\overline{Cap^{RU}}$/TrueMaxCap	99.54% (0.41%)	99.08% (0.82%)	98.63% (1.23%)	98.17% (1.64%)	97.71% (2.05%)

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 $MPRMCap_3$ shrinkage 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 7, it is also seen that GenCo 3 attains much higher mean net earnings than GenCo 5. The reasons for this are similar to the

Table 7: Mean outcomes (with standard deviations) on day 1000 for GenCo daily net earnings (DNE), and for GenCo 3 and GenCo 5's reported supply offers (a^R,$b^R Cap^{RU}$), when GenCo 3 and GenCo 5 use VRE learning to exercise both economic and physical capacity withholding. Results are reported for a fixed $MPRMCap_5$ shrinkage value of 0.70 for GenCo 5 and for a range of possible $MPRMCap_3$ shrinkage values for GenCo 3.

MPRMCap for GenCo 3	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 $\overline{Cap^{RU}}$ (TrueMaxCap=520)	517.66 (2.03)	515.58 (4.22)	512.72 (6.70)	509.25 (8.07)	505.92 (10.67)
GenCo 5 $\overline{Cap^{RU}}$ (TrueMaxCap=600)	507.00 (60.18)	508.50 (59.66)	508.50 (59.66)	511.50 (60.82)	508.50 (61.95)
GenCo 3 $\overline{Cap^{RU}}$/TrueMaxCap	99.55% (0.39%)	99.15% (0.81%)	98.60% (1.29%)	97.93% (1.55%)	97.29% (2.05%)
GenCo 5 $\overline{Cap^{RU}}$/TrueMaxCap	84.50% (10.03%)	84.75% (9.94%)	84.75% (9.94%)	85.25% (10.14%)	84.75% (10.33%)

reasons discussed in Section 6.2.

9. Comparisons of Experimental Findings

This section summarizes the detailed experimental findings for economic and physical capacity withholding reported in Sections 6 through 8.

9.1. Comparison of Cases with One GenCo Learner

From Table 2, Table 5 and Table 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 $\mathrm{MPRMCap}_3$ shrinkage value, these gains are substantially smaller.

9.2. Comparison of Cases with Two GenCo Learners

Co-learning findings for GenCo 3 and the relatively small and inexpensive GenCo 1 are reported in Table 3 and Figures 7-9. 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 of 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.

Co-learning experimental findings for GenCo 3 and the relatively large but inexpensive baseload GenCo 5 are presented in Tables 4-7 and Figures 10-12. 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 in combination with 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 co-learning experiments, 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 branch 1-2 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.

9.3. Overall Summary of Experimental Findings

The experiments reported in this study 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 for 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 all experiments reported in this study, the inexpensive but small GenCo 1 located at Bus 1 is persistently non-marginal. Consequently, 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 of GenCo 5 located at the neighboring Bus 5. If GenCo 5 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.

10. Concluding Remarks

The wholesale electric power market experiments conducted in this study provide relatively simple but empirically relevant test cases for examining claims about weak emergence, both positive and negative.

One issued noted by Bedau (1997) is whether weak emergence has any practical usefulness. The co-learning patterns stressed in this study — correlated GenCo supply offer behaviors and correlated GenCo net earnings outcomes – are weakly emergent in the sense of Bedau (1997). These co-learning patterns could provide useful information for real-world operators and regulators interested in understanding the myriad complicated ways in which market power can be directly and indirectly exercised in wholesale electric power markets.

For example, they reveal the existence of complementarity effects in which capacity withholding exercised by a relatively large and favorably located GenCo, such as GenCo 3, induces capacity withholding by other relatively large GenCos, such as GenCo 5. The resulting impacts on the net earnings of smaller GenCos, such as GenCo 1, can then be positive (a complementarity effect) or negative (a substitution effect) depending on how aggressively the larger GenCos exercise capacity withholding and the extent to which this withholding results in branch congestion affecting power flow on the grid.

Another issue noted by Bedau (1997) is predictability. By definition, the weakly emergent macrostates of a system S are fully determined given the microdynamics of S together with all impinging external conditions, including initial conditions. In what sense, then, are simulations really needed to understand these macrostates?

The 5-bus test cases computationally implemented in this study are deterministic systems. By construction, the co-learning patterns that emerge during the implementation for any one of these systems are entirely determined by the specification of the initial system microstate. This specification includes: the methods and attributes of the GenCos, the LSEs, and the ISO (decision-making agents); the methods and attributes of the day-ahead market (institutional agent); the methods and attributes of the transmission grid (physical agent); and initial seed values for all pseudo-random number generators included among these agents' methods.

The co-learning patterns arising in these test case implementations are "explained" ex post in this study by close examination of the underlying transmission grid configuration, market rules, GenCo true cost and capacity attributes, and GenCo learning capabilities that appear to advantage some GenCos over others. In addition, some use is made of *synchronic explanation*, i.e., current co-learning patterns are "explained" in terms of the current (synchronic) interactions among the rivalrous co-learning GenCos.

However, the actual causal linkages between the initial system microstates and the resulting co-learning patterns in these 5-bus test cases are so exceedingly complex it appears impossible that anyone could succeed in accurately predicting them in advance of experimentation. Moreover, the intensive nature of the experiments that would be needed to understand these linkages with confidence is daunting; the experiments conducted in the current study are suggestive but by no means definitive.

The saving grace here is that the 5-bus test cases, while highly simplified, capture important features of real-world wholesale electric power systems. As detailed in Weidlich and Veit (2008), agent-based researchers around the world are building empirically-grounded testbeds for the study of such systems, sharing the computational burden and speeding the real-world benefits. It seems likely that the concept of weakly emergent patterns will play an increasingly important role in conveying our hard-won understanding of these and other complex real-world systems critical for both social welfare and national security.

APPENDICES: TECHNICAL MATERIALS

A. GenCo Cost and Net Earnings Functions: Technical Details

For each day D, the *supply offer* s_i^R chosen by GenCo i to report to the ISO for use in each hour H of the day-ahead market for day D+1 consists of a linear *reported marginal cost function*

$$\mathrm{MC}_i^R(p_{Gi}) \;=\; \mathrm{a}_i^R \;+\; 2\mathrm{b}_i^R \cdot p_{Gi} \quad (\$/MWh) \tag{1}$$

defined over a *reported generation capacity interval*

$$0 \;\leq\; p_{Gi} \;\leq\; \mathrm{Cap}_i^{RU} \quad (MW) \tag{2}$$

for the generation of real power p_{Gi}. The expression $\mathrm{MC}_i^R(\mathrm{p}_{Gi})$ in (1) 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. The reported marginal cost functions (1) must lie either on or above GenCo i's *true marginal cost function*

$$\mathrm{MC}_i(p_{Gi}) \ = \ \mathrm{a}_i \ + \ 2\mathrm{b}_i \cdot p_{Gi} \quad (\$/MWh) \ \ , \tag{3}$$

where $\mathrm{a}_i > 0$ and $\mathrm{b}_i > 0$. Also, GenCo i's reported maximum generation capacity Cap_i^{RU} in (2) must lie within its true generation capacity interval:

$$0 \ \ \leq \ \ \mathrm{Cap}_i^{RU} \ \ \leq \ \ \mathrm{Cap}_i^{U} \ \ (MW) \ \ , \tag{4}$$

where $\mathrm{Cap}_i^U > 0$.

Thus, for the study at hand, a reported supply offer s_i^R for any GenCo i takes the form of a reported marginal cost function (1) that can be summarized by a vector (a_i^R,b_i^R) determining its ordinate a_i^R and slope $2\mathrm{b}_i^R$, together with a reported value Cap_i^{RU} for its maximum generation capacity. Henceforth such supply offers will be abbreviated in the form $\mathrm{s}_i^R = (\mathrm{a}_i^R, \mathrm{b}_i^R, \mathrm{Cap}_i^{RU})$.

At the beginning of any planning period, a GenCo's *avoidable costs* consist of the operational costs that it can avoid by shutting down production together with the portion of its fixed (non-operational) costs that it can avoid by taking appropriate additional actions such as asset re-use or re-sale. In order for production to proceed, revenues from production should at least cover avoidable costs. In the present study the GenCos do not incur start-up/shut-down or no-load costs, and all of their fixed costs are assumed to be *sunk*, i.e., non-avoidable. Consequently, the *avoidable cost function* $\mathrm{C}_i^a(\mathrm{p}_{Gi})$ for each GenCo i for any hour H is given by the integral of its true hourly marginal cost function:

$$\mathrm{C}_i^a(p_{Gi}) = \int_0^{p_{Gi}} \mathrm{MC}_i(p)dp = a_i p_{Gi} + b_i[p_{Gi}]^2 \quad (\$/h) \tag{5}$$

where p_{Gi} lies between 0 and Cap_i^U.

Suppose GenCo i, located at bus $k(i)$, is dispatched at level p_{Gi}(H,D) at price $\mathrm{LMP}_{k(i)}$(H,D) for hour H of the day-ahead market for day D+1. The *revenues* due to GenCo i for all 24 hours of day D+1, settled at the end of day D, are

$$\mathrm{Rev}_i(D) = \sum_{H=00}^{23} LMP_{k(i)}(H,D) \cdot p_{Gi}(H,D) \quad (\$) \tag{6}$$

Net earnings are defined as revenues minus avoidable costs. Let the avoidable costs incurred by GenCo i on day D for any hour H of day D+1 based on its day-D dispatch p_{Gi}(H,D) be denoted by C_i^a(H,D). Then the net earnings of GenCo i for all 24 hours of day D+1, realized on day D, are

$$\mathrm{NE}_i(D) \;=\; \mathrm{Rev}_i(D) - \sum_{H=00}^{23} \mathrm{C}_i^a(H,D) \quad (\$) \tag{7}$$

Finally, as will be seen in Section C, we make use of estimates MaxDNE_i for each GenCo i's maximum possible daily net earnings derived from its action domain AD_i assuming "competitive" marginal-cost pricing (sales price = reported marginal cost). Specifically,

$$\mathrm{MaxDNE}_i \;=\; 24h * \left(\max_{s_i^R \in AD_i} \left[HNE(s_i^R) \right] \right) \; (\$), \tag{8}$$

where the hourly net earnings function $\mathrm{HNE}(\mathrm{s}_i^R)$ (\$/h) is given by

$$HNE(s_i^R) \;=\; [MC_i^R(Cap_i^{RU}) * Cap_i^{RU}] - C_i^a(Cap_i^{RU}) \; . \tag{9}$$

Note that $\mathrm{MaxDNE}_i > 0$ if $\mathrm{s}_i^{\mathrm{true}} \equiv (\mathrm{a}_i,\mathrm{b}_i,\mathrm{Cap}_i^U)$ is included in GenCo i's action domain AD_i, a requirement imposed in the following subsection.

B. GenCo Learning: Technical Details

GenCo learning is implemented using a variant of a stochastic reinforcement learning method developed by Roth and Erev (1995) based on human-subject experiments, hereafter referred to as the *Variant Roth-Erev (VRE) learning method*. 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.

Each GenCo i has available an *action domain AD_i* consisting of a finite number of possible reported supply offers $\mathrm{s}_i^R = (\mathrm{a}_i^R,\mathrm{b}_i^R,\mathrm{Cap}_i^{RU})$, hereafter referred to as *actions*. The action domain AD_i is tailored to GenCo i's own particular true cost and capacity attributes as follows: (a) it only contains reported marginal cost functions (1) lying on or above GenCo i's true marginal cost function (3); and (b) it contains GenCo i's true marginal cost function defined over GenCo i's true generation capacity interval, i.e., it contains $\mathrm{s}_i^{\mathrm{true}}$

$\equiv$ (a_i,b_i,Cap_i^U). However, the action domains are constructed so as to ensure equal cardinalities and similar densities across all GenCos to avoid favoring some GenCos over others purely through action domain construction.[13]

The remainder of this section describes how an arbitrary GenCo i goes about using the VRE learning method to select actions s_i^R from its action domain AD_i to submit to the ISO for the day-ahead energy market on successive days D, starting from an initial day D=1. As will be seen below, the only relevant attribute of AD_i for implementation of VRE learning is that it has finite cardinality. Consequently, letting $M_i \geq 1$ denote the cardinality of AD_i, it suffices to index the actions in AD_i by $m = 1,...,M_i$.

The *initial propensity* of GenCo i to choose action $m \in AD_i$ is given by $q_{im}(1)$ for m = $1,...,M_i$. AMES(V2.05) permits the user to set these initial propensity levels to any real numbers. However, the assumption used in this study is that GenCo i's initial propensity levels are all set equal to some common value $q_i(1)$, as follows:

$$q_{im}(1) \quad = \quad q_i(1) \;\text{ for all actions } m \in AD_i \tag{10}$$

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

$$p_{im}(D) = \frac{\exp(q_{im}(D)/T_i)}{\sum_{j=1}^{M_i} \exp(q_{ij}(D)/T_i)} \;,\; m \in AD_i \tag{11}$$

In (11), 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 (11) become increasingly peaked over the particular actions m having the highest propensity values q_{im}(D), thereby increasing the probability that these actions will be chosen.

At the end of day D, the current propensity q_{im}(D) that GenCo i associates with each action m in AD_i is updated in accordance with the following

[13] A complete technical explanation of this important action domain construction can be found in Appendix B of Li et al. (2009).

rule. Let m' denote the action *actually* selected and reported into the day-ahead market by GenCo i in day D. Also, let $\mathrm{NE}_{im'}$(D) denote the *actual* daily net earnings (7) attained by GenCo i at the end of day D as its settlement payment for all 24 hours of the day-ahead market for day D+1. Then, for each action m in AD_i,

$$q_{im}(D+1) \quad = \quad [1-r_i]q_{im}(D) \; + \; \mathrm{Response}_{im}(D) \quad , \tag{12}$$

$$\mathrm{Response}_{im}(D) = \begin{cases} [1-e_i] \cdot NE_{im'}(D) & \text{if} \quad m = m' \\ \\ e_i \cdot q_{im}(D)/[M_i - 1] & \text{if} \quad m \neq m', \end{cases} \tag{13}$$

where[14] $r_i \in [0,1]$, $e_i \in [0,1)$, and $m \neq m'$ implies $\mathrm{M}_i \geq 2$. The introduction of the *recency parameter* r_i in (12) acts as a damper on the growth of the propensities over time. The *experimentation parameter* e_i in (13) permits reinforcement to spill over to some extent from a chosen action to other actions to encourage continued experimentation with various actions in the early stages of the learning process.

C. Calibration of GenCo Learning: Technical Details

This appendix section provides technical details in support of the discussion in Section 5 regarding the sweet-spot calibration of VRE learning parameters for each learning GenCo.

C.1. Learning Calibration for Pure Economic Capacity Withholding Experiments

In the pure economic capacity withholding experiments reported in Section 6, each learning GenCo i makes daily use of the VRE learning method to adjust the ordinate and slope parameters a_i^R,b_i^R of its reported marginal

[14]In the original Roth-Erev method, the term $q_{im}(D)$ in (13) is instead given by $\mathrm{NE}_{im'}$(D). However, as explained in Nicolaisen et al. (2001), in this case there is no updating of propensities when net earnings outcomes are zero, e.g., due to a failure to be dispatched. This can result in prolonged mushing around in the early stages of learning when GenCos are trying to learn appropriate supply offers, with subsequent losses of net earnings and reductions in market efficiency. The substitution of $q_{im}(D)$ for $\mathrm{NE}_{im'}$(D) in equation (13), introduced in Nicolaisen et al. (2001) to avoid this zero-updating problem, resulted in dramatic improvements in both GenCo net earnings and in market efficiency.

cost function (1) in pursuit of increased net earnings. The action domain AD_i for each learning GenCo i includes 100 possible reported supply offers of the form $s_i^R = (a_i^R, b_i^R, Cap_i^U)$, constructed by crossing 10 distinct values for a_i^R with ten distinct values for b_i^R.

The VRE recency and experimentation learning parameters r_i and e_i for each learning GenCo i are fixed at common levels $r = 0.04$ and $e = 0.96$. These values for r_i and e_i have yielded relatively high mean net earnings for the GenCos in intensive sensitivity experiments conducted by ISU researchers for a variety of market contexts with VRE-learning traders, including 3-bus and 5-bus wholesale power market experiments with pure economic capacity holding conditional on various settings for the initial propensity levels $q_i(1)$ and temperature levels T_i for each GenCo i [cf. Pentapalli (2008)].

Given these sweet-spot values for r_i and e_i, we then conducted intensive parameter sweeps to determine sweet-spot settings for $q_i(1)$ and T_i for each GenCo i for the 5-bus test case with each GenCo able to learn to exercise pure economic capacity withholding. More precisely, we systematically tested a range of positive values for α_i and β_i, defined as follows:

- $\alpha_i = q_i(1)/\text{MaxDNE}_i$, where $q_i(1)$ is the initial propensity level for GenCo i appearing in (10), a measure of GenCo i's net earnings aspirations at the beginning of the initial day 1, and MaxDNE_i denotes GenCo i's (positive) estimated maximum possible daily net earnings as determined in (8);

- $\beta_i = q_i(1)/T_i$, where T_i is the temperature parameter for GenCo i appearing in (11).

As detailed in Appendix A, the (positive) estimate MaxDNE_i is exogenously determined from the structural aspects of GenCo i's action domain AD_i. Consequently, any specification of positive α_i and β_i values for GenCo i determines unique positive $q_i(1)$ and T_i values for GenCo i.[15]

Figure 4 depicts experimental findings for mean total GenCo daily net earnings attained for the 5-bus case with pure economic capacity holding

[15]In our earlier study Li et al. (2008), identical initial propensity and temperature levels were set for all learning GenCos: namely, $q(1) = 6000$ and $T = 1000$. This was unsatisfactory since the "prior anticipated net earnings" $q(1)$ were then set commonly across GenCos with different costs and locations without regard for their different net earnings opportunities.

under various positive commonly-set values for

$$\alpha \ = \ \frac{q_i(1)}{\mathrm{MaxDNE}_i} \ , \quad \beta \ = \ \frac{q_i(1)}{T_i} \ , \ i = 1, \ldots, 5 \ . \tag{14}$$

An interesting pattern is immediately evident in Fig. 4. The (α,β) combinations associated with the highest mean net earnings outcomes lie along a nonlinear ridge line that traverses from (1, 100) in the northwest corner to (1/24, 2) in the south-central region.

In all of the pure economic capacity withholding experiments reported in Section 6, the values $(\alpha, \ \beta) = (1, 100)$ found to achieve the relatively highest mean total GenCo daily net earnings for pure economic capacity withholding are used as sweet-spot settings for the VRE α and β parameters in (14) for each learning GenCo i. This, in turn, determines sweet-spot settings for the VRE initial propensity and temperature parameters $\mathrm{q}_i(1)$ and T_i for each learning GenCo i.

C.2. Learning Calibration for Pure Physical Capacity Withholding Experiments

In the pure physical capacity withholding experiments reported in Section 7, the following value ranges are used for the GenCos' *Minimum Possible Reported Max Capacity (MPRMCap)* fractions (i.e., their shrinkage values):

- 0.75 to 0.95 for GenCo 1;
- 1.00 for GenCo 2;
- 0.95 to 0.99 for GenCo 3;
- 1.00 for GenCo 4;
- 0.70 to 0.95 for GenCo 5.

If $\mathrm{MPRMCap}_i < 1.00$ for GenCo i, its action domain AD_i includes 30 possible reported supply offers of the form $\mathrm{s}_i^R = (\mathrm{a}_i,\mathrm{b}_i,\mathrm{Cap}_i^{RU})$ for thirty distinct possible reported maximimum generation capacity settings Cap_i^{RU}. The 30 settings for Cap_i^{RU} are equally spaced across the interval ranging from the lower bound $\mathrm{MPRMCap}_i \cdot \mathrm{Cap}_i^U$ to the upper bound Cap_i^U, where Cap_i^U denotes GenCo i's true maximum generation capacity. Alternatively, if $\mathrm{MPRMCap}_i = 1.00$ for GenCo i, its action domain AD_i includes only one possible reported supply offer, $s_i^{\text{true}} = (a_i, b_i, Cap_i^U)$. That is, AD_i includes only GenCo i's true marginal cost function (3) defined over its true generation capacity interval (4), implying GenCo i does not have learning capabilities.

It follows from this action domain construction that our pure physical capacity withholding experiments involve at most three learners: GenCo 1, GenCo 3, and GenCo 5. For each of these potential learners the VRE recency and experimentation learning parameters r_i and e_i are fixed at the common values r = 0.04 and e = 0.96, as in Appendix C.1. However, as indicated in Table 8, preliminary pure physical capacity withholding experiments were conducted to determine sweet-spot values for the VRE α and β parameters in (14) for these three potential learners.[16]

Table 8: GenCo action domain and learning parameter settings for pure physical capacity withholding experiments with three potential learners: GenCos 1, 3, and 5.

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)

Figure 5 depicts pure physical capacity withholding experimental findings for mean GenCo 3 net earnings on day 1000 under alternative (α, β) specifications set commonly for GenCos 1, 3, and 5 with MPRMCap=0.95 for all three GenCos. The net earnings of GenCo 3 are chosen as the criterion for the calibration of (α, β) for our pure physical capacity withholding experiments because in most hours this large supplier turns out to be a *pivotal*

[16]For a full explanation of the action domain parameters appearing in Table 8 that determine the construction of the GenCos' action domains, see Appendix B of Li et al. (2009).

supplier, i.e., its generation capacity is essential for meeting fixed demand.

In all of the pure physical capacity withholding experiments reported in Section 7, the values $(\alpha,\ \beta) = (1/24,\ 100)$ found above to achieve the relatively highest mean daily net earnings for GenCo 3 are used as sweet-spot settings for the VRE α and β parameters in (14) for each learning GenCo i. This, in turn, determines sweet-spot settings for the VRE initial propensity and temperature parameters $\mathrm{q}_i(1)$ and T_i for each learning GenCo i.

C.3. Learning Calibration for Combined Capacity Withholding Experiments

In the combined economic and physical capacity withholding experiments reported in Section 8, only two types of learning are considered. Either GenCo 3 is the only learning GenCo (capable of both economic and physical capacity withholding), or GenCo 3 and GenCo 5 are the only learning GenCos (each capable of both economic and physical capacity withholding). Consequently, the action domains AD_i are constructed as follows.

The action domain AD_3 for GenCo 3, always a learner, includes $3000 = 10{\times}10{\times}30$ possible reported supply offers of the form $\mathrm{s}_3^R = (\mathrm{a}_3^R,\mathrm{b}_3^R,\mathrm{Cap}_3^{RU})$. This action domain is constructed as the cross-product of ten values for the reported ordinate parameter a_3^R, ten values for the reported slope parameter b_3^R, and 30 values for the reported maximum generation capacity Cap_3^{RU}. The reported ordinate/slope values for GenCo 3 are the same as used for GenCo 3 in the pure economic withholding experiments (see Appendix C.1) and the reported maximum generation capacity values for GenCo 3 are the same as used for GenCo 3 in the pure physical capacity withholding experiments (see Appendix C.2).

The action domain AD_5 for GenCo 5 in experiments for which GenCo 5 is a learner (hence $\mathrm{MPRMCap}_5 < 1.00$) includes 3000 possible reported supply offers, constructed similarly to AD_3. The action domain AD_i for each non-learning GenCo i consists of only one possible reported supply offer, $\mathrm{s}_i^{\mathrm{true}} = (\mathrm{a}_i,\ \mathrm{b}_i,\ \mathrm{Cap}_i^U)$. That is, AD_i consists of GenCo i's true marginal cost function (3) defined over its true generating capacity interval (4).

For each combined economic and physical capacity withholding experiment, the VRE learning parameters (r, e, α β) are set commonly across all learning GenCos at the particular values (0.04, 0.96, 1, 100) determined in Appendix C.1. Ideally, a separate calibration exercise for these four VRE learning parameters should be undertaken for the combined case to see how the sweet-spot region for the combined case geometrically relates to the sweet-spot regions for each of the pure cases; this is a topic for future research.

Acknowledgement: The authors are grateful for exceptionally constructive and helpful comments received from the editors and two anonymous referees. This work was supported in part by a grant from the Electric Power Research Center.

References

Anderson, P.W., 1972. More is different. Science 177(4047), 393–396.

Auyang, S., 1998. Foundations of Complex-System Theories in Economics, Evolutionary Biology, and Statistical Physics. Cambridge, UK: Cambridge University Press.

Baldick, R., Hogan, W., 2002. Capacity constrained supply function equilibrium models of electricity markets: Stability, non-decreasing constraints, and function space iterations. University of California Energy Institute PWP-089. Online: Accessible at www.ucei.berkeley.edu/ucei/PDF/pwp089.pdf

Bedau, M., 1997. Weak emergence. In: Tomberlin, J. (Ed.), Philosophical Perspectives: Mind, Causation, and World. Vol. 11, Malden, MA: Blackwell, 375–399.

Branner, B., 1989. The Mandelbrot set. In: Devaney, R.L., Keen, L. (Eds.). Chaos and Fractals: The Mathematics Behind the Computer Graphics. Proceedings of Symposia in Applied Mathematics, Vol. 39, Providence, RI: American Mathematical Society, 75-105.

Borrill, P., Tesfatsion, L., 2011. Agent-based modeling: The right mathematics for the social sciences? In: Davis, J.B, Hands, D. W., (Eds.). Elgar Recent Economic Methodology Companion, Northhampton, MA: Edward Elgar Publishing.

Bulow, J., Geanakoplos, J.D., and Klemperer, P., 1985. Multimarket oligopoly: Strategic substitutes and complements. Journal of Political Economy 93, 488–511.

Corning, P.A., 2002. The re-emergence of "emergence": A venerable concept in search of a theory. Complexity 7(6), 18–30.

Dechenaux, E., Kovenock, D., 2007. Tacit Collusion and capacity withholding in repeated uniform price auctions. RAND Journal of Economics 38(4), 1044-1069.

Delgado, J., 2002. Emergence of social conventions in complex networks. Artificial Intelligence 141, 171–185.

Dessalles J.L., Ferber J., Phan D., 2008. Emergence in agent-based computational social science: Conceptual, formal and diagrammatic analysis. In: Yang, A., Shan, Y. (Eds.), Intelligent Complex Adaptive Systems, Idea Group Incorporated (IGI) Global, Chapter 9, 255–299. Online: Accessible at www.gemas.msh-paris.fr/dphan/papers/DessallesFerberPhan2007.pdf

Devaney, R.L., Keen, L. (Eds.), 1989. Chaos and Fractals: The Mathematics Behind the Computer Graphics. Proceedings of Symposia in Applied Mathematics, Vol. 39, Providence, RI: American Mathematical Society.

Gilbert, N., 1995. Emergence in social simulations. In: Gilbert, N., Conte, R. (Eds.). Artificial Societies: The Computer Simulation of Social Life. London: UCL Press, 144-156.

Gode, D.K., Sunder, S., 1993. Allocative efficiency of markets with zero intelligence traders: Market as a partical substitute for individual rationality. Journal of Political Economy 101, 119–137.

Gode, D.K., Sunder, S., 1997. What makes markets allocationally efficient?. Quarterly Journal of Economics 112(2), 603–630.

Harper, D.A., Endres, A.M., 2010. The anatomy of emergence, with a focus on capital formation. Working Paper, Department of Economics, New York University, April.

Hayek, F.A., 1948. Individualism and Economic Order. Chicago, IL: University of Chicago Press.

Hodgson, G.M., 2000. The concept of emergence in social science: Its history and importance. Emergence 2(4), 65–77.

IEEE, 2011. Task Force on Open Source Software for Power Systems, Institute of Electrical and Electronic Engineers. Online: Accessible at ewh.ieee.org/cmte/psace/CAMS_taskforce/index.htm

Kirman, A., 1993. Ants, rationality, and recruitment. Quarterly Journal of Economics 108(1), 137-156.

Kittock, J.E., 1995. Emergent conventions and the structure of multi-agent systems. In: Nadel, L., Stein, D. (Eds.), 1993 Lectures in Complex Systems, SFI Studies in the Sciences of Complexity, Vol. VI, Santa Fe Institute, Addison-Wesley, Reading, MA, 507–521.

Kuperberg, M., 2006. The two faces of emergence in economics. Presented at Emergent Systems: A Discussion, Center for Science in Society, Bryn Mawr College, July 18. Online: Accessible at www.econ.iastate.edu/tesfatsi/TwoFacesOfEmergence.MKuperberg.pdf

Lally, J., 2002. Financial transmission rights: Auction example, Section 6 in Financial Transmission Rights Draft 01-10-02, m-06 ed., ISO New England, Inc.

Laughlin, R.B., Pines, D., Schmalian, J., Stojković, B.P., Wolynes, P., 2000. The middle way. Proceedings of the National Academy of Sciences 97, 32–37.

Laughlin, R.B., 2002. The physical basis of computability. Computing in Science and Engineering 4(3), 27–30.

Lewis, P., 2011. Emergent properties in the work of Friedrich Hayek. Journal of Economic Behavior and Organization, to appear.

Li, H., Tesfatsion, L., 2009a. Capacity withholding in restructured wholesale power markets: An agent-based testbed study, Proceedings of the IEEE Power Systems Conference & Exposition (PSCE), Seattle, WA. Online: www.econ.iastate.edu/tesfatsi/CapacityWithhold.LiTesPSCE2009.pdf

Li, H., Tesfatsion, L., 2009b. The AMES wholesale power market testbed: A computational laboratory for research, teaching, and training. Proceedings of the IEEE Power and Energy Society General Meeting, Calgary, Alberta, CA. Online: www.econ.iastate.edu/tesfatsi/AMESTestBed.2009IEEEPESGM.pdf

Li, H., Tesfatsion, L., 2009c. Development of open source software for power market research: The AMES testbed", Journal of Energy Markets 2(2), 111–128.

Li, H., Tesfatsion, L., 2011. ISO Net Surplus Collection and Allocation in Wholesale Power Markets under LMP. IEEE Transactions on Power Systems 26(2), 627–641.

Li, H., Sun, J., Tesfation, L., 2008. Dynamic LMP response under alternative price-cap and price-sensitive demand scenarios. Proceedings of the IEEE Power and Energy Society General Meeting, Pittsburgh, PA.

Li, H., Sun, J., Tesfatsion, L., 2009. Separation and volatility of locational marginal prices in restructured wholesale power markets. Economics Working Paper #09009, Iowa State University, Ames, Iowa. Online: Accessible at www.econ.iastate.edu/research/publications/viewabstract.asp?pid=13075

Nicolaisen, J., Petrov, V., 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.

O'Connor, R., Wong, H.Y., 2006. Emergent properties. In: Zalta, E.N. (Ed.). The Stanford Encyclopedia of Philosophy. Online: Accessible at http://plato.stanford.edu/entries/properties-emergent/

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, Ames, Iowa.

Roth, A.E., Erev, I., 1995. Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior 8, 164–212.

Schelling, T.S., 1978. Micromotives and Macrobehavior. New York, NY: W.W. Norton & Co.

Shahidehpour, M., Yamin, H., Li, Z., 2002. Market Operations in Electric Power Systems. New York:IEEE/Wiley-Interscience.

Shoham, Y., Tennenholtz, M., 1994. Co-learning and the evolution of social activity. Technical Report. Stanford University.

Smith, A., 1937. An Inquiry into the Nature and Causes of the Wealth of Nations. New York, NY: Cannan Edition, American Modern Library Series.

Squazzoni, F. (Ed.), 2006. Epistemological Aspects of Computer Simulation in the Social Sciences. Lecture Notes in Artificial Intelligence, Vol. 5466, Springer-Verlag Berlin, Heidelberg.

Stephan, A., 1999. Varieties of emergentism. Evolution and Cognition 5(1), 49–59.

Tellidou, A., 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., 2010. AMES Wholesale Power Market Testbed Homepage, Iowa State University. Online: Accessible at www.econ.iastate.edu/tesfatsi/AMESMarketHome.htm

Tesfatsion, L., Judd, K.L., 2006. Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics. Handbooks in Economics Series. Amsterdam:North-Holland, 831–880.

Tirole, J., 2003. The Theory of Industrial Organization. Fourteenth Printing. Cambridge, MA: MIT Press.

Urbano, P., Balso, J., Antunes, L., Moniz, L., 2009. Force versus majority: A comparison in convention emergence efficiency. In: Hübner, J., Matson, E., Boissier, O., and Dignum, V. (Eds.), Coordination, Organizations, Institutions, and Norms in Agent Systems IV. Lecture Notes in Computer Science, Vol. 5428, Springer Berlin/Heidelberg, 48–63.

Weidlich, A., Veit, D., 2008. A critical survey of agent-based wholesale electricity market models. Energy Economics 30(4), 1728–1759.