

2010

Electric Power Market Modeling with Multi-Agent Reinforcement Learning

Nathanael K. Miksis

University of Massachusetts Amherst

Follow this and additional works at: <https://scholarworks.umass.edu/theses>



Part of the [Industrial Engineering Commons](#), and the [Operational Research Commons](#)

Miksis, Nathanael K., "Electric Power Market Modeling with Multi-Agent Reinforcement Learning" (2010). *Masters Theses 1911 - February 2014*. 494.

Retrieved from <https://scholarworks.umass.edu/theses/494>

This thesis is brought to you for free and open access by ScholarWorks@UMass Amherst. It has been accepted for inclusion in Masters Theses 1911 - February 2014 by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

ELECTRIC POWER MARKET MODELING WITH
MULTI-AGENT REINFORCEMENT LEARNING

A Thesis Presented

By

NATHANAEL K. MIKSIS

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of
MASTER OF SCIENCE IN INDUSTRIAL ENGINEERING AND OPERATIONS RESEARCH
September 2010
Industrial Engineering & Operations Research

© Copyright by Nathanael K. Miksis 2010

All Rights Reserved

ELECTRIC POWER MARKET MODELING WITH
MULTI-AGENT REINFORCEMENT LEARNING

A Thesis Presented

by

NATHANAEL K. MIKSIS

Approved as to style and content by:

Erin Baker, Chair

Ana Muriel, Member

Victor Lesser, Member

Donald Fisher, Department Chair

Department of Mechanical and Industrial Engineering

DEDICATION

To my supportive parents who always encouraged me in all my endeavors.

ACKNOWLEDGEMENTS

I would like to thank many people who have helped me and encouraged me to pursue research that became this thesis. I will be ever grateful to my advisor, Erin Baker, for her constructive feedback and support to see this venture through, and to my committee members Ana Muriel and Victor Lesser. I would also like to thank Mario DePillis, my good friend and my mentor at ISO New England Inc. and to all the other incredibly capable people I met and learned from there. Robert Entriken of the Electric Power Research Institute (EPRI) went above and beyond, providing the supportive feedback and encouragement that at just the right times. Ekundayo Shittu was the best office-mate anyone could ask for when I was at the University.

I am grateful for the funding support that came from the University of Massachusetts Amherst and ISO New England Inc. Funding for research done at Électricité de France, which provided a foundation for some of the work presented here, also was provided by EPRI.

And finally, last but certainly not least, to my family and my friends for their unwavering encouragement.

ABSTRACT

ELECTRIC POWER MARKET MODELING WITH MULTI-AGENT REINFORCEMENT LEARNING

SEPTEMBER 2010

NATHANAEL K. MIKSIS, M.S., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Erin Baker

Agent-based modeling (ABM) is a relatively new tool for use in electric power market research. At heart are software agents representing real-world stakeholders in the industry: utilities, power producers, system operators, and regulators. Agents interact in an environment modeled after the real-world market and underlying physical infrastructure of modern power systems. Robust simulation laboratories will allow interested parties to stress test regulatory changes with agents motivated and able to exploit any weaknesses, before making these changes in the real world. Eventually ABM may help develop better understandings of electric market economic dynamics, clarifying both delineations and practical implications of market power.

The research presented here builds upon work done in collateral fields of machine learning and computational economics, as well as academic and industry literature on electric power systems. We build a simplified transmission model with agents having learning capabilities, in order to explore agent performance under several plausible scenarios. The model omits significant features of modern electric power markets, but is able to demonstrate successful convergence to stable profit-maximizing equilibria of adaptive agents competing in a quantity-based, available capacity model.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	v
ABSTRACT	vi
LIST OF TABLES	ix
LIST OF FIGURES	xi
CHAPTER	
1. INTRODUCTION	1
1.1. Primer on Electric Power Markets	3
1.2. Pressing Questions in Electric Power Market Regulation and Design	8
1.3. How Agent-Based Modeling Has Been Used in Electric Power Market Applications	15
2. AGENT-BASED MODELING: THE CONTEXT OF THIS RESEARCH	18
2.1. Challenges and Opportunities	23
3. THE ELECTRICITY MARKET MODEL	30
3.1. Where This Research Fits: What’s New Under the Sun?	31
3.2. Predecessor Model	32
3.3. Current Market Model Description	36
3.3.1. The Power System Model	39
3.3.2. Least-Cost Dispatch	41
3.3.3. Pricing and Congestion	42
3.3.4. Agents	44
3.4. Implementing the Model in Java	47

4. EXPLORING THE MODEL.....	49
5. USING THE MODEL	51
5.1. Scenario 1.....	53
5.2. Scenario 2.....	59
5.3. Scenario 3.....	61
5.4. Using The Model – Conclusion	63
6. CONCLUSION AND WHAT’S NEXT	64
APPENDIX A: EXPOSITION OF NODAL PRICES	68
APPENDIX B: DATA.....	71
BIBLIOGRAPHY.....	81

LIST OF TABLES

Table	Page
1. Generator Shift Factors (Nodes in Rows, Lines in Columns)	68
2. Transmission System Data.....	71
3. Defaults Scenarios' Power System Parameters	71
4. Scenario 1 Cleared All-Bid-In, Portfolio Results	72
5. Scenario 1 Cleared All-Bid-In, Nodal Prices.....	72
6. Scenario 1 Cleared All-Bid-In, Transmission Line Flows.....	72
7. Scenario 1 Cleared All-Bid-In, Nodal MWs.....	73
8. Scenario 1, Cleared Equilibrium Joint Strategy, Agent Quantity Offers.....	73
9. Scenario 1, Cleared Equilibrium Joint Strategy, Portfolio Results.....	73
10. Scenario 1 Cleared Equilibrium Joint Strategy, Nodal Prices	74
11. Scenario 1, Cleared Equilibrium Joint Strategy, Transmission Line Flows.....	74
12. Scenario 1, Cleared Equilibrium Joint Strategy, Nodal MWs	74
13. Scenario 2, Cleared All-Bid-In, Portfolio Results	75
14. Scenario 2, Cleared All-Bid-In, Nodal Prices.....	75
15. Scenario 2, Cleared All-Bid-In, Transmission Line Flows.....	75
16. Scenario 2, Cleared All-Bid-In, Nodal MWs.....	76
17. Scenario 2, Cleared Equilibrium Joint Strategy, Agent Quantity Offers.....	76
18. Scenario 2, Cleared Equilibrium Joint Strategy, Portfolio Results.....	76
19. Scenario 2, Cleared Equilibrium Joint Strategy, Nodal Prices	77
20. Scenario 2, Cleared Equilibrium Joint Strategy, Transmission Line Flows.....	77

21. Scenario 2, Cleared Equilibrium Joint Strategy, Nodal MWs	77
22. Scenario 3, Cleared All-Bid-In, Portfolio Results	78
23. Scenario 3, Cleared All-Bid-In, Nodal Prices.....	78
24. Scenario 3, Cleared All-Bid-In, Transmission Line Flows.....	78
25. Scenario 3, Cleared All-Bid-In, Nodal MWs.....	79
26. Scenario 3, Cleared Equilibrium Joint Strategy, Agent Quantity Offers.....	79
27. Scenario 3, Cleared Equilibrium Joint Strategy, Portfolio Results.....	79
28. Scenario 3, Cleared Equilibrium Joint Strategy, Nodal Prices	80
29. Scenario 3, Cleared Equilibrium Joint Strategy, Transmission Line Flows.....	80
30. Scenario 3, Cleared Equilibrium Joint Strategy, Nodal MWs	80

LIST OF FIGURES

Figure	Page
1. Q-Agent and \$10 Static Offer Opponent, 1,000 MW Load.....	34
2. Q-Agent and \$10 Static Offer Opponent, 1,400 MW Load.....	35
3. Q-Agent and \$10 Static Offer Opponent, 1,600 MW Load.....	36
4. Modeled Power System Diagram	39
5. Line Flow Demonstration with Node 3 Injection	40
6. Market Model Interface	50
7. System Diagram Showing Line Limits and Loads	53
8. Scenario 1, Nodal LMPs for a Single Simulation Run	55
9. Scenario 1, Load-Weighted System LMP for Three Simulation Runs.....	57
10. Scenario 1, Aggregate Supply Offered.....	58
11. Scenario 1, Individual Agents' Profits over the Simulation Run.....	59
12. Scenario 2, Nodal LMPs for a Single Simulation Run	60
13. Scenario 2, Aggregate Supply Offered.....	61
14. Scenario 3, Nodal LMPs for a Single Simulation Run	62
15. Scenario 3, Aggregate Supply Offered.....	63
16. Scenario 1 Equilibrium Joint Strategy, Power Line Flows % of Capacity.....	68

CHAPTER 1

INTRODUCTION

This thesis reports on a research project to build an adaptive agent-based model (ABM) and apply it to the restructured electric power industry. We use as a basis a learning algorithm developed within the field of reinforcement learning (a cross-disciplinary field originating out of computer science and machine learning) to model electric power generating companies in a simplified multi-node transmission-constrained power system, similar in kind if not in scale to real-world power systems (such as that operated by ISO New England Inc.). There are myriad challenges ahead for researchers in this field, and these are detailed in this proposal. However, equally as important is the potential value that robust and demonstrably-accurate simulation tools hold for applications to electric power market modeling.

Early literature in the reinforcement learning field has already shown that optimality proofs exist for agents engaged in online learning in a stationary environment (Sutton and Barto, 1998). The challenges of using online learning algorithms in non-stationary environments, of which most electric power market models are examples, are described below¹. In this paper, we first provide an overview of the significant features of modern electricity markets that make modeling research particularly valuable (Section 1.1). We then cover relevant contemporary research being done on more realistic market models (including multi-node transmission systems and more adaptive agents competing simultaneously in the same market) and some background

¹ A stationary task is not necessarily deterministic, but one in which the effect of an action is a sample of a fixed probability distribution. Through trial and error, a person or an agent can learn that the expected result of a coin flip is 50% heads and 50% tails and will be correct in believing this distribution to be true, but many tasks involve probability distributions that are not stationary. A task in which the effect of an action depends on the action taken by another agent who is also learning through trial and error is one example.

on agent-based modeling (Section 1.3 and Chapter 2). After that, we outline first the agent-based electric power market model presented in Miksis, 2006, which was a predecessor to the current model. Then we detail the new model in Section 3.3, which incorporates a 5-node transmission model, with thermal line limits and congestion pricing (or locational marginal pricing)². In contrast to most contemporary research, and the model presented in Miksis, 2006, the agents in the new model compete on quantity rather than price.³

Because quantity-based competition has received less attention than price-based models in the electric power ABM field, while market manipulation through physical withholding remains a significant potential in existing competitive markets, we hope to demonstrate the utility of this line of research in the field by answering several questions:

1. Can a community of adaptive agents competing in a quantity-based market model achieve equilibrium under various initial conditions?
2. Can players with plants at multiple locations (on either side of a transmission constraint) discover withholding policies that cause congestion, essentially raising prices at the import-constrained node?
3. Does the market as a whole supply surplus capacity above load or is the average excess capacity margin in equilibrium close to zero?

² See http://en.wikipedia.org/wiki/Electricity_market#Bid-based.2C_security-constrained.2C_economic_dispatch_with_nodal_prices

³ Banal-Estanol and Micola, 2009, survey agent-based electricity market models and highlight the shortcomings of price-based competition (Bertrand models). They mention that neither Bertrand nor Cournot (pure quantity-based) models are ideally suited to electricity markets because real-world markets allow variable prices over tranches of capacity. We have attempted to address this by modifying traditional quantity-based competition to introduce these variably-priced tranches of capacity.

Several researchers have incorporated additional features of real world power markets that we were unable to incorporate, but are possible as an extension to this work, such as demand-side bidding and a multi-settlement (Day-Ahead and Balancing) market. Many pressing questions about the effects of multi-settlement and elastic demand on market efficiency have real world implications and modeling these features with adaptive agents could contribute to policy discussions, particularly when assessment of costs-benefits of implementation is being considered. The next section introduces in broad form modern electric power markets, some history and some distinguishing features.

1.1. Primer on Electric Power Markets

Here we outline the salient features of modern restructured power systems that motivate the research and provide an overview of recent modeling efforts. Several recent published works provide excellent overviews on the most critical avenues of industry research as well as the state-of-the-art modeling and experiment efforts, including Conzelman, et al, 2004, Ventosa, et al, 2005, Weidlich and Veit, 2008, Nanduri and Das, 2009.

In the last couple of decades, restructuring in the electric power industry has fundamentally changed the environment in which every interested party operates, from regulators to system operators to generation and transmission owners to load-serving power companies and consumers. One thing has become clear in the move toward a competitive structure: Electricity is unlike any other commodity, and the existing economic and financial models are largely inapplicable to the study of it (Bunn and Oliveira, 2001; Ventosa, et al, 2005; Weidlich and Veit, 2008), while problems absent from or present only to limited degree in other

industries appear regularly in the electric power industry, such as price volatility and extreme market power (Stoft, 2002).

A comprehensive introduction to the history and economic dynamics of restructured power markets is beyond the scope of this paper (interested readers are referred to Stoft, 2002 or Hunt, 2002). That said an overview of a sort is called for. The following features characterize electric power markets:

1. Real-time load balancing (supply and demand must match at all times, while even day-ahead load forecasting, like all forecasting, is subject to error; these errors can have significant real-time market impacts; load balancing is also intimately related to other factors affecting reliability, including voltage and frequency).
2. Undeveloped demand-side participation (while supply is at most times highly elastic, demand is not) (Bunn and Martocchia, 2008).
3. Non-storability (economical large-scale storage technologies, other than pumped storage hydro facilities, which are limited by topography, do not exist).
4. Physical fragility (thousands of interconnected physical components make contingency analysis and very conservative contingency coverage, a crucial role of system operators).

5. High market concentration (even after divestiture efforts, the generation sector can still be considered a semi-competitive oligopoly with a competitive fringe; Entriken and Wan, 2005; Somani and Tesfatsion, 2008)⁴.
6. High integration of multiple related markets (fuel markets, capacity and ancillary services markets, bilateral contracting, as well as both day-ahead and real-time energy markets; emissions markets are becoming increasingly relevant, too).
7. Temporal disconnect between investment decisions and revenue streams (with extreme price volatility and uncertainty with regard to regulatory changes, forecasting revenues to recover costs of a twenty, thirty or fifty year investment is difficult).
8. Non-convexity of costs (total costs are characterized by extremely high fixed investment costs and relatively low variable costs).
9. Mutable regulatory regimes.

These nine features combine to create an electric power sector that is complex and poorly understood in terms of the economic dynamics, and motivate the development of novel sophisticated modeling tools for operations (for system operators), policy-making (for regulators), decision-support (for parties with vested interests in the industry) and academic research (Amin, 2002; Czernohous, et al, 2003; Koritarov, 2004; Sun and Tesfatsion, 2007; Bunn and Martoccia, 2008; Weidlich and Veit, 2008; Nanduri and Das, 2009). While the physical

⁴ A partial exception would be systems with large vertically integrated utilities, with generation as well as transmission and distribution. In these cases, generation may still be concentrated, though a significant portion of capacity will be controlled by entities without the profit-maximizing incentive of higher electricity prices.

characteristics of power systems are generally well understood, the questions of efficiency and social welfare that are raised by imperfect competition and market power as it is exacerbated by the physical and economic complexity remain. Consequently, research into market behavior under imperfect competition has generally become the focus of study du jour (Bunn and Martoccia, 2008). Of particular interest is whether prices reflect what could be expected of a fully competitive market (which we know electricity is not), and so robust, demonstrably accurate simulation tools for benchmarking to real-world data are a particularly sought-after goal.

To briefly summarize the physical operational challenges and how the nine characterizing features outlined above interact, at all times supply must be balanced with demand (#1), as electricity is still largely a non-storable commodity except at prohibitively high costs (#3), and demand is highly inelastic (#2). Complicating the picture further is the fact that the power grid is essentially a single large interconnected machine, with many thousands of moving parts, dynamic constraints and almost uncountable potential sources of failure (contingencies, in industry parlance), which, if unmitigated, can lead to cascading system failure (blackouts; #4). Consequently, system operators spend significant amounts of time on contingency analysis (CA), and operate the system with sufficient redundancy to be able to weather the first and second largest contingencies (generator or transmission line trips) without risking systemic failure. Additionally, real-time conditions (both forecast and not) can essentially create high market concentration in a small area (or system-wide in the case of a heat wave or cold snap during which real-time demand approaches system capacity), if transmission import capability is reduced through contingency or planned outage (#4 + #5).

The long planning horizon and high fixed costs (#7 and #8) along with the uncertainty and complexity that come from multiple interacting markets (#6), makes forecasting and investment decisions challenging for market participants. Also, in the short history of industry restructuring (Orders 888 and 889, accelerating nascent existing restructuring in the United States, were passed by the Federal Energy Regulatory Commission in 1996; formation of Independent System Operators or ISOs followed in 1997), regulators and policy-makers have shown a tendency to refine/tweak/enhance the markets, contributing to uncertainty for market participants. For more on the continuing debates surrounding market structure, readers are referred to (Hirst and Hadley, 1999; Besser, Farr and Tierney, 2002; Borenstein and Holland, 2002; Stoft, 2002; Joskow, 2003; De Vries and Heijnen, 2008).

As mentioned above, whether competitive markets achieve the goals set forth in the laws governing electricity and natural gas markets (the Federal Power Act and subsequent enacted laws) remains an open and contentious question. A consequence of this has been that significant resources are regularly dedicated to refining market rules. Thus the lack of certain robust proofs that competition in electricity leads necessarily to just and reasonable (a legal threshold frequently cited) outcomes both motivates the study of modeling tools and complicates the development of them.

Of special recent interest has been the development of what has been termed a “smart grid.” According to a National Energy Technology Laboratory report on “the Modern Grid,” a central component to a smart grid is advanced simulation capabilities to support system operator training under dynamic conditions (NETL, 2007). In this context, the ability to simulate both the learning and acting capabilities of all actors in a power system that is changing is critical to enhance system operator situational awareness for greater reliability (U.S.-Canada Power System

Outage Task Force, 2005; Overbye and Wiegmann, 2005). This could be considered a second parallel motivation, as robust simulation of economic dynamics becomes increasingly crucial under changing industry conditions (distributed smart-grid technologies and automated sophisticated protection systems throughout the transmission system).

1.2. Pressing Questions in Electric Power Market Regulation and Design

Increasingly, the question of whether current market designs provide the proper incentives to ensure levels of investment needed for a reliable electricity system has surrounded debates about the future of power systems under liberalized energy markets⁵. The novelty of these debates is not surprising, as electricity was long considered a good that required high levels of centralized coordination in both short and long-term decision-making in order to, as the saying goes, “Keep the lights on.” It has long been understood that the electric power system is a unique and especially delicate network requiring close constant attention and centralized operation. Thus, the concept of reliability in industry literature has long been recognized. However, responsibility for maintaining reliability was more clearly established and assigned under the former paradigm of highly regulated, vertically-integrated utilities, and reliability was ensured through investments in new and upgraded generation assets based on technical calculations of loss of load probability (or LOLP), with remuneration for capacity investments determined by state regulatory agencies (Viscusi et al, 2000).

In recent years, the question of whether market operation can indeed provide the right incentives for maintaining sufficient levels of available electric power generating capacity has

⁵ Andrew Ford (2002), showed empirically and through computer simulation how electricity markets are prone to boom-bust cycles similar to other commodity markets. This phenomenon is especially dangerous in electricity simply because of the complexity and fragility of the system as a whole. When demand outstrips supply, the integrity of the system can be compromised, price volatility notwithstanding.

come to the fore in industry trade journals, among academic researchers and industry participants (Jaffe & Felder, 1996; Besser, Farr & Tierney, 2002; De Vries & Hakvoort, 2004; Baldick et al, 2005; Cramton & Stoft, 2005). There remains no consensus on the answer, and real world experiences in market implementation vary significantly worldwide. Generally, proponents of the position that a separate mechanism is needed to ensure adequacy maintain that electricity is a commodity with particularities that prevent a pure spot market from simultaneously providing enough revenue for producers and shielding consumers from extreme price volatility (Besser, Farr and Tierney, 2002).

The question of how adequacy is defined is also a contentious issue, particularly with regard to the fact that capacity adequacy in many other industries is not an issue of public concern. Hirst and others have argued that adequacy should not be predetermined, but should be discovered implicitly by market interactions (Hirst and Hadley, 1999; Rochlin, 2004). This position holds that energy reserves, like energy itself, can be a private good. The argument goes that electricity reserve, which is generation capacity that is available on very short notice, fails the test of a public good. It can be bought and sold privately, if spot energy prices are allowed to fluctuate unconstrained. If given the opportunity end-users who currently are not exposed to time-of-use-based prices for electricity would begin to treat electricity like other consumption commodities, and the efficiencies of the market model would improve today's power systems.

Another position holds that the definition of resource adequacy really has not changed qualitatively from the days of integrated and regulated monopolies. Those in this camp generally maintain that reliability is a public good, one that the energy-only market will always underprovide. They highlight that while capacity and reserves are private goods, the extra service they provide in the form of increased system security must be remunerated through another channel

beyond the energy spot market (Jaffe & Felder, 1996; Farr and Felder, 2005). A general rule of thumb used in the United States is that put forward by the North American Electric Reliability Council (NERC), which maintains that adequacy is defined as the level of capacity necessary to ensure no more than one day of outages (lost or curtailed load) in ten years.

In some of the original literature on the subject of electricity markets, it was shown that a model consisting of an energy-only spot market for electricity could provide sufficient incentives for capacity investment and would result in both short and long term efficiency (Schweppe, 1978; Caramanis, 1982; Caramanis et al, 1982). Later work in the field built on this foundation economic model and many today argue that a pure spot market can result in a socially optimal outcome (Shuttleworth, 1997; Hirst and Hadley, 1999). Stoft (2002, 2003) maintains that while the pure model could work with defined energy price caps, determining the proper cap is difficult and setting it improperly could do more harm than no caps at all.

Many others have built upon this contention and concluded that, while the model would work given the assumptions made, energy-only markets cannot adequately ensure capacity levels necessary for system reliability due to certain fundamental characteristics of electric power systems such as are discussed below (Jaffe and Felder, 1996; Besser, Farr and Tierney, 2002; Bidwell and Henney, 2004; De Vries and Hakvoort, 2004; Cramton and Stoft, 2005). Some have highlighted that the complexities surrounding liberalized markets in modern power systems require that markets (for energy, ancillary services and capacity, where it is included) be carefully designed in concert, while the real necessity of capacity mechanisms of any form is an open question (Baldick et al, 2005). Jaffe and Felder (1996) propose that installed capacity requirements contribute to reliability, by lowering the estimated loss of load probability (the

likelihood that demand will not be met), and that this contribution cannot be reflected and valued in the spot market.

The specific characteristics that are highlighted which can lead to market failure vary. Baldick et al (2005) provide an overview of the particular characteristics of electricity markets that make application of standard economic theory to market design difficult. They and others mention that many of the assumptions made in consideration of economic models for power markets cannot be realistically made. Jaffe and Felder put forward the idea that there exist certain significant externalities to production and consumption of electricity that distort the market's operation; They argue that the spot energy market cannot reflect the changes to system reliability that result from capacity investments and load changes, and thus those who lower overall reliability are not charged while those who contribute are not remunerated. This idea of market externalities and free-ridership has been the subject of hot debate (Shuttleworth, 1997; Jaffe and Felder, 1996; Hirst, 1999; Rochlin 2004). Physical characteristics of power systems and incumbent technologies have also been highlighted as contributing factors in market failure, among them the inherited system of rate-based retail pricing which removes demand participation from the market (Stoft, 2002).

According to many authors, one specific characteristic of electric power systems that distinguishes electricity from other commodities is the inability for significant economical storage. Without the ability to store electricity, consumers who require service cannot choose to purchase excess when prices are low and forego purchases when prices rise. In this way, the primary rationing mechanism of markets cannot function and consumers may be subject to extreme price volatility (Creti and Fabra, 2003). In light of this, most power systems have been

built and operated so that consumers do not face real time prices, but pay a preset rate for energy largely without regard to time-of-use.

There is general consensus that the lack of robust demand-side participation in electricity markets causes considerable problems for market operation (Hirst and Hadley, 1999; Fraser, 2001; Borenstein and Holland, 2002; Stoft, 2002; Creti & Fabra, 2003; Joskow, 2003; Rochlin 2004; Baldick et al, 2005; Cramton & Stoft, 2005). When the idealized market model is proposed as a proper mechanism for pricing and distribution of resources in most contexts, a general assumption is made that both supply and demand participate, expressing their preferences in terms of prices and quantities. In power systems historically, on the other hand, most consumers have received service insulated from the real-time circumstances in the generation process.

In testimony before the Federal Energy Regulatory Commission, Joskow (2003) highlights certain attributes of an ideal energy market. Topping the list are the ability for consumers to see and have the ability to respond to spot market prices, to express their valuation of reliability of service in market transactions, have available various financial and contractual tools to manage their risks and have the incentive to use them efficiently. For the most part, these features are missing from power systems. Stoft (2002) and others argue that there are a number of reasons why most consumers do not make their market decisions in real time, including metering technology and the inherited retail service structure highlighted above. Fraser (2001) argues that unresponsive demand is the primary factor necessitating “socialized reliability solutions,” or capacity margins, because consumers are unable to express their value of reliability in the marketplace.

Other authors have cited the implementation of price caps in many markets (suppressing legitimately high on-peak prices), risk aversion among investors and the exercise of market power in concentrated systems as contributing to the failure of the spot market model to provide sufficient available capacity (Oren, 2003; Vazquez et al, 2002; De Vries and Hakvoort, 2004). Risk aversion, described by De Vries and Hakvoort as discouraging otherwise economic investments, can be related to many underlying factors, including uncertainty surrounding future revenue streams and possible regulatory and legislative action. Price caps are a feature in many energy markets designed to shield consumers from extreme scarcity rents that could be charged in shortage times. Many authors have suggested that this cap on prices makes recovery of fixed costs in energy markets impossible, especially so because the electricity industry has such large fixed costs relative to operating/variable costs.

Given the widespread belief among many in the field that energy-only markets cannot be relied upon to ensure generation adequacy, a number of solutions have been proposed. The most common type has been a form of capacity payment designed to provide an administratively determined proxy value of available capacity to the marketplace, and to provide producers with revenue to cover fixed costs. In its most basic form, the capacity payment is intended to address the shortfall in fixed-cost-recovery created by a non-market determination of capacity adequacy above the market equilibrium (Felder and Farr, 2005). Other proposals have included capacity obligations placed on producers or on load, reliability contracts, energy options, mothball reserves, and capacity subscriptions (Vazquez et al, 2002; Stoft, 2002; Doorman, 2003; Chao and Wilson, 2004; Oren, 2003). The creation of forward reserve markets in the New England power system has been another effort to provide additional revenue to generators that provide a

premium capacity good in the form of quickly dispatchable generation capability held in reserve (Cramton, Chao and Wilson, 2005).

Cramton and Stoft (2005) and others make the case that a capacity mechanism is necessary in most deregulated power systems because of the negligible levels of demand participation in the market for energy. They point out that past efforts at implementing capacity markets, when they were done, were often fatally flawed and that well-designed capacity markets can serve to ensure generation adequacy and remain politically feasible when coupled with energy price caps.

Experiences with implemented energy markets have varied across the globe. Some power systems have relied entirely on price signals and investment incentives from the spot market to provide the proper investments in generation capacity. Most European power systems operate this way, as does the Australian system. In the U.S., the most visible experience with energy-only markets was the widely publicized string of price spikes and rolling blackouts in California in the summer of 2000.⁶ Weare, 2003, in a report on the California crisis, estimated that total costs to the state amounted to between \$40 billion and \$45 billion, or 3.5% of annual gross state output.

In the northeastern U.S., the markets implemented in 1999 and modifications presently under consideration have included some type of extra mechanism to ensure generation capacity investment (“installed capacity” or ICAP) as well as investment in peaking capacity (“forward

⁶ It should be noted, however, that other factors contributed besides market structure: While consensus on the causes remains somewhat elusive, one contributing was that utilities who had been compelled to divest of many of their generation assets were forced through circumstances to procure energy for their customers and in some cases unexpected new customers on unfavorable terms in a chaotic environment. Weare, 2003, is an excellent source of more information.

reserve market” or FRM) used to ensure energy reserves that the pure market would undersupply (Cramton and Stoft, 2005).

Because private decisions about investment in a deregulated industry are made by market participants based on the predicted profitability of a plant (using calculations of expected fixed and variable costs and forecast revenue streams), addressing whether current and proposed energy and ancillary service markets provide the right incentives is critical. The original literature on electricity pricing in spot markets was predicated on certain assumptions of underlying market fundamentals. Debate has focused in recent years on whether many of these assumptions are valid. Because of the centrality of electricity to economic growth and prosperity, as well as the potential risks posed by insecure power systems, the concern of whether implemented market models can ensure socially optimal levels of reliable service is central to policy discussions. The next section introduces some contemporary applications of agent-based modeling to electric power markets.

1.3. How Agent-Based Modeling Has Been Used in Electric Power Market Applications

Ventosa, et al, 2005 distinguishes three types of modeling commonly used in the study of electric power markets: optimization models, equilibrium models and simulation (or agent-based) models. Although there has been some overlap, this categorization serves an important purpose in offering a foundation for researchers working in a field that has occasionally lacked commonality of methods to facilitate reproducibility (this criticism has been lodged by Weidlich and Veit, 2008).

Optimization models restrict the analysis to the decision process of a single firm within a complex market, and are the least explored of the three types. The second type, equilibrium models, attempts to solve a mathematical program representing the decisions of all participants in closed form. These are typically either Cournot (quantity-based) or Supply Function Equilibrium (quantity and price-based) models, and have received significant attention both independently and as part of benchmarking efforts for the third type, simulation (Klemperer and Meyer, 1989; Hobbs and Helman, 2004; Waltman and Kaymak, 2008). Because simulation is the approach to be used here, interested readers are referred to Ventosa, et al, 2005 for more literature on non-simulation electric power market modeling.

The third modeling approach, simulation, here referred to synonymously as agent-based modeling (ABM), or sometimes as agent-based computational economics (ACE), has received significant attention for its general framework and the ability to surmount many of the computational limitations of equilibrium models (Fagiolo, et al, 2007). Many equilibrium models suffer from the inability to model more than very simplified transmission systems and few active decision-makers (Hobbs and Helman, 2004). On the other hand, ABM, while potentially a very powerful modeling tool usable for almost every conceivable application related to power market analysis and transmission systems of complexity approaching the real world, has some crucial weaknesses (which are all-to-infrequently emphasized in most ABM literature), and has received its own share of criticism. Most often, these are criticisms of the modelers, whose results are presented without a) benchmarked results or b) robust proofs of the relevance of the results (Weidlich and Veit, 2008). On the other hand, with the spread of ABM concepts among applied economics researchers, greater attention has recently been paid to these omissions, including renewed efforts at benchmarking results against real-world observed

phenomena (Fagiolo, et al, 2007) and defining common frameworks for easy comparison (Midgley, et al, 2007; Marks, 2007). Leombruni and Richiardi, 2005, offer a useful summary of remaining issues with ABM methodology in their exploration of why it hasn't received greater attention from top-ranking journals.

CHAPTER 2

AGENT-BASED MODELING: THE CONTEXT OF THIS RESEARCH

The project below is an application of agent-based modeling using reinforcement learning (a popular machine learning approach). As described above, electric power market modeling has taken three general forms, although there has been some overlap: Optimization models, equilibrium models and simulation models (Ventosa, et al, 2005). Simulation models, of which the proposed research project is one, are useful for situations in which the complexity of the desired model makes equilibrium analysis impractical. In a simulation model, firms are represented as autonomous agents, where the structure of the agents is general (much research in the collateral field of machine learning regards agent structures and the performance of these agents under various types of environments; see below for greater depth of coverage on agent models, or Kaelbling, Littman and Moore, 1996 or Sutton and Barto, 1998). A fourth model that has received little attention and is omitted mention in Ventosa, et al, 2008, is human-subject experiments, in which human players take on the role of a generation firm. A possible reason for this for this omission, and a motivating factor in agent-based modeling, is that the learning curve for operating a portfolio of assets even in a much simplified market is steep for human-subjects, making experiments of this form of limited utility (Weidlich and Veit, 2008).⁷

Agent-based modeling is a simulation technique particularly suitable to the study of complex systems, such as electric power markets, for which analytical, closed form, solutions do not exist or are impractical to discover. While there are several ABM approaches that have been

⁷ This isn't to say that the learning curve for a computer agent is much shorter, but an ideal adaptive agent, lacking in heuristics that humans can use, for better or worse, may be able to learn the complex tasks faster due to computational speed advantages.

applied to modeling electricity markets, the most prevalent use adaptive agents with some type of a reinforcement learning algorithm for online learning. Reinforcement learning is significantly founded on Thorndike's Law of Effect, in which those actions an agent (or animal) takes which result in immediate positive reward will be repeated more often and vice versa. A drawback to ABM, which has been mentioned throughout the literature, is precisely that proofs of optimality or even demonstration of realism through benchmarking are difficult and these limitations have received little attention relative to the substantial reported work in the field (Weidlich and Veit, 2008)⁸.

However, a counterbalancing factor is that much ABM research borrows lessons and insights from substantial existing research in the fields of artificial intelligence and reinforcement learning (Bunn and Oliveira, 2001; Bagnall and Smith, 2005; Veit, et al, 2006). An example in the field is Waltman and Kaymak, 2008, in which they demonstrate both analytically and through ABM in a Cournot market that Q-Learning agents collude to raise prices above a competitive equilibrium.⁹

At the root of ABM are the agents that act autonomously according to their own individual algorithms and the rules of the model. The general structure involves an agent that in a single stage or iteration of the simulation takes an action according to its policy, receives information and a reward (or calculates the reward itself from the environment message passed to it) and updates its policy by which future actions will be selected. ABM is especially useful

⁸ Benchmarking is difficult for another reason, namely that industry data are often hard to retrieve, and that market outcomes are significantly impacted by operating procedures which, though they attempt to balance reliability and cost reasonably, are constantly evolving and differ from control area to control area.

⁹ Occasionally, the term Q-Learning is used synonymously with learning that simply utilizes a general reinforcement update algorithm. Strictly speaking, Q-Learning relates more to dynamic programming, in which reward values are backpropagated from values of subsequent states and actions to those of antecedent states and actions. Q-Learning particularly is "off-policy", contrasted with "on-policy" learning, but this distinction is beyond the scope of this report. See Sutton and Barto, 1998 for more (available free online at <http://webdocs.cs.ualberta.ca/~sutton/book/ebook/index.html>)

for studying complex systems that are not easily analyzed with traditional modeling techniques, in particular in the case of economic systems, where the real subjects of interest are often the emergent macro-phenomena that result from interactions between hundreds, thousands or millions of economic actors. While laboratory game experiments using human participants as generators can discover market flaws at risk for exploitation, there is a limit to how complex a system can be analyzed in these economic experiments. Especially in the case of real-world power systems, suppliers often control multiple plants in different locations on the transmission grid (often in multi-settlement or multi-stage markets with more than one product) and the learning curve for a human playing the part of a supplier with a portfolio of these plants can be steep, often prohibitively so.

Adaptive-agent-based modeling is especially suited to analyzing electricity markets (with certain important caveats, high among them the exponential relationship between environment complexity and task complexity) because participating agents can often quickly search through their policy-space for successful strategies even given a potentially very complex problem to address. Many obstacles exist to reaching this goal of building agents able to learn strategies in environments of great complexity (such as the non-stationary nature of an environment composed of more than one learning agent), and some of these are detailed in this report, but overall experience to date has shown promise in this field. The motivation for the work described in this report comes from the observation that the field of agent design is relatively new and many potential avenues of research into agent design (and the suitability of particular designs to electric power market modeling) have not yet been exhaustively explored.

Generally there have been two approaches researchers have taken to ABM over the last couple of decades. One has been to attempt to design agents that approximate the behavior of

human-subjects (Erev and Roth, 1995), while the other has disregarded benchmarking agent behavior to those of humans and explored emergent phenomena in environments with many agents following simple heuristics in decision-making (Bunn and Bower, 1999, for example). The results of these early experiments showed the potential of ABM to model complex phenomena emerging from simple foundations. Approaching the field of ABM from the direction of benchmarking, Axelrod (1987) used genetic algorithms to model agents participating in an iterated prisoner’s dilemma. Erev and Roth (1998) built upon this work with agents participating in simple games in a study on using learning algorithms to predict how humans learn to play matrix games¹⁰ with mixed-strategy equilibria. More recently, Abdallah and Lesser (2008) have provided further analysis of the performance of a community of reinforcement learning agents, evaluating the convergence capabilities of several learning algorithms on some standard matrix-payoff games with pure and mixed-strategy Nash Equilibria. Nicolaisen, et al (2001) explores applications of a modified Roth-Erev (MRE) reinforcement learning algorithm in electricity market simulations, and has the advantage of using agent models that have been demonstrated to realistically approximate human behavior in simple applications in relatively complex power systems (incorporating a multi-node transmission model; Somani and Tesfatsion, 2008). Other applications of ABM to electricity markets are covered in (Bower and Bunn, 1999; Bunn and Oliveira, 2001; Bagnall and Smith, 2005). Bagnall and Smith, 2005, applies ABM to

¹⁰ Littman, 1994, describes a matrix game, in the context of game theory, as simply a deterministic reward function $R_{i,j}$ for the first of two agents choosing action i and their opponent choosing action j . And offers the zero-sum game “rock, paper, scissors” in their exploration of Q-Learning in a multi-agent context:

(Agent 1 reward, Agent 2 reward)	Rock	Paper	Scissors
Rock	(0, 0)	(-1, 1)	(1, -1)
Paper	(1, -1)	(0, 0)	(-1, 1)
Scissors	(-1, 1)	(1, -1)	(0, 0)

the UK electricity markets under the New Electricity Trading Agreement (or NETA) using relatively sophisticated learning classifier system agents that incorporate a genetic algorithm applied to encoded rules of behavior.

Agents that can develop policies through trial-and-error interaction with their environments are especially interesting for use in agent-based modeling. Well-designed adaptive agents have the potential to independently discover strategies in complex environments beyond what can be reasonably hard-coded into heuristic agents or expected of human game-participants. In this way, some of the key benefits of human-participant experimental economics (discovery of unforeseen strategies/policies) can be incorporated into more robust modeling tools for real world applications. On the other hand, sufficiently complex adaptive agents can retain the ability for direct optimization (when such programming subroutines are part of their action sets) in decision-making that human participants cannot (Bagnall and Smith, 2005 describe their process of agent-design clearly; Dariani, et al, 2008 treat precise design of both adaptive and non-adaptive agents with several parts, of which a value function for determining a policy can be one). This capability is especially relevant for ABM applications to electricity markets because a successful power supplier strategy may require close coordination between generating plants controlled by the agent at different locations on the power grid. Relatively sophisticated agents capable of solving these complex problems using mathematical programming with equilibrium constraints (MPECs) have been implemented in modeled electricity markets (an example of combining optimization modeling with simulation; Entriken and Wan, 2005). There may be significant promise in combining agent learning with these existing mathematical programming tools available to an agent, but the author is unaware of this sort of application to date.

Essentially an adaptive agent competing in a modeled market is faced with the task of discovering strategies that will maximize their profits (or in some cases attempt to optimize with respect to multiple objectives) by searching the space of potential policies through online interaction with their environment (other agents and the market clearing mechanism or market module). The complexity of agents that have been implemented in market models varies widely, as do the broad agent structures used.

In early work on reinforcement learning, optimality proofs were demonstrated for an agent learning in a stationary environment (not necessarily deterministic, but stochastic with fixed probability distributions). When the task faced by a learning agent is non-stationary (which it is whenever a second learning agent is part of the market environment), the optimality proofs provided in many reinforcement learning agent research (Bertsekas and Tsitsiklis, 1996; Sutton and Barto, 1998) may not exist.

2.1. Challenges and Opportunities

There are many challenges to developing robust electric power system models with adaptive agents, many of which were described above. Perhaps the greatest challenge involves building confidence in the results. Section 1.1 introduced the most relevant complicating factors of electric power markets and provided some treatment of the motivations for this research. Here we summarize the greatest challenges facing ABM researchers in exploring dynamics of realistic electric power markets in their models.

1. Electric power systems are complex.
2. The dynamics of competitive electricity markets are poorly understood.

3. Analytical solutions are mostly or completely unattainable for optimal behavior in even relatively simple models of electric power systems (this is especially true for multi-agent models and for other models incorporating non-stationary phenomena).
4. Human-participant experiments are of limited use in electric power market models, due to the complexity of the task an agent with a portfolio of generating plants participating in multiple inter-connected markets faces.
5. Agent-based modeling is a relatively new tool for modeling economic systems, but offers promise as a compromise environment for evaluating market structures and participant behavior.
6. Most models lack realism (real world power markets, as described above, are complex, interconnected, multi-faceted and mutable).
7. Lack of realism of the agents (even the most sophisticated agents have limitations that become pronounced as the complexity of the task environment increases).
8. Proprietary nature of power system data.
9. Non-scalability of existing agent learning models to realistic scenarios. Even those multi-agent reinforcement learning models which demonstrate convergence in non-stationary environments have only been applied to very simple games.

This research project is motivated by several interconnected factors, and is an attempt to address as many of these issues as possible. First, there is a well-understood need to have

electric power market modeling tools that can reliably and quickly approximate real-world conditions and predict market behavior under variable conditions. To a significant extent, research to date has either been devoted to modeling the physical power system as accurately as possible and omitting strategic participant behavior (including state estimators used in ISO system operations, and generator commitment and dispatch algorithms are examples), or has sacrificed the detail of most physical power systems to focus on agent behavior. Some exceptions that attempt both include work reported in Conzelman, et al, 2004, Bagnall and Smith, 2005, Sun and Tesfatsion, 2006, and Somani and Tesfatsion, 2008. The model detailed here incorporates, as mentioned before, a multi-node transmission model with locational marginal prices, features central to the Standard Market Design put forward by the Federal Energy Regulatory Commission in its early orders mandating a move towards open competitive electric power markets. Other features that will be necessary to incorporate into future models include multi-settlement systems, non-linear commitment and dispatch algorithms, dynamic load, demand-side participation and more.¹¹

Second, there continues to be insufficient acknowledgement of the particular strengths and weaknesses of reinforcement learning-based agents used in ABM research. This has likely been caused by the inaccessibility of much reinforcement learning research to those who are not theorists in the field, combined with a possibly unjustified confidence among electric power market modeling practitioners in the abilities of these agents in complex multi-agent

¹¹ Veit, et al, 2006 do incorporate these features into a model similar to the one used in our research, but the complexity introduced there is not explicitly addressed for its impact on the dimensionality of the agents' tasks. In general learning agents must sample from their action set a very large number of times (positively proportional to the number of possible environment states) in order to learn a task. While non-stationary task environments, such as two interacting learning agents comprise, admittedly do not have convergence or optimality proofs, as the ratio of possible environment states to the number of rounds run in a simulation increases, so does the chance that particular states are never encountered, which intuitively adversely impacts learning in a multi-agent environment.

environments (this author has been one). While recent attempts have been made to demonstrate analytically the applicability of Q-Learning agents (such as Waltman and Kaymak, 2008) far too little attention has been paid to robustly demonstrate the advantages of adaptive agents in electric power market models.

Any attempt at modeling a complex environment must begin with an understanding of both the capabilities and limitations of the tools available as well as a vision of the endgame (what features an ideal product would have). The obstacles to implementing learning algorithms in models such as this one, that are well known in the reinforcement learning community, include the challenges of representing continuous and high-dimensional environments for an agent to develop realistic situational awareness (see “the curse of dimensionality”; Bertsekas and Tsitsiklas, 1996), and the significant amount of time, in terms of both modeling rounds and computing time, needed for optimal behavior to be learned (in cases where optimality proofs exist) or convergence to be achieved (if it does at all). For these reasons researchers in electric power market modeling commonly sacrifice some aspects of an ideal agent structure for tractability (with good reason, although often these sacrifices are left unacknowledged). A frequent sacrifice made is to omit state awareness and take advantage of evidence provided in earlier work on the realism of agent-learning and behavior with respect to human learning (Erev and Roth’s work on predicting human game playing with reinforcement learning agents is widely cited, and widely used; Sun and Tesfatsion, 2006, Veit et al, 2006, Somani and Tesfatsion, 2008). However, without state awareness, agents are essentially blind to market conditions.¹² A compensating effort that is common is to benchmark results of a multi-agent model with real-

¹² This shortcoming is a factor in the present version of the market model, but state awareness (and all the challenges accompanying it) is a feature that will have to be incorporated in future multi-agent reinforcement learning research on electric power markets.

world data, although comprehensive attribution of specific agent features and modeling assumptions to simulation results is lacking (Fagioli, et al, 2008). Another challenge is the limited availability of much power system data, as they are owned by private entities including system operators, transmission owners and electric utilities.

The reinforcement learning literature is particularly valuable for its attention to categorizing task environments that agents face. From the simple Markov decision process (MDP, which has the Markov property) to partially observable Markov decision processes (POMDPs) and non-stationary POMDPs (the latter two of which don't have the Markov property), the limitations imposed on the research are explicitly acknowledged. From this literature, it is possible to glean insights into the task that a representative agent in a market model environment faces (Littman, 1994 does a good job of this). In most research applications that utilize reinforcement learning, the task environment isn't explicitly described (Xiong, Okuma and Fujita, 2004; Rahimiyan and Mashhadi, 2008). This is the case with almost all ABM models cited here. This makes the extension of work such as Abdallah and Lesser, 2008, (in which the task environment is explicitly described in terms of joint reward functions) to complex environments difficult, though in their paper, the authors cite some promising directions being explored in multi-agent learning applications in iterative matrix games involving a general number of agents, beyond the two that Abdallah and Lesser use (Tuyls, et al, 2006).

Some promising recent research has been done in the RL field on so-call multi-agent reinforcement learning (MARL) applications (of which this research project is one; Krause, et al, 2005; Abdallah and Lesser, 2008). In these lines of research, the non-stationary stochastic nature of the single-agent's task (reward function) is placed within the context of a community of agents, in which case the payoff to all agents is a deterministic function of the agents' collective

actions (a mapping from joint-policy space to a vector of rewards to each agent).¹³ Abdallah and Lesser report on work they've done to explore learning processes and convergence of various learning algorithms in non-stationary n-armed bandit problems (the single-agent perspective of multi-agent games with known pure or mixed-strategy Nash Equilibria). The primary difficulty with borrowing too much from this line of exploration is that many of the most interesting applications are ones where payoff functions of joint-policies are not clear, particularly in the as-yet unapplied electric market model incorporating realistic features of existing power systems (multi-node transmission system with locational market pricing, multi-settlement rules, ancillary services, stochastic loads, price-responsive demand, etc.).

Often, in much of the electric power market ABMs, explicit framing of the task environment is left out because either classification is difficult or would be unhelpful even if the class were known. This is the case with almost ABM to date applications (present company included) that use reinforcement learning. Reinforcement learning has been proven to discover optimal behavior only under extremely restrictive assumptions, including properly-calibrated parameter adjustments and a stationary environment (Sutton and Barto, 1998), some that are necessarily missing from all simulations in which more than one adaptive agent is used (such as Xiong, Okuma and Fujita, 2004; Rahimiyan and Mashhadi, 2008). This is because the presence of another adaptive agent makes the environment non-stationary and non-Markov.

To the author's knowledge, little of the most recent research in reinforcement learning (such as Abdallah and Lesser, 2008) has been utilized in agent-based modeling of electric power

¹³ The task environment in the model described below is technically a matrix game, although with 2^{25} possible combinations of actions, explicit framing of the joint-reward function under varied topologies and plant-node assignments is prohibitive. The otherwise excellent model presented in Veit, et al, contains an [estimated] $3.25 * 10^{128}$ possible joint actions and environment states, making it impractical for agents to fully learn tasks, unless generalization is used in the learning algorithms (see Sutton and Barto, 1998 for very good treatment of generalization in reinforcement learning).

markets. This is a primary motivator for the research below, although the omission of state representation leaves something to be desired in modeling power markets, particularly those models that would incorporate multi-settlement market features and inter-temporal constraints.

There is a necessary trade-off when you move from a task environment with a reinforcement learning agent that is embedded in an environment having the Markov property to one that doesn't. When you leave the Markov property behind, you enter the realm where success is measured by convergence and benchmarking against real-world data, and this is often tenuous.

CHAPTER 3

THE ELECTRICITY MARKET MODEL

Green and Newberry (1992) and others have made the case for using quantity-based or supply function modeling, as equilibrium results of Bertrand (price-based analysis) models show competitive outcomes in even very concentrated markets, which is counterintuitive and contradicted by results from other models (Banal-Estanol and Micola, 2009). Another case that can be made for using quantity-based or supply function-based models is that they have greater realism: In most extant competitive electric power markets, the exercise of market power through price-manipulation is closely monitored and mitigated through robust rules and market monitors endowed with non-trivial referral authority. Less well understood is the effect of physical withholding on market outcomes, particularly because generator outages and reductions are a common occurrence, even unplanned ones (or “forced outages”). Therefore, absent blatant misrepresentation or an incriminating paper trail, anticompetitive manipulation of physical parameters is challenging to catch and to prove, although penalties for misrepresentation of an asset’s availability (whether in-service when it is not, or vice versa) are nontrivial. For this reason, we constructed a market model in which agents compete on quantity rather than price (an advantage is that the path here isn’t entirely untrod; Waltman and Kaymak, 2008, also explored using Q-Learning agents in a Cournot model, as did Veit, et al, 2006). The market model contains a clearing engine built upon an open-sourced mathematical optimization package (LP Solve¹⁴), which has a multi-node transmission system with locational marginal pricing (LMP).

¹⁴ lpsolve.sourceforge.net/5.5/

The motivation for this research is to further development of flexible, scalable and reliable market models. The promise of these goals are models that can help training of system operators, regulatory/policy-making, refining of market rules and general understanding of the short-run and long-run dynamics of competitive electric power markets under general initial conditions.

3.1. Where This Research Fits: What's New Under the Sun?

The agent-based model we developed and describe below finds good company among those that have been reviewed and reported on in the contemporary literature on modeling electricity markets. Inspired by early results of agent-based electric power market modeling (such as Bunn and Bower, 1999) and reinforcement learning applications in stationary and cooperative tasks (Littman, 1994; Sutton and Barto, 1998, Nicolaisen, et al, 2001, and others), development of linearized direct-current optimal power flow models (Chao and Peck, 1996), first-hand experience with system operations¹⁵ and recent work exploring convergence of multi-agent reinforcement learning models in matrix games (Abdallah and Lesser, 2009), the model described here attempts to incorporate lessons from disparate fields to both identify avenues of improvement to the existing body of electric power ABM research, and demonstrate tractability in a simplified case. The intention is to contribute to this exciting field in several ways:

First, the economic model used in this research project closely resembles the real-world conditions of the power industry we are attempting to simulate. As Banal-Estanol and Micola, 2009, and others have noted, real-world generator owners typically bid their units in piecewise tranches of capacity, as our agents do. Additionally, operating costs of most widely-used

¹⁵ Some of the work reported here was done during a period of time when the author was also an employee of ISO New England Inc.'s Internal Market Monitoring Unit (INTMMU). No non-public information is disclosed here.

generator technologies are known to be within a narrow range around a linear relationship with fuel prices. With most power systems incorporating robust price-mitigation measures, market manipulation through price manipulation is both difficult and rare. A more practical and difficult to detect method of manipulation involves capacity withholding, either through complete plant outages, partial plant reductions or misrepresentation of operating parameters. It is our contention that more research on quantity-based models is appropriate.

Second, as several other researchers, but not all, have also done, we incorporate a transmission system into the model. We show that this is both practical and the agents successfully converge to equilibria that take advantage of transmission congestion to increase profits.

Third, although we were not able to incorporate those newly developed learning algorithms that have demonstrated success in complex non-cooperative games where others, such as the Q-Learning update algorithm we use, fail, the model is immensely flexible for future extension. As additional market features are incorporated into agent-based models (day-ahead commitment and dispatch, with real-time balancing, ancillary services and demand-participation), it is expected that equilibria (if existence is shown) will be more difficult for agents to discover, in which case new and different learning algorithms will be needed.

3.2. Predecessor Model

The first version of the market model described in this report consisted of a simple electric power market with two or more agents competing to supply fixed demand by submitting price offers for their full capacity up to an offer cap of \$20¹⁶. There was no transmission model

¹⁶ These results were earlier presented in Miksis, 2006.

(like the current version described below contains). The decision process of the adaptive agents (to distinguish from the fixed-action agent that submitted a single static price offer every round) was to modify their previous round price offer by one of several multiplier factors (0.5 to 2). Consistent with previous research in reinforcement learning agents in a stationary task, initial simulations showed that an adaptive agent could successfully discover optimal behavior in several load scenarios. In each round, the market module solicited actions from each agent and ran a least-cost optimization to match supply with demand. Three scenarios were run and reported on:

The first scenario involved a single adaptive Q-Learning agent (with a cost of \$0) competing against a static opponent offering a constant price. Each agent had a fixed capacity of 1,000 MW, while load was also 1,000 MW. In this case, the agent offering the lower price would clear and serve the entire load, while the other would not.



Figure 1. Q-Agent and \$10 Static Offer Opponent, 1,000 MW Load

The adaptive agent's profit function is piecewise linear, increasing from an offer of \$0 (\$0 profit) to \$9.99 (\$9,999 profit) and dropping to \$0 profit at offers higher than \$10,000 (tie-breaking was 50/50 random, so at \$10 exactly, expected profit was \$5,000). The results are shown in Figure 1; the non-linear series is the adaptive agent's offer series (moving average to address noise during the simulation). The agent learned relatively quickly to just underbid the static price agent.




Figure 2. Q-Agent and \$10 Static Offer Opponent, 1,400 MW Load

Figure 2 shows the second scenario in which load is raised to 1,400 MW; the other parameters are unchanged. While both agents will necessarily clear and the uniform clearing price will be set to the higher offer, the agent that sets price will clear only 400 MW. Therefore, the choice for the adaptive agent is really between offering at the price cap and receiving a profit of \$8,000 ($\$20 * 400 \text{ MW}$) versus offering below \$10 and clearing its full capacity, for a profit of \$10,000. As the figure shows, the agent quickly learned to make the second choice.



Figure 3. Q-Agent and \$10 Static Offer Opponent, 1,600 MW Load

Figure 3 shows the third scenario, in which load is raised to 1,600 MW. In this case, the optimal behavior is to bid at the price cap of \$20 (profit = \$12,000) rather than bidding low (profit = \$10,000). The agent successfully learned this. The results showed the promise of an adaptive agent using a Q-Learning algorithm in a stationary scenario with a uniform clearing price market structure similar to real-world electric power markets.

3.3. Current Market Model Description

The new version of the market model used in this research is significantly more complex than that used in Miksis, 2006. It is built with the same fundamental design as the real-time scheduling, pricing and dispatch (SPD) algorithms used by power system operators to manage least-cost dispatch of electric power generating stations and other assets; it is similarly comprised

of a transmission system, loads and generating assets, with physical and economic parameters. At a basic level, power transmission systems are composed of nodes and lines, each of which connect two or more nodes (infrequently more than two), while assets are located at nodes on the grid. In reality, most power systems are made up of hundreds or thousands of nodes, lines, generators and other assets. The model used in the simulations described later is a much-simplified representation of an actual power system, including only 5 nodes, 6 lines and 25 generators (5 agents controlling 5 generators each). Figure 4 shows the system used in this research project.

As described above, power systems are extremely fragile machines. In order to maintain reliable delivery of power, system operators must perform least-cost dispatch optimization while respecting many physical constraints: Generators must be dispatched at or below their maximum capacity (and occasionally there are nonlinearities imposed by the binary on/off decisions with a unit operating at zero, at a minimum output, somewhere above minimum, or at maximum), power lines must not transfer electricity in excess of their thermal limits and load at every node on the system must be matched with supply (Kirchhoff's current law declares that nodal balance is respected: the sum of injections, line inflows, line outflows and withdrawals equals zero). Because electric power flows through a network according to physical laws, and for the most part cannot be directed from point to point along a particular path, these laws must be translated into constraints in the dispatch model, so that power injections and withdrawals at each node are made to ensure power flows respect all thermal transmission line limits. A key point is that in a transmission network a unit of power (MW) injected at one point, conventionally called a source, and withdrawn at another, called a sink, will flow partially along every possible path between the two locations; how much energy flows along each path is determined by the line resistance (in

the case of direct-current, or DC, models) or impedance (in the case of alternating-current, or AC, models)¹⁷. To aid in the dispatch of the system, these line resistances for a given network are used to pre-calculate factors, known as power transfer distribution factors (“PTDFs”), for each triplet (source node/power line/sink node) to determine how much (0 to 1, or 0% to 100%) of a MW injected at the source and withdrawn at the sink flows along the line.

A further simplifying process decomposes this factor into two others, called generator shift factors (or sometimes just shift factors), representing first, the PTDF from the sink to a chosen reference node and second the PTDF from reference node to the sink (the chosen reference node uniquely determines every GSF, but does not change the resulting power flows or prices). This is done to simplify the calculation of power flows when modeling line constraints.¹⁸

While real-world power systems must be dispatched respecting hundreds or thousands of constraints (including not only thermal line limits, but also generator capacity constraints, upward and downward ramp limitations, post-contingency reliability line transfer limits¹⁹, stability/voltage limits and more), the model used in this research contains only 25 generator capacity constraints and 12 thermal line limit constraints (lines are bi-directional, so a constraint must impose lower and upper bound flows), plus the energy balance equality constraint.

¹⁷ The market model used in this research is, for practical reasons, built upon a DC system.

¹⁸ A GSF for a node-line pair allows the linear program to respect line constraints using a single GSF coefficient on each decision variable; because of the system energy balance equality constraint, the dispatch will ensure that each MW injected and “sent” to the reference node is matched by an offsetting MW “sent” from the reference node to the sink/withdrawal node.

¹⁹ We have already described the fragility of power systems. Because power flows adjust almost instantaneously in response to changes in the system, if a line “trips” out of service, the power that previously flowed on that line is directed along other paths. This can overload these lines, which can lead to cascading failure as protection systems will automatically trip lines that are overloaded to avoid permanent equipment damage. Consequently, power flow on lines are limited not to their actual thermal limit rating, but to the level at which the lines that would pick up their flows could carry if the line in question did trip. For this reason, in industry parlance, line constraints are sometimes referred to as “line x onto line y”.

The simulations performed as part of this research are repeated single-round runs of the market model, which means that inter-temporal ramp constraints are not needed. Additionally, the dispatch levels of generators are lower-bounded at zero, so that the non-linear binary commitment decision is not a part of the model. With these restrictions, the problem is linear, and as such, is easy and quick to solve with off-the-shelf linear programming packages.

3.3.1. The Power System Model

Having described the general components of a simplified power system model, we next introduce the model used in this research: first, the physical system, then the market mechanics (3.3.2 and 3.3.3) and lastly, the agents (3.3.4). As mentioned, there are 5 nodes in the system model, with 6 lines connecting them in a network (the nodes are circles and the transmission lines are squares):




Figure 4. Modeled Power System Diagram

The resistances used for the lines were borrowed from a public system model used in training tools by the ISO New England Inc. and PJM regional transmission organizations.²⁰ Table 2 in Appendix B shows the relevant data and the generator shift factors derived from the resistances. The data in the table are not significant unless understood in context. What they mean, for example, is that a MW injected at Node 3 and withdrawn at the reference flows through the network as shown in Figure 5 (reading down the 3→2 column, we see that 16% flows in the direction from Node 3 to Node 4 along Line 5; because of nodal balance at Node 4, this 16% continues along Line 2 to Node 0, etc.; lines are directional only for the purpose of calculating power flows).

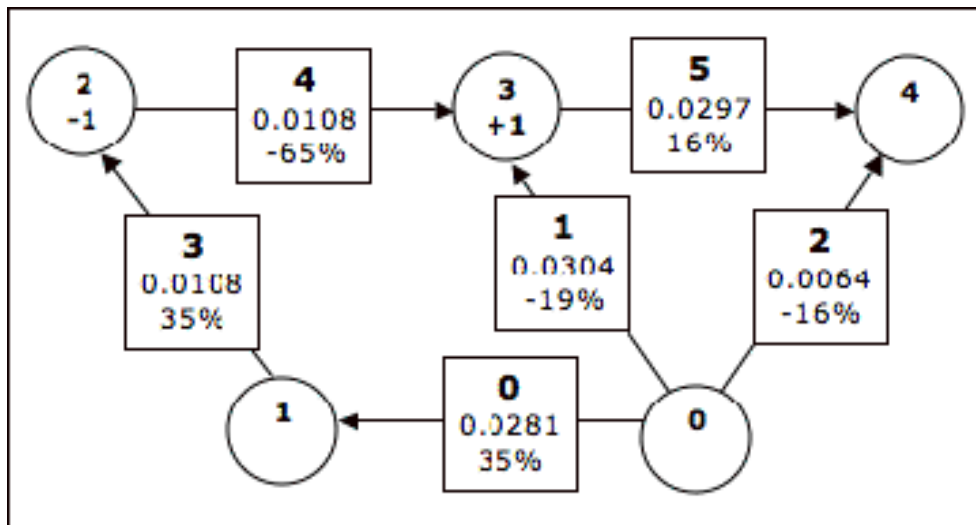


Figure 5. Line Flow Demonstration with Node 3 Injection

Because GSFs are additive, to calculate flows on Line 5 caused by an injection at Node 3 and a withdrawal at Node 4, we add 16% ($GSF_{Node3,Line5}$) to 32% ($-1 * -32\% = -1 * -GSF_{Node4,Line5}$, because we are withdrawing instead of injecting at node 4). This equals the PTFD

²⁰ See <http://www.iso-ne.com/support/training/5bus/index.html>.

on Line 5 from an injection at Node 3 and a withdrawal at Node 4. It is then clear that, although GSFs are uniquely determined by the choice of the reference node, the choice is arbitrary for purposes of calculating power flows and prices.²¹

Now that we have an understanding of the mechanics of the 5-node model used in the simulations, we will discuss the three other components of the complete market model, the least-cost dispatch, pricing and congestion, and the agents.

3.3.2. Least-Cost Dispatch

The problem of minimizing total dispatch cost is linear. Each decision variable corresponds to output of a generator, while the cost coefficient of each variable is the price offer of that unit (used for brevity synonymously with generator; more on this in the “Agents” section, 3.3.4.). The objective function is therefore:

$$\sum_{a,p} c_{a,p} x_{a,p} \quad a \in A \text{ (set of agents; 0 to 4); } p \in P \text{ (set of plants; 0 to 4)} \quad (\text{i})$$

The most important constraint in all power systems is the energy balance equality constraint (supply ***MUST*** equal demand at all times). Therefore, the first constraint is:

$$\sum_{a,p} x_{a,p} = \sum_n \text{nodalLoad}_n \quad (\text{ii})$$

As \vec{c} , $\overline{\text{nodalLoad}}$ is also an input to the optimization, determining as implied, the load (electric power demand) at each of the 5 nodes n . The next set of constraints ensure that the output of each unit (a,p) is not greater than its capacity $\text{capacity}_{a,p}$.

$$x_{a,p} \leq \text{capacity}_{a,p} \quad \forall a \in A, p \in P \quad (\text{iii})$$

²¹ This is not strictly true in AC systems, where losses are modeled. Litvinov et al, 2004 discuss the choice and consequence of reference nodes.

Finally, there are the line limits (maximum flow in either direction, modelled as positive and negative flows in the convention used above), which are similarly inputs to the simulation. The variable nr is used to denote a non-reference node $\in (0,1,3,4)$, as we omit the reference node from line flow calculations. The matrix GSF is shown in Table 2 (Appendix B).

$$\sum_{nr} GSF_{line,nr} NI_{nr} \leq lineLimit_{line} \quad \forall line \in L \text{ (set of lines; 0 to 5)} \quad (iv)$$

$$\sum_{nr} GSF_{line,nr} NI_{nr} \geq -lineLimit_{line} \quad \forall line \in L \quad (v)$$

NI_{nr} is the net injection (possibly negative, in which case it's a net withdrawal) at the non-reference node nr .

$$NI_{nr} = \sum_{a,p} x_{a,p} n_{a,p,nr} - nodalLoad_{nr} \quad \forall nr \in (0,1,3,4) \quad (vi)$$

$$\begin{aligned} \text{where } n_{a,p,nr} &= 1 \text{ if plant } p \text{ of agent } a \text{ is located at node } nr \\ &= 0 \text{ otherwise} \end{aligned} \quad (vii)$$

The last set of constraints is that which ensures non-negativity; although in real-world power systems there are assets that can become net consumers (pump storage hydroelectric facilities and flywheels are two examples), the units in our model are assumed to be traditional thermal power plants, dispatchable from zero to capacity:

$$x_{a,p} \geq 0 \quad \forall a \in A, \quad p \in P \quad (viii)$$

Altogether, (i) through (viii) comprise the least-cost dispatch algorithm.

3.3.3. Pricing and Congestion

In most modern power systems, prices are determined nodally (so-called locational marginal pricing, or LMP). The term ‘‘locational’’ indicates that prices differ across the power

system, while marginal indicates that the system uses uniform pricing²²; the price at any location is exactly equal to the marginal change in the objective function caused by an incremental MW delivered there (net withdrawal off of the system).²³

In a DC model without losses, the normal case where no transmission line constraints are binding has the marginal cost to deliver an additional MW anywhere on the grid equal to the shadow price of the energy balance equation, which is always binding (it is often denoted λ). Intuitively, this will be equal to the cost coefficient of the unit providing the last increment of energy (in the case of no binding transmission constraints), equivalently the most expensive unit online; this unit is known in the industry as the marginal unit.

In cases where one or more transmission lines are binding, there will be more than one marginal unit. It is a truism that, except in the degenerate case of identically-priced and identically-located units, there will always be one greater number of marginal units than the number of binding transmission constraints. The shadow prices of binding transmission constraints are traditionally labelled μ . The formula for calculating LMPs is:

$$\text{LMP}_{nr} = \lambda + \sum_k \mu_k \text{GSF}_{k,nr} \quad \forall nr \in (0,1,3,4) \quad (\text{ix})$$

where $\text{GSF}_{k,nr}$ is the generator shift factor corresponding to the line whose upper or lower limit is binding and non-reference node nr . For the reference node, the LMP is always equal to λ (or “system λ ”), the marginal cost to deliver an additional MW to the reference node (there is no congestion component of LMP at the reference).

²² Xiong, Okuma and Fujita, 2004 have addressed a comparison between pay-as-bid and uniform pricing in their multi-agent simulation research. Kahn, et al, 2001 provides an excellent overview of the comparison.

²³ It will be noted that revenues collected by generators will not equal the objective function. The objective of least-cost dispatch is not to minimize the total electricity bill paid by consumers, but to minimize production costs.

Two interesting phenomena worth noting are observed with locational marginal pricing. First is that, in cases of transmission congestion, the system λ will not equal the marginal cost of a single marginal unit, but will be a function of the costs of all marginal units (determined by system topology, and equivalently shift factors). Secondly, occasionally the LMP at a location will be higher than the highest price offer submitted by any unit²⁴.

3.3.4. Agents

Agents are arguably the most important component of this market model. Conceptually, an agent is an autonomous software agent with an interface for interacting with the outside world, in this case the modelled power system and market clearing mechanism. In our context, each agent represents a power generating company with a portfolio of 5 plants. The agent interacts with the market in two ways: by taking action and by receiving reward. It would be a stretch to assign volition or awareness to the agents in our context, but through the action-selection mechanism and the value-function-update mechanism, the agents, over iterations of the market, are designed to behave in ways that increase their profits.

Each agent controls 5 plants, each situated (possibly independently of each other) at one of the 5 nodes. The default capacity of every plant is 500 MW, while the costs of the plants, in increasing cost order are \$0, \$40, \$80, \$150 and \$600/MW, respectively. These parameters (location, capacity and cost) are modifiable by the modeller at the beginning of each simulation; a simulation is set of a predetermined number of rounds).

²⁴ This is well-known in the power industry and results from the fact that the possibility exists that maintaining nodal balance when delivering an additional MW to a particular node will require adjusting dispatch elsewhere on the grid such that cheap generation in several locations must be replaced with more expensive generation; as the locational marginal price equals the change to the total objective function, this incremental energy can be very expensive when several MW (or more) of expensive generation displace cheap generation)

Agents have a set of 32 actions, consisting of all combinations of their 5 plants either in-service or out-of-service (2^5). In contrast to most ABM research, the agents here are competing on quantity rather than price. The price-offer for each unit is fixed and equal to the marginal cost (for purposes of calculating profit; if a \$600 unit clears and the LMP at the node is \$600, the profit is \$0). Because determining quantity offered is the decision of each agent, it is therefore possible that the agents collectively may offer insufficient capacity to meet load, or offer their portfolios in a way that forces an infeasible solution with respect to the thermal line limit constraints. When this happens, proxy generators clear to supply energy at an arbitrarily high price so that the market solves (the feasible region is never empty), while the agents receive \$0 profits for the round.²⁵

3.3.4.1. Agent Value Set

The value function, in the fashion of Q-Learning, is $Q(a): \mathbf{A} \rightarrow \mathbf{R}$. $Q(a)$ holds a scalar value for each action, which in the limit will approach the expected value of that action²⁶, where convergence occurs. In our model, $Q(a)$ for every a is set to a high initial positive value, to encourage exploration (this implies an optimistic estimate of expected values for actions, by which agents are purposefully “disappointed” with actions they try and will continue to take actions they haven’t tried until overall values fall into line with experience).

²⁵ Strictly speaking, the LP described in the preceding section omits this proxy-unit feature. This discrepancy is analogous to real-world system operations, in which an infeasible dispatch solution will be disregarded by system operators, and the dispatch will be rerun with new inputs (as in Eastern RTOs), or the violating constraint will be relaxed with a penalty (as in California).

²⁶ As mentioned previously, Bertsekas and Tsitsiklis 1996 and Sutton and Barto, 1998 have shown that an action’s value in a stationary task will converge in the limit to the true expected value to the agent of taking that action.

In the reinforcement learning literature, there are two basic ways in which agents use their action-values to select actions: soft-max or epsilon-greedy²⁷. In the first case, soft-max, the agent selects one combination of their portfolio of plants in or out of service (each combination corresponding to one of their 32 actions) by sampling from an action-probability distribution, derived from their action-values. The probability of selecting any action a is:

$$P(a) = \frac{e^{Q(a)/\tau}}{\sum_{b=1}^{32} e^{Q(b)/\tau}} \quad (\text{x})$$

where τ is a so-called “temperature” parameter (used in reinforcement learning).

$$\sum_{a=1}^{32} P(a) = \sum_{a=1}^{32} \frac{e^{Q(a)/\tau}}{\sum_{b=1}^{32} e^{Q(b)/\tau}} = 1$$

Obviously, regardless of τ , the probabilities sum to 1. In reinforcement learning applications, in order for the theoretical optimality to be achieved, τ is adjusted downward by a scaling factor less than 1 throughout the simulation to gradually shift action selection from exploration (high τ implies a flatter action-probability distribution) to exploitation (a low τ implies a peakier distribution).²⁸

The second way agents may pick actions is the epsilon-greedy method: Simpler than soft-max, agents select their highest-valued action with probability 0.9 and sample from a uniform distribution of their actions with probability 0.1. The method used here is epsilon-greedy. In addition to requiring less computation, performance was better (agents discovered more profitable stable equilibria more often) with epsilon-greedy.

²⁷ The method used in Somani and Tesfatsion, 2008, is a special case of soft-max, with $\tau = 1$.

²⁸ In the stationary task case, according to Sutton and Barto 1998, optimality is guaranteed when certain criteria are met, one of which is that the temperature must be reduced properly to ensure that the optimal action is always chosen when the Q-function asymptotically approaches the true expected values of the actions.

Because the joint reward function is deterministic, the task resembles a matrix game, though the number of possible combinations of the five agents' actions exceeds 33.5 million ($= 32^5$). After the market clears, the agents are returned the cleared results of their portfolio, which plants cleared and what prices were paid. The agents then calculate their reward as a function of the total profit (a scalar multiplier in our case profit / 1000000; this has no practical impact except to avoid a chance of a computational error when the soft-max action-selection method was used with very small τ ; we initially tried this, but eventually went with epsilon-greedy). Having just taken action a , and calculated a reward r from their profit, the individual agent's update algorithm is:

$$Q(a) \leftarrow (1 - \alpha)Q(a) + \alpha r \quad (\text{ix})$$

This update algorithm has been widely used in agent-based models of electricity markets, since it was popularized by Erev and Roth (1998), in their comparison of reinforcement learners with human learners in experimental games. As agent-based models incorporate additional features that have the effect of complicating the models and facing modelers with the potential for non-convergence, other reinforcement learning algorithms (such as Abdallah and Lesser's weighted-policy learning) should be explored, as convergence of a community of agents to one or several Nash-Equilibria has been demonstrated in certain contexts when convergence has failed with the Q-learning update algorithm above.

3.4. Implementing the Model in Java

The model described above was implemented as a stand-alone Java application built within the Eclipse Integrated Development Environment (IDE)²⁹. In addition to standard Java

²⁹ See <http://www.eclipse.org/> for more information.

packages, we also used an open-source mathematical programming package (LP Solve) for optimization.³⁰ The graphical user interface (GUI) is shown in Figure Figure 6.

The application has two primary functions. First, “Run Scenario Round” allows the modeler to clear a single round of the market with manual inputs (setting not only the overall parameters: loads, line limits, unit capacities, unit locations and marginal costs, but also agents’ bids). This feature is useful for constructing scenarios in preparation for a simulation run.

The second feature is to run a simulation using inputs. In this case, all the parameters except the “Agent Offer Schedules” are used in the simulation (the agent offers are determined by the agents, as described above; changing these values in the interface does not alter the ability of each agent to offer their full portfolio).

In the present model version, the system topology and the numbers of agents and units (also called plants or generators), are hard-coded. Adding the flexibility of inputting these parameters was considered but was discarded because the marginal benefit to the research project was deemed less than the added cost of changing the model. Future steps in this line of research could include adding this feature. In fact, to make this research more applicable to real-world power systems, expanding the number of nodes, lines and agents will be critical. A key component will be to incorporate dynamic calculation of shift factors from an easy system-topology GUI design tool.

During simulation runs, data are saved to csv files in the default working directory (the target working directory can be changed through editing the Java code). Analysis was performed with the help of a second open-source program called R.³¹

³⁰ LP Solve was chosen because it is an established package of optimization tools with encouraging performance (many 100,000s of rounds are practical on a standard laptop or desktop personal computer). For more information, see <http://lpsolve.sourceforge.net/5.5/>.

CHAPTER 4

EXPLORING THE MODEL

Early research results presented in Miksis, 2006, showed that adaptive agents in a stationary electric power market model task successfully discovered optimal behavior. The challenges to achieving convergence in a more complex and multi-agent reinforcement learning application are significant, and as others have shown, convergence is not guaranteed, even in relatively simple matrix games (Abdallah and Lesser, 2008).

Although many of the most interesting aspects of the market model are seen through the simulation feature, the single round mode of the model supports intuition development as well as sensitivity analysis after simulation results are gathered.

Figure 6 shows the market model interface that is used for both modes. There are 8 windows for simulation inputs and one control panel. Clockwise from top left are: 1) the control panel, 2) the offer schedules for each of the 25 power plants (used only for the single-round mode), 3) the node locations of each of the 25 plants, 4) the thermal limits on each of the 6 transmission lines, 5) the nodal (demand) loads, 6) the marginal costs for each of the 25 plants, 7) the capacities of each plant (defaulted to 500) and finally, 8) inputs controlling the simulation.

The simulation inputs are the number of rounds to run, the learning rate for the agents, and the frequency with which the model should output market results to the set of csv files and the frequency with which information should be printed to the Java console (helpful for debugging). Using the single-round mode, it is possible to develop an intuitive understanding of how the model clears the market.

³¹ See <http://www.r-project.org/>.

The Market Model Interface consists of several panels for configuring simulation parameters:

- Control Panel:** Contains buttons for "Create Market", "Run Market", "Run Scenario Round", and "Run Many Rounds".
- Inputs Panel:** Contains input fields for:
 - Rounds: 10000
 - Learning Parameter: 0.05
 - Output Print Frequency: 1000
 - Console Print Frequency: 1000
- Agent Offer Schedules:** A 5x5 grid of values, all set to 1.
- Plant Locations:** A 5x5 grid of values: 4, 3, 2, 1, 0.
- Agent Capacities:** A 5x5 grid of values, all set to 500.
- Agent Marginal Costs:** A 5x5 grid of values: 0, 40, 80, 150, 600.
- Line Limits:** A vertical list of five input fields with values: 1500, 1500, 2500, 1000, 500.
- Nodal Loads:** A vertical list of five input fields with values: 3500, 2500, 1000, 1000, 500.

Figure 6. Market Model Interface

CHAPTER 5

USING THE MODEL

As described in the introduction, this research was motivated by 3 fundamental questions:

1. Can a community of adaptive agents competing in a quantity-based market model achieve equilibrium under various initial conditions?
2. Can players with plants at multiple locations (on either side of a transmission constraint) discover withholding policies that cause congestion, essentially raising prices at the import-constrained node?
3. Does the market as a whole supply surplus capacity above load or is the average excess capacity margin in equilibrium close to zero?

In order to answer these questions, we first developed a set of 3 scenarios with various distributions of plants across the system and recorded how the market cleared an “all-in” situation, in which no units are withheld. For each scenario, we then ran several simulations with the agents competing and learning. These scenarios well-represent two issues of concern in modern competitive power systems: The exercise of system and local market power. The first issue relates to firms’ ability to unilaterally or collectively withhold some capacity such that prices are maintained above a level consistent with a competitive market. Several system operators in the U.S. implement market power tests at the system level which allows mitigation actions to be taken if one or a group of firms are pivotal (demand cannot be met without at least some of their energy).

The second issue, local market power, is the ability to affect prices in a local area caused by network transmission constraints. System operators also have mechanisms in place to test for and mitigate market power exercised at the local level. The agents in our model are motivated to find profit-maximizing strategies by any means possible. We demonstrate that when highest profits are achieved by withholding capacity and raising the system clearing price to \$600, the agents converge to this point. In the case where the most profitable joint-strategy is to strategically withhold particular units at particular locations to take advantage of network constraints, the agents successfully find this as well. In order for agents to converge to a joint-strategy that takes advantage of network constraints, they must have both the means to cause congestion and the incentive to benefit from this action. If either element is missing, strategic congestion will not result. Scenario 1 has both elements, and so we see higher prices result. Scenarios 2 and 3 are structured such that overall capacity is withheld to raise system price, but there is no congestion.

The data appendix contains market clearing output for all of these scenarios in both the all-bid-in case and the strategic competition case. Scenarios are generically defined by thermal line limits, nodal loads, plant capacities, plant marginal costs and plant-to-node assignments. All scenarios use the line thermal limits and nodal loads as shown in Table 3, as well as the default capacity and marginal costs parameters described earlier.

These values were chosen to represent a power grid with defined import-constrained and export-constrained regions: Those with higher nodal demands represent load centers, such as large metropolitan areas (Boston is a good example in the New England power system), which generally also have higher generation costs and limited import capability, relative to demand; Those nodes with lower loads represent more rural regions where generation costs are lower and

the limitation is transmission capacity sufficient to export the energy that can potentially be generated at that node (Maine in the New England power system is an example).

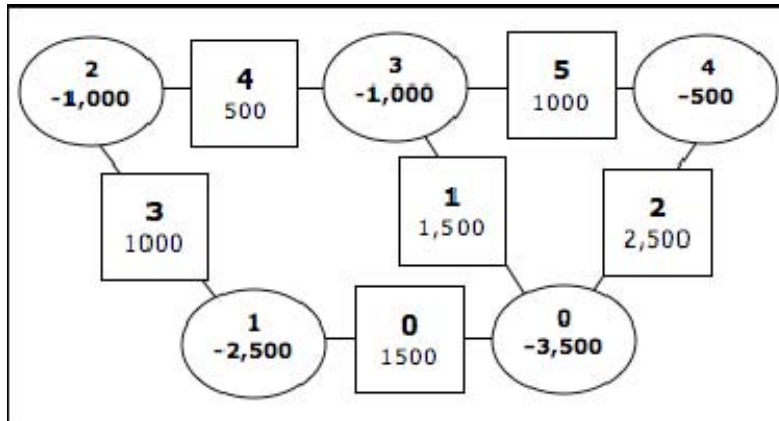


Figure 7. System Diagram Showing Line Limits and Loads

The scenarios explored below differ in the distribution of plants to nodes. While the potential avenues of experimentation with the model developed here are numerous, only a subset was chosen for practical reasons. These were chosen for their ability to showcase strategic behavior.

5.1. Scenario 1

In the first scenario, the plants are distributed such that each agent has one plant at each node, and the most expensive units are located in the higher-load areas. All agents have identical distributions of their plants among the nodes, with all \$600/MW plants at Node 0, \$150/MW plants at Node 1, etc. Appendix B Table 4 to Table 7 show market-cleared results in the case when all units are bid in (no units are out of service). The baseline least-cost dispatch in this case has congestion, because Line 3 (Node 1 to Node 2) is congested. Consequently, as the data tables show, Agent 0's Plant 1 is marginal and partially dispatched at the reference node at \$80,

while Agent 3's Plant 3 is marginal and partially dispatched at Node 1 at \$150. Prices at the other nodes are determined as explained above, by the system energy component ($\$80 =$ marginal generation unit cost at the reference node, Node 2) + shift factors * the shadow price on the binding transmission line (Line 3).

Scenario 1 was then simulated with the agents competing to meet load. This scenario attempts to answer Questions 1 and 2, whether equilibrium is reached, and whether agents strategically bid to take advantage of transmission constraints. Simulations show that the agents successfully achieve one of 20 (symmetrical) equilibrium joint policies in Scenario 1, involving a permutation of the offer schedules shown in Table 8.

Because the agent's plant distributions are identical, any of the permutations are equally likely (which permutation is achieved depends on the randomly-determined path of experiences and learning of the agents). As the table shows, only 4 of the 25 units are withheld, but profits are substantial. Figure 8 shows time series of a single 100,000 round simulation (the series represent averages of 100 rounds, so as to eliminate noise), in which rounds 45,000 and higher showcase the equilibrium profit-maximizing case shown in Table 8.




Figure 8. Scenario 1, Nodal LMPs for a Single Simulation Run

In the all-in market run, the highest LMP is \$150.00 at Node 1, one of the two nodes with relatively high load, and relatively expensive generation. Table 10 shows the prices resulting from [any permutation of] the offer schedule shown above in Table 8.³²

The highest nodal price is at the reference node (note that marginal congestion components can be negative, even though congestion itself always increases total costs), at \$1,955.79. A detailed derivation of this phenomenon is explained in Appendix A.

³² Table 9 to Table 12 show market output data for the one of the symmetrical permutations of the agents' equilibrium joint strategies.

The agents converge to a permutation of this equilibrium joint strategy because it is the strategy that maximizes total agent profit and it is a Nash-Equilibrium. No agent can unilaterally increase its profits by deviating (while three agents are making less profit than the best-off and second-best-off agents, they make even less if they break the stable joint strategy; which agents end up in which role is random). It should be noted that convergence was not achieved within 100,000 rounds in every simulation.

Figure 9 shows price results from 3 simulation runs of the model with the parameter set we have called “Scenario 1”. Clearly, after about 15,000 to 20,000 rounds, the agent’s have begun to assign significantly different values to their actions, demonstrated by the reduction in noise in the prices. It is clear that there are generally two regimes the agents converge to, and a regime switch occurs at some point in each simulation.




Figure 9. Scenario 1, Load-Weighted System LMP for Three Simulation Runs

It is notable that convergence occurs at different points in the simulation in each run. In some other runs of the model (not shown), we observed that convergence did not occur at all in 100,000 rounds. This indicates that there may be a non-zero probability of regime switching to the higher pricing and profit state never occurring.

By looking at the bid behavior over time, we can gain insights as well into whether the agents tend to offer capacity close to load levels (excess capacity ≈ 0), near full capacity (12,500) or somewhere in between. We see in Figure 10 that dynamic adjustments in market supply occur throughout the simulation, even while overall prices remain high. The series shows both average aggregate supply over all rounds and only those which cleared competitively. This

highlights how, after the initial learning phase, the agents proceed to rarely bid such that the market doesn't clear (the case when a proxy generator is dispatched and revenues are \$0 for all agents). Note that the chance of submitting a random offer (as opposed to the action with the highest expected value) remains 10%, and so the "Approved Rounds" series of aggregate MW supply offers is below 10,500 (the result of optimal behavior, with 4 units out of service of 500 MW each).

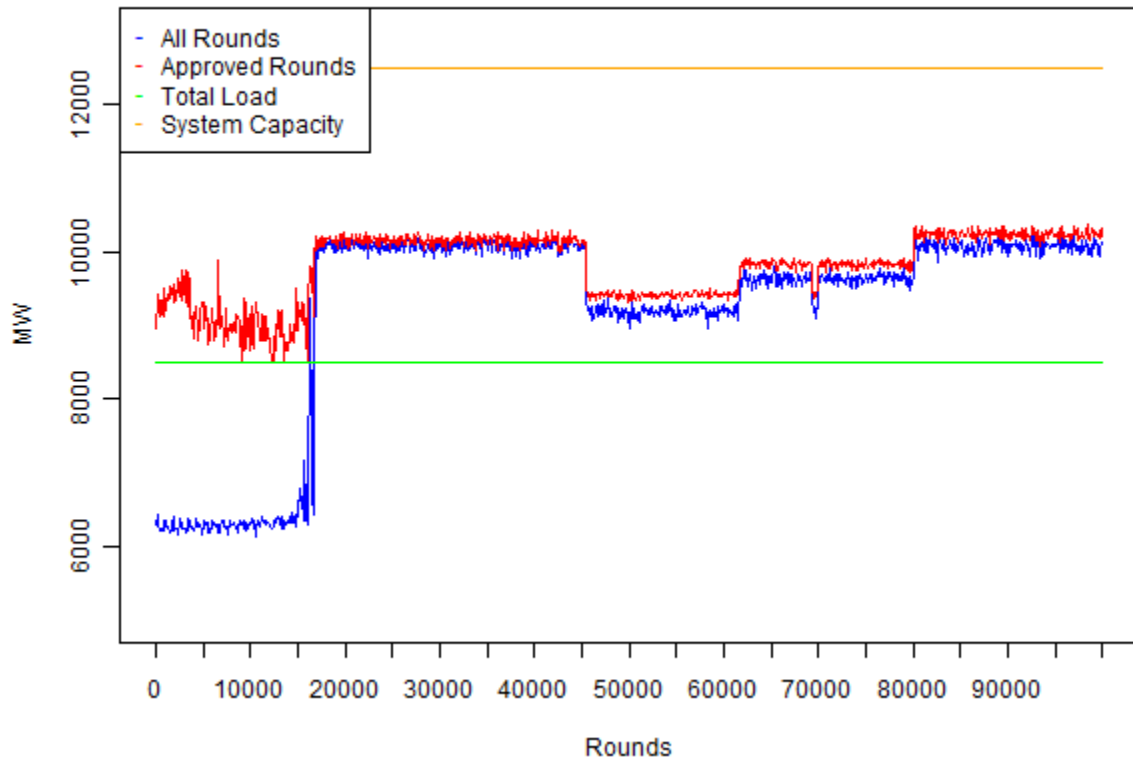


Figure 10. Scenario 1, Aggregate Supply Offered

Figure 11 shows the individual agents' profits over time, demonstrating that even after overall high prices have been achieved there remains some jostling among the agents for their piece of the profits.




Figure 11. Scenario 1, Individual Agents' Profits over the Simulation Run

5.2. Scenario 2

In the next scenario we explored distributed the agents' plants non-identically among the nodes, in which all of each agent's plants were at a single node (Agent 0 at Node 0, etc.). The results of the single-round all-bid-in market clearing are shown in Table 13 to Table 16. In this scenario, the all-bid-in least-cost dispatch does not have any congestion, and prices at all nodes

are equal to the energy component (also shadow price of the energy balance equation, or the marginal price of the marginal unit) of \$150. Figure 12 shows the nodal prices over the simulation run of Scenario 2. The competitive equilibrium market results can be found in the appendix (Appendix B, Table 17 to Table 21).

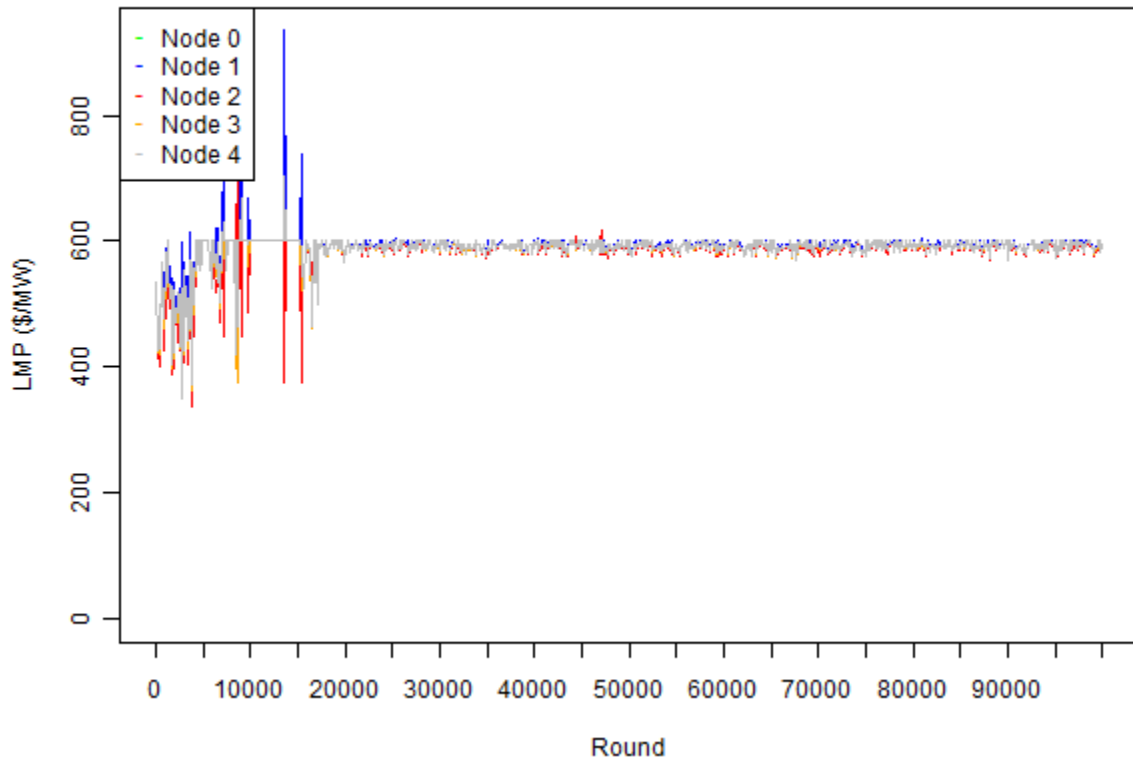


Figure 12. Scenario 2, Nodal LMPs for a Single Simulation Run

We do not see the same pattern of congestion causing price separation in the equilibrium solution for Scenario 2. However, due to selective withholding, the agents do maintain a price point at \$600/MW, compared with \$150/MW in the all-bid-in case. Figure 13, below, shows aggregate supply, graphed with total load and system capacity. We see that, in answer to

Question 3, the agents withhold enough capacity to hold prices above the baseline least-cost all-bid-in case.




Figure 13. Scenario 2, Aggregate Supply Offered

5.3. Scenario 3

The third scenario we examined was the inverse of Scenario 1. In this case, each agent had a single plant at each node, except the distributions were in the reverse order (\$600/MW at Node 4, \$150/MW at Node 3, etc.). The key result here was that, while all-bid-in least-cost dispatch resulted in congestion, the equilibrium competitive solution did not (See Table 22 and Table 27). As in Scenario 2, the agents successfully withheld capacity to maintain an average

system price of \$600/MW, but this was not more profitable than a solution involving strategically-caused congestion.




Figure 14. Scenario 3, Nodal LMPs for a Single Simulation Run

Similarly to the simulation run of Scenarios 1 and 2, we see a switch in aggregate supply offered somewhere between rounds 85,000 and 90,000, though prices do not change.




Figure 15. Scenario 3, Aggregate Supply Offered

5.4. Using The Model – Conclusion

In all 3 of the scenarios we ran, the 5 agents established baseline price outcomes significantly higher than the “all-bid-in” competitive solution. In one scenario, the agents utilized the system topology and distribution of generating units among the nodes to maintain a load-weighted system average price significantly above even the highest price offer: Absent the continued exploration caused by the epsilon value, which caused deviations from the learned greedy behavior, every round after learning converged in Scenario 1 would have had prices as shown in Table 10 (a weighted system price of \$974.49, 62% higher than the highest unit costs).

CHAPTER 6

CONCLUSION AND WHAT'S NEXT

In the last two decades, many examples of agent-based models of competitive electric power markets have been developed, varying along several dimensions, including sophistication of the agents, complexity of the underlying transmission system and incorporation of features unique to modern electric power systems. As with many new ventures in a young field, there remain significant trade-offs that must be made, explicitly or not, between complexity/realism and practicability. Simple agents in a complex environment may not behave as agents in an analogous real-world situation would, which can lead to results that, while interesting, do not carry the gravity of results with optimality proofs, or those with strong evidence of convergence to equilibria that can be favorably benchmarked against real-world phenomena. This carries the risk that interesting conclusions are not as compelling as they could be. On the other hand, sophisticated agents with learning capabilities may not fully learn a task with multiple dimensions in terms of interacting products (gas and electricity, electricity and ancillary services, etc.), overlapping and nested time periods (multi-year contracting, day-ahead and spot balancing), or a large number of assets with granular action sets, in which case convergence may not occur.

Several researchers in the field have commented that agent-based models suffer from insufficient robustness necessary for the ABM approach to gain wider acceptance, particularly among economists and regulators. Part of the motivation for this research was to incrementally build a model with enough realism to demonstrate phenomena that are in fact of great concern to stakeholders in the electric power industry (market power in transmission systems with

congestion). We have succeeded in doing this, demonstrating, although without rigorous parallel analytical solutions (like Waltman and Kaymak, 2008, have), that profit-seeking agents both 1) quickly find and favor feasible solutions and 2) often converge in a relatively small number of rounds to profit-maximizing equilibria. In all 3 of the scenarios we ran, the 5 agents established baseline price outcomes significantly higher than the “all-bid-in” competitive solution. In one scenario, the agents utilized the system topology and distribution of generating units among the nodes to maintain a load-weighted system average price significantly above even the highest price offer (a load-weighted system price of \$974.49, 62% higher than the most expensive units’ costs).

In the work presented in this thesis, we have attempted to answer the three questions posed in the introduction:

1. Can a community of adaptive agents competing in a quantity-based market model achieve equilibrium under various initial conditions?
2. Can players with plants at multiple locations (on either side of a transmission constraint) discover withholding policies that cause congestion, essentially raising prices at the import-constrained node?
3. Does the market as a whole supply surplus capacity above load or is the average excess capacity margin in equilibrium close to zero?

We can answer the first two affirmatively, under the conditions presented in each of the scenarios (with the caveat that the 2nd and 3rd scenarios did not have the underlying generating plant distribution among the nodes in which an equilibrium solution with congestion was the most profitable outcome). The third question is actually not very significant in the scenarios

presented in this work, as the agents clearly had little incentive to withhold capacity down to the level of load once the system energy price reached the maximum attainable in the uncongested case. For the congested case, clearly aggregate supply is not the deciding factor in profitability, but it is the interplay between localized capacity and transmission constraints.

Future steps in this line of research must include consideration of how to incorporate environments with increased complexity while appreciating the limitations of learning agents. Additional surveys of contemporary multi-agent reinforcement learning literature may lead toward insights on how to move beyond the limitations presently identified of more complexity and less tractability vs. less complexity and more tractability. Neural networks and other general function approximation techniques may support sufficient learning speeds in high-dimensional tasks where generalization is necessary.

Immediate steps to proceed from the current state of the research presented here will include enhancing the model to support dynamic system topology creation through a GUI or input files in order to explore more realistic transmission congestion scenarios, and more robust experimentation to explore the distribution of rounds-to-convergence among larger numbers of identical scenario runs with different pseudo-random number seeds. Introducing demand-side bidding and a multi-settlement system may be possible to the extent that the costs in terms of added joint-action space are made explicit and convergence remains achievable; here more than anywhere else, what we don't yet know of contemporary multi-agent systems research may be a silver bullet of sorts in surmounting the trade-off highlighted above.

The potential for research in the direction begun in this thesis is both academically engaging and practical for applications to real-world concerns. Flexible, scalable, customizable models of market interactions support ultimate what-if scenario analysis. Combined with an

underlying transmission model and generation portfolio with sufficient realism, the models have the potential to be run parallel to transmission planning studies, structural market power estimates, and market rule-making to stress-test scenarios before real-world implementation. Regulators and other stakeholders have repeatedly expressed concern about the potential exercise of market power in competitive electricity markets, particularly in regions where transmission bottlenecks are prone to creating extreme market power where there are only a few independent suppliers available to meet demand in extreme circumstances. However, rules written to test for, catch and punish anticompetitive behavior, such as strategic withholding, are complex and imperfect. The Federal Energy Regulatory Commission has determined to ensure just and reasonable electric rates through design and implementation of competitive generation markets and open-access transmission systems coupled with careful regulations tailored to the uniqueness of electric power markets. Using all available research tools to examine market outcomes where agents seek to exploit profitable market weaknesses without regard to regulatory repercussions can offer an added layer of protection to market monitoring and mitigation procedures of system operators, particularly in light of the complex pricing effects strategic behavior can have on an interconnected power grid.

APPENDIX A

EXPOSITION OF NODAL PRICES

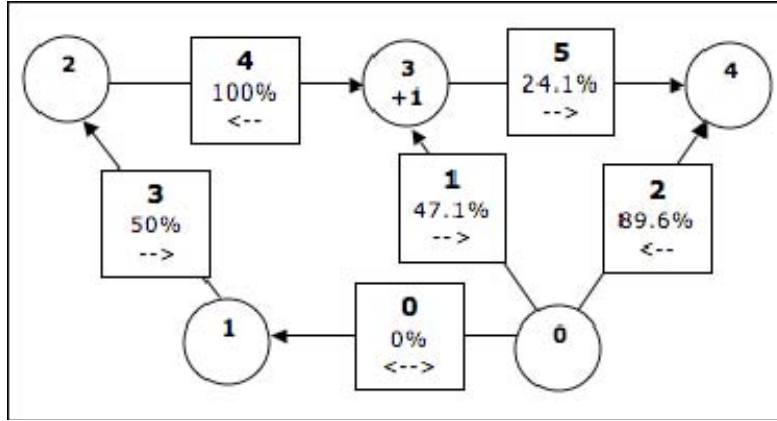


Figure 16. Scenario 1 Equilibrium Joint Strategy, Power Line Flows % of Capacity

1. Line 4, running from Node 2 to Node 3, is congested in the direction of the reference Node (2). Consequently, in order to deliver a single MW to Node 2, we need to find a combination of redispatches that will not increase flow on this line.
2. Note that the GSF of every non-reference node to Line 4 is negative, meaning unilaterally increasing generation anywhere will exacerbate the problem (violate the constraint).

Table 1. Generator Shift Factors (Nodes in Rows, Lines in Columns)

	0	1	2	3	4	5
0	0.54	0.25	0.21	0.54	-0.46	-0.21
1	-0.13	0.07	0.06	0.87	-0.13	-0.06
3	0.35	-0.19	-0.16	0.35	-0.65	0.16
4	0.51	0.17	-0.68	0.51	-0.49	-0.32

3. To solve this problem, we need to increase generation by a net of 1 MW at two nodes by increasing output at one of them by 1 more than we decrease output at another.³³ However, there is another limitation. The node at which we increase generation must have a smaller (absolute) shift factor with Line 4 than the one where we decrease generation.
4. The candidates for increasing generation are a plant at Node 3 or a plant at Node 0 (notice that those plants in service at the other nodes have all been dispatched to their maximum capacity).
5. Because the GSF for Node 0 and Line 4 (-0.46) is absolutely smaller than the GSF for Node 3 and Line 4 (-0.65), we will increase output at Node 0 and decrease it at Node 4.
6. Consider that meeting the incremental MW at Node 2 from Node 0 will cause Line 4 to go 0.46 MW over its limit. Therefore, we need to further increase output at Node 1 and balance it with a decrease in output at Node 3. But by how much?
7. The effect of replacing a single MW at Node 3 with one from Node 0 has the net effect of decreasing flow in the offending direction on Line 4 by 0.19 MW ($-0.46 - (-0.65) = 0.19$).

³³ PowerWorld is a software package that is extremely useful for generating intuition of power flows. See <http://www.powerworld.com/> for more information.

8. Additionally, the marginal change in the objective function of this single MW substitution is \$560 (\$600 cost at Node 0 - \$40 saved at Node 3). Therefore, we need to incur $\$560 / 0.19$ to decrease flow on Line 4 by a full 1 MW.
9. However, we don't need a full MW, we only need 0.46 MW. $0.46 * \$560 / 0.19 = \1355.79 . This equals the price difference between the source of that incremental MW (Node 0) and the reference node (Node 2)!³⁴

³⁴ It is important to remember that the choice of a reference node will not change LMPs, but it will change congestion components and loss components in AC systems. Therefore, in an energy-only power system, the reference node is arbitrary, but when financial settlements are made based on components of LMP (in AC systems, congestion and losses), the choice is not longer arbitrary. See Litvinov, et al, 2004.

APPENDIX B

DATA

Table 2. Transmission System Data

		Node to Reference	0→2	1→2	3→2	4→2
Line	Resistance	Line Flow % (Generator Shift Factors)				
0	0.0281		54%	-13%	35%	51%
1	0.0304		25%	7%	-19%	17%
2	0.0064		21%	6%	-16%	-68%
3	0.0108		54%	87%	35%	51%
4	0.0297		-46%	-13%	-65%	-49%
5	0.0297		-21%	-6%	16%	-32%

Table 3. Defaults Scenarios' Power System Parameters

Line	Limit		Node	Load
0	1,500		0	3,500
1	1,500		1	2,500
2	2,500		2	1,000
3	1,000		3	1,000
4	500		4	500
5	1,000			

Table 4. Scenario 1 Cleared All-Bid-In, Portfolio Results

Agent	Plant 0			Plant 1			Plant 2			Plant 3			Plant 4		
	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M
0	500	\$121.03	0	500	\$108.16	0	252.9	\$80	1	500	\$150	1	0	\$123.45	0
1	500	\$121.03	0	500	\$108.16	0	500	\$80	1	500	\$150	0	0	\$123.45	0
2	500	\$121.03	0	500	\$108.16	0	0	\$80	0	500	\$150	1	0	\$123.45	0
3	500	\$121.03	0	500	\$108.16	0	500	\$80	1	247.1	\$150	1	0	\$123.45	0
4	500	\$121.03	0	500	\$108.16	0	500	\$80	1	0	\$150	0	0	\$123.45	0

Table 5. Scenario 1 Cleared All-Bid-In, Nodal Prices

LMPs			
Node	Energy	Congestion	Price
0	\$80	\$43.45	\$123.45
1	\$80	\$70	\$150
2	\$80	\$0	\$80
3	\$80	\$28.16	\$108.16
4	\$80	\$41.03	\$121.03

Table 6. Scenario 1 Cleared All-Bid-In, Transmission Line Flows

Line Flows		
Line	MW Flow	Capacity (+/-)
0	-247.1	1500
1	-872.7	1500
2	-2380.2	2500
3	-1000	1000
4	-247.1	500
5	380.2	1000

Table 7. Scenario 1 Cleared All-Bid-In, Nodal MWs

Nodal Net Loads				
Node	Total Generation	Competitive Generation	Load	Net Import
0	0	0	3500	-3500
1	1747.1	1747.1	2500	-752.9
2	1752.9	1752.9	1000	752.9
3	2500	2500	1000	1500
4	2500	2500	500	2000

Table 8. Scenario 1, Cleared Equilibrium Joint Strategy, Agent Quantity Offers

Agent	Plant 0	Plant 1	Plant 2	Plant 3	Plant 4
0	500	500	0	500	500
1	500	500	0	500	500
2	500	500	500	0	500
3	500	500	0	500	500
4	500	500	500	500	500

Table 9. Scenario 1, Cleared Equilibrium Joint Strategy, Portfolio Results

Agent	Plant 0			Plant 1			Plant 2			Plant 3			Plant 4		
	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M
0	500	\$511.58	0	500	\$40	1	0	\$1,955.79	0	500	\$1,572.63	0	52.6	\$600	1
1	500	\$511.58	0	500	\$40	0	0	\$1,955.79	0	500	\$1,572.63	0	500	\$600	0
2	500	\$511.58	0	447.4	\$40	1	500	\$1,955.79	0	0	\$1,572.63	0	0	\$600	0
3	500	\$511.58	0	500	\$40	1	0	\$1,955.79	0	500	\$1,572.63	0	0	\$600	0
4	500	\$511.58	0	500	\$40	1	500	\$1,955.79	0	500	\$1,572.63	0	0	\$600	0

Table 10. Scenario 1 Cleared Equilibrium Joint Strategy, Nodal Prices

LMPs			
Node	Energy	Congestion	Price
0	\$1,955.79	(\$1,355.79)	\$600.00
1	\$1,955.79	(\$383.16)	\$1,572.63
2	\$1,955.79	\$0.00	\$1,955.79
3	\$1,955.79	(\$1,915.79)	\$40.00
4	\$1,955.79	(\$1,444.21)	\$511.58

Table 11. Scenario 1, Cleared Equilibrium Joint Strategy, Transmission Line Flows

Line Flows		
Line	MW Flow	Capacity (+/-)
0	0	1500
1	-706.8	1500
2	-2240.5	2500
3	-500	1000
4	-500	500
5	240.5	1000

Table 12. Scenario 1, Cleared Equilibrium Joint Strategy, Nodal MWs

Nodal Net Loads				
Node	Total Generation	Competitive Generation	Load	Net Import
0	552.6	552.6	3500	-2947.4
1	2000	2000	2500	-500
2	1000	1000	1000	0
3	2447.4	2447.4	1000	1447.4
4	2500	2500	500	2000

Table 13. Scenario 2, Cleared All-Bid-In, Portfolio Results

Agent	Plant 0			Plant 1			Plant 2			Plant 3			Plant 4		
	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M
0	500	\$150.00	0	500	\$150.00	0	500	\$150.00	0	0	\$150.00	0	0	\$150.00	0
1	500	\$150.00	0	500	\$150.00	0	500	\$150.00	0	11.5	\$150.00	1	0	\$150.00	0
2	500	\$150.00	0	500	\$150.00	0	500	\$150.00	0	488.5	\$150.00	1	0	\$150.00	0
3	500	\$150.00	0	500	\$150.00	0	500	\$150.00	0	0	\$150.00	0	0	\$150.00	0
4	500	\$150.00	0	500	\$150.00	0	500	\$150.00	0	500	\$150.00	0	0	\$150.00	0

Table 14. Scenario 2, Cleared All-Bid-In, Nodal Prices

Node	Energy	Congestion	Price
0	\$150.00	\$0.00	\$150.00
1	\$150.00	\$0.00	\$150.00
2	\$150.00	\$0.00	\$150.00
3	\$150.00	\$0.00	\$150.00
4	\$150.00	\$0.00	\$150.00

Table 15. Scenario 2, Cleared All-Bid-In, Transmission Line Flows

Line	MW Flow	Capacity (+/-)
0	-11.5	1500
1	-409.2	1500
2	-1579.3	2500
3	-1000	1000
4	-11.5	500
5	79.3	1000

Table 16. Scenario 2, Cleared All-Bid-In, Nodal MWs

Node	Total Generation	Competitive Generation	Load	Net Import
0	1500	1500	3500	-2000
1	1511.5	1511.5	2500	-988.5
2	1988.5	1988.5	1000	988.5
3	1500	1500	1000	500
4	2000	2000	500	1500

Table 17. Scenario 2, Cleared Equilibrium Joint Strategy, Agent Quantity Offers

Agent	Plant 0	Plant 1	Plant 2	Plant 3	Plant 4
0	500	500	500	0	500
1	500	500	500	0	500
2	500	500	500	0	500
3	500	500	500	0	500
4	500	500	500	500	0

Table 18. Scenario 2, Cleared Equilibrium Joint Strategy, Portfolio Results

Agent	Plant 0			Plant 1			Plant 2			Plant 3			Plant 4		
	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M
0	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0	0	\$600.00	0
1	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0	500	\$600.00	1
2	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0	0	\$600.00	0
3	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0	0	\$600.00	0
4	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0

Table 19. Scenario 2, Cleared Equilibrium Joint Strategy, Nodal Prices

Node	Energy	Congestion	Price
0	\$600.00	\$0.00	\$600.00
1	\$600.00	\$0.00	\$600.00
2	\$600.00	\$0.00	\$600.00
3	\$600.00	\$0.00	\$600.00
4	\$600.00	\$0.00	\$600.00

Table 20. Scenario 2, Cleared Equilibrium Joint Strategy, Transmission Line Flows

Line	MW Flow	Capacity (+/-)
0	-75	1500
1	-375	1500
2	-1550	2500
3	-575	1000
4	-75	500
5	50	1000

Table 21. Scenario 2, Cleared Equilibrium Joint Strategy, Nodal MWs

Node	Total Generation	Competitive Generation	Load	Net Import
0	1500	1500	3500	-2000
1	2000	2000	2500	-500
2	1500	1500	1000	500
3	1500	1500	1000	500
4	2000	2000	500	1500

Table 22. Scenario 3, Cleared All-Bid-In, Portfolio Results

Agent	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M
0	500	\$129.54	0	500	\$94.00	0	500	\$80.00	0	315.4	\$150.00	1	0	\$132.77	0
1	500	\$129.54	0	500	\$94.00	0	500	\$80.00	0	500	\$150.00	0	0	\$132.77	0
2	500	\$129.54	0	500	\$94.00	0	500	\$80.00	0	0	\$150.00	0	0	\$132.77	0
3	500	\$129.54	0	500	\$94.00	0	500	\$80.00	0	500	\$150.00	0	0	\$132.77	0
4	500	\$129.54	0	500	\$94.00	0	184.6	\$80.00	1	0	\$150.00	0	0	\$132.77	0

Table 23. Scenario 3, Cleared All-Bid-In, Nodal Prices

Node	Energy	Congestion	Price
0	\$80.00	\$49.54	\$129.54
1	\$80.00	\$14.00	\$94.00
2	\$80.00	\$0.00	\$80.00
3	\$80.00	\$70.00	\$150.00
4	\$80.00	\$52.77	\$132.77

Table 24. Scenario 3, Cleared All-Bid-In, Transmission Line Flows

Line	MW Flow	Capacity (+/-)
0	-684.6	1500
1	-394.9	1500
2	79.5	2500
3	-684.6	1000
4	500	500
5	420.5	1000

Table 25. Scenario 3, Cleared All-Bid-In, Nodal MWs

Node	Total Generation	Competitive Generation	Load	Net Import
0	2500	2500	3500	-1000
1	2500	2500	2500	0
2	2184.6	2184.6	1000	1184.6
3	1315.4	1315.4	1000	315.4
4	0	0	500	-500

Table 26. Scenario 3, Cleared Equilibrium Joint Strategy, Agent Quantity Offers

Agent	Plant 0	Plant 1	Plant 2	Plant 3	Plant 4
0	500	500	0	500	500
1	500	500	500	0	500
2	500	500	0	500	500
3	500	500	500	0	500
4	500	500	500	500	500

Table 27. Scenario 3, Cleared Equilibrium Joint Strategy, Portfolio Results

Agent	Plant 0			Plant 1			Plant 2			Plant 3			Plant 4		
	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M	MW	Price	M
0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0	500	\$600.00	0	0	\$600.00	0
1	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0	500	\$600.00	1
2	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0	500	\$600.00	0	0	\$600.00	0
3	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0	0	\$600.00	0
4	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	500	\$600.00	0	0	\$600.00	0

Table 28. Scenario 3, Cleared Equilibrium Joint Strategy, Nodal Prices

Node	Energy	Congestion	Price
0	\$600.00	\$0.00	\$600.00
1	\$600.00	\$0.00	\$600.00
2	\$600.00	\$0.00	\$600.00
3	\$600.00	\$0.00	\$600.00
4	\$600.00	\$0.00	\$600.00

Table 29. Scenario 3, Cleared Equilibrium Joint Strategy, Transmission Line Flows

Line	MW Flow	Capacity (+/-)
0	-365.0	1500
1	-345.0	1500
2	-290.0	2500
3	-365.0	1000
4	135	500
5	290.0	1000

Table 30. Scenario 3, Cleared Equilibrium Joint Strategy, Nodal MWs

Node	Total Generation	Competitive Generation	Load	Net Import
0	2,500	2,500	3,500	-1,000
1	2,500	2,500	2,500	0
2	1,500	1,500	1,000	500
3	1,500	1,500	1,000	500
4	500	500	500	0

BIBLIOGRAPHY

- Abdallah, S. and V. Lesser, (2008), "A Multiagent Reinforcement Learning Algorithm with Non-linear Dynamics," *Journal of Artificial Intelligence Research* 33(2008):521-549.
- Amin, M., (2002), "Restructuring the Electric Enterprise: Simulating the Evolution of the Electric Power Industry with Intelligent Adaptive Agents," In *Market Analysis and Resource Management*, eds. A. Faruqui and K. Eakin, March 2002.
- Axelrod, R. (1987). "The Evolution of Strategies in the Iterated Prisoner's Dilemma." In L. Davis, ed., *Genetic Algorithms and Simulated Annealing*, Research Notes in Artificial Intelligence. Pitman Publishing, London.
- Bagnall, A. J. and G. D. Smith (2005), "A Multi-Agent Model of the UK Market in Electricity Generation." *IEEE Transactions on Evolutionary Computation*, 9(5). October.
- Baldick, R., U. Helman, B. F. Hobbs and R. P. O'Neill, (2005), "Design of Efficient Generation Markets." In the *Proceedings of the IEEE*, 93(11):1998-2012. November.
- Bertsekas, D. P. and J. N. Tsitsiklis. (1996). *Neuro-Dynamic Programming*. Athena Scientific.
- Besser, J. G., J. G. Farr and S. F. Tierney, (2002), "The Political Economy of Long-Term Generation Adequacy: Why an ICAP Mechanism is Needed as Part of Standard Market Design." In *The Electricity Journal*, 15(7):53-62. August/September.
- Bidwell, M. and A. Henney. 2004. "Will the New Electricity Trading Arrangements Ensure Generation Adequacy?" *The Electricity Journal* 17(7):15-38. August/September.
- Borenstein, S. and S. P. Holland (2002). "Investment Efficiency in Competitive Electricity Markets With and Without Time-Varying Retail Prices." *The Center for the Study of Energy Markets Working Paper*, University of California Energy Institute, November.
- Bower and Bunn (1999), "A Model-Based Comparison of Pool and Bilateral Market Mechanisms for Electricity Trading" *Energy Markets Group*, London Business School.
- Bunn, D. W. and M. Martocchia, (2008), "Analyzing the Price-Effects of Vertical and Horizontal Market Power with Agent Based Simulation," *IEEE Power and Energy Society General Meeting – Conversion and Delivery of Electrical Energy in the 21st Century*, 2008.
- Bunn, D. W. and F. S. Oliveira, (2001), "Agent-Based Simulation – An Application to the New Electricity Trading Arrangements of England and Wales," *IEEE Transactions on Evolutionary Computation*, 5(5):493-503.
- Caramanis, M. C. (1982), "Investment Decisions and Long-Term Planning Under Electricity Spot Pricing." *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-101(12):4640-4648.

- Caramanis, M. C., R. E. Bohn and F. C. Schweppe (1982), "Optimal Spot Pricing: Practice and Theory." IEEE Transactions on Power Apparatus and Systems, Vol. PAS-101(9):3234-3245.
- Chao, H and R. Wilson. 2004. "Resource Adequacy and Market Power Mitigation via Option Contracts" Electric Power Research Institute.
- Conzelmann, G., M. J. North, G. A. Boyd, R. R. Cirillo, V. Koritarov, C. M. Macal, P. R. Thimmapuram and T. D. Veselka, (2004), "Simulating strategic market behavior using an agent-based modeling approach. In Proceedings, 6th IAEE European Conference, Zurich, Switzerland.
- Cramton P., H-P. Chao, R. Wilson (2005), "A Review of the Proposed Reserve Markets in New England." Market Design, Inc. White Paper (<http://www.cramton.umd.edu/auction-papers.htm>). January.
- Cramton, P. and S. Stoft (2005), "A Capacity Market That Makes Sense." The Electricity Journal, 18(7):43-54. August/September.
- Creti, A. and N. Fabra (2003), "Capacity Markets for Electricity." Industrial Organization, EconWPA (<http://ideas.repec.org/p/wpa/wuwpio/0309002.html>). November.
- Czernohous, C., W. Fichtner, D. J. Veit and C. Weinhardt, (2003), "Management Decision Support Using Long-term Market Simulation," Information Systems and e-Business Management, accessed online at (<http://www.im.uni-karlsruhe.de/Default.aspx?PageId=355&lang=en>) on May 24, 2009.
- De Vries, L. and P. Heijnen, (2008), "The Impact of Electricity Market Design Upon Investment Under Uncertainty: The Effectiveness of Capacity Mechanism," Utilities Policy, 16(2008):215-227.
- De Vries, L. J. and Hakvoort R. A. (2004), "The Question of Generation Adequacy in Liberalised Electricity Markets." INDES Working Paper No. 5, Center for European Policy Studies. March.
- Doorman, G. L. (2003), "Capacity Subscription and Security of Supply in Deregulated Electricity Markets." Presented at the Research Symposium on European Electricity Markets, The Hague. September.
- Duffy, J., (2006), "Agent-Based Models and Human Subject Experiments," in Handbook of Computational Economics, Volume 2, Eds. L. Tesfatsion and K. L. Judd, 2006, Elsevier.
- Entriiken, R. and S. Wan, (2005), "Agent-Based Simulation of an Automatic Mitigation Procedure," Proceedings of the 38th Hawaii International Conference on System Sciences, 2005, IEEE.

- Epstein, J.M. and R. Axtell (1996), *Growing Artificial Societies: Social Science from the Ground Up*. The Brookings Institution/MIT Press.
- Erev, I. and A. E. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *The American Economic Review*, 88(4):848-881. September.
- Fagiolo, G., C. Birchenhall and P. Windrum, (2007), "Empirical Validation in Agent-based Models: Introduction to the Special Issue," *Computational Economics*, 2007(30):189-194.
- Farr, J. G. and F. A. Felder (2005), "Competitive Electricity Markets and System Reliability: The Case for New England's Proposed Locational Installed Capacity Market." *The Electricity Journal*, 18(8):22-33, October.
- Federal Energy Regulatory Commission (referred in-text as FERC Order 888), (1996), "Promoting Wholesale Competition Through Open Access Non-discriminatory Transmission Services by Public Utilities; Recovery of Stranded Costs by Public Utilities and Transmitting Utilities," Final Order No. 888, (Docket Nos. RM95-8-000 and RM94-7-001), (accessed at <http://www.ferc.gov/legal/maj-ord-reg/land-docs/order888.asp>, on June 2, 2009).
- Federal Energy Regulatory Commission (referred in-text as FERC Order 889), (1996), "Open Access Same-Time Information System (formerly Real-Time Information Networks) and Standards of Conduct," Final Order No. 888, (Docket Nos. RM95-8-000 and RM94-7-001), (accessed at <http://www.ferc.gov/legal/maj-ord-reg/land-docs/order888.asp>, on June 2, 2009).
- Ford, A., (2002), "Boom and Bust in Power Plant Construction: Lessons from the California Electricity Crisis," *Journal of Industry, Competition and Trade*, 2(1/2):59-74.
- Fraser, H. (2001). "The Importance of an Active Demand Side in the Electricity Industry." *The Electricity Journal*. November.
- Hirst, E. and S. Hadley (1999), "Maintaining Generation Adequacy in a Restructuring U.S. Electricity Industry." ORNL/CON-472, Oak Ridge National Laboratory, Oak Ridge, TN.
- Hobbs, B. F. and U. Helman, (2004), "Complementarity-Based Equilibrium Modeling for Electric Power Markets," in *Modeling Prices in Competitive Electricity Markets*, ed. D. W. Bunn, John Wiley & Sons, Ltd.
- Hunt, S., (2002), *Making Competition Work in Electricity*, Wiley, March 22, 2002.
- Jaffe, A. B. and F. A. Felder (1996), "Should Electricity Markets Have a Capacity Requirement? If So, How Should it be Priced?" *The Electricity Journal*, 9(10):52-60. December.

- Joskow, 2003. "Comments of Professor Paul Joskow before the Federal Energy Regulatory Commission." FERC Testimony, docket no. RM01-12-000, FERC. January.
- Kahn, A. E., P. C. Cramton, R. H. Porter and R. D. Tabors, 2001, "Uniform Pricing or Pay-as-Bid Pricing: A Dilemma for California and Beyond," *The Electricity Journal*, 14(6):70-79. July.
- Kaelbling, L. P., M. L. Littman and A. W. Moore, (1996), "Reinforcement Learning: A Survey," *Journal of Artificial Intelligence Research*, 4:237-285.
- Klemperer, P. D. and M. A. Meyer, (1989), "Supply Function Equilibria in Oligopoly under Uncertainty," *Econometrica*, 57(6):1243-1277.
- Koritarov, V. S, (2004), "Real-World Market Representation with Agents: Modeling the Electricity Market as a Complex Adaptive System with an Agent-Based Approach," *IEEE Power and Energy Magazine*, July/August, 2004.
- Leombruni, R. and M. Richiardi, (2005), "Why are Economists Sceptical about Agent-Based Simulations?" *Physica A*, 355(2005):103-109.
- Li, F. and R. Bo, (2007), "DCOPF-Based LMP Simulation: Algorithm, Comparison With ACOPF, and Sensitivity," *IEEE Transactions on Power Systems*, 22(4):1475-1485.
- Littman, M. L., (1994), "Markov Games as a Framework for Multi-agent Reinforcement Learning," In *Proceedings of the Eleventh International Conference on Machine Learning (1994)*, pp. 157-163.
- Litvinov, E., T. Zheng, G. Rosenwald and P. Shamsollahi, 2004, "Marginal Loss Modeling in LMP Calculation," *IEEE Transactions on Power Systems*, 19:880-888. May.
- Marks, R. E., (2007), "Validating Simulation Models: A General Framework and Four Applied Examples," *Computational Economics*, 30(2007):265-290.
- Midgley, D., R. Marks and D. Kunchamwar, (2007), "Building and Assurance of Agent-Based Models: An Example and Challenge to the Field," *Journal of Business Research*, 60(2007):884-893.
- Miksis, N., (2006), "Program on Technology Innovation: Optimal Bidding in Power Markets: Q-Learning for Adaptive Agents," EPRI, Palo Alto, CA and EDF R&D, 1016229 (<http://www.epriweb.com/public/000000000001016229.pdf>).
- Nanduri, V. and T. K. Das, (2009), "A Survey of Critical Research Areas in the Energy Segment of Restructured Electric Power Markets," *Electrical Power and Energy Systems*, 31(2009):181-191.

- Nicolaisen, Petrov and Tesfatsion, (2001), "Market Power and Efficiency in a Computational Electricity Market with Discriminatory Double-Auction Pricing." IEEE Transactions on Evolutionary Computation, 5(5). October.
- Oren, S. (2003), "Ensuring Generation Adequacy in Competitive Electricity Markets." UCEI Energy Policy and Economics Working Paper 007.
- Overbye, T. J. and D. A. Weigmann, 2005, "Reducing the Risk of Major Blackouts Through Improved Power System Visualization," Submitted to 15th Power System Computation Conference (PSCC), Liege, Belgium. August.
- Rochlin, C. (2004). "Resource Adequacy Requirement, Reserve Margin and the Public Goods Argument." The Electricity Journal. April.
- Rahimiyan, M. and H. R. Mashhadi, (2008), "Supplier's Optimal Bidding Strategy in Electricity Pay-As-Bid Auction: Comparison of the Q-Learning and a Model-Based Approach," Electric Power Systems Research, 78(2008):165-175.
- Schweppe, F. C. (1978), "Power Systems '2000': Hierarchical Control Strategies." IEEE Spectrum, July 1978.
- Shuttleworth, G. (1997), "Getting Markets to Clear." The Electricity Journal, Letter to the Editor, 10(3):2. April.
- Somani, A. and L. Tesfatsion, (2008), "An Agent-Based Test Bed Study of Wholesale Power Market Performance Measures," Forthcoming
- Stoft, S. (2002), Power System Economics: Designing Markets for Electricity. Wiley-IEEE Press.
- Stoft, S. (2003), "The Demand for Operating Reserves: Key to Price Spikes and Investment." IEEE Manuscript. (<http://stoft.com/metaPage/lib/Stoft-2003-spikes-invest-IEEE.pdf>). January
- Sun, J. and L. Tesfatsion, (2007), "Dynamic Testing of Wholesale Power Markets: An Open-Source Agent-Based Framework," Computational Economics, 30(2007):291-327.
- Sutton, R. and A. Barto. (1998). "Reinforcement Learning: An Introduction." MIT Press, Cambridge, MA (also available: <http://www.cs.ualberta.ca/~sutton/book/the-book.html>).
- Tuyts, K., P. J. T. Hoen and B. Vanschoenwinkel, (2006), "An Evolutionary Dynamical Analysis of Multi-Agent Learning in Iterated Games," Autonomous Agents and Multi-Agent Systems, 12(2006):115-153.
- U.S.-Canada Power System Outage Task Force, (2004), "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations," accessed at (<http://www.iwar.org.uk/cip/resources/blackout-03/>) on 5/16/09.

- Vázquez, C., M. Rivier and I. J. Pérez-Arriaga (2002), "A Market Approach to Long-Term Security of Supply." *IEEE Transactions on Power Systems*, 17(2):349-357.
- Veit, D. J., A. Weidlich, J. Yao and S. S. Oren, (2006), "Simulating the Dynamics in Two-Settlement Electricity Markets via an Agent-Based Approach," *International Journal of Management Science and Engineering Management*, 1(2006):83-97.
- Ventosa M, A. Baillo, A. Ramos and M. Rivier (2005), "Electricity market modeling trends," *Energy Policy*, 33(2005):897-913.
- Viscusi, W. K., J. M. Vernon and J. E. Harrington (2001), "Economics of Regulation and Antitrust." MIT Press, Cambridge, MA.
- Waltman, L. and U. Kaymak, (2008), "Q-Learning Agents in a Cournot Oligopoly," *Journal of Economic Dynamics & Control*, 32(2008):3275-3293.
- Watkins, C. (1989), "Learning From Delayed Rewards", Ph.D thesis, Cambridge University.
- Weare, C. (2003), "The California Electricity Crisis: Causes and Policy Options." Report Publishes by the Public Policy Institute of California. (retrieved at http://www.ppic.org/content/pubs/report/R_103CWR.pdf)
- Weidlich, A. and D. Veit, (2008), "A Critical Survey of Agent-Based Wholesale Electricity Market Models," *Energy Economics*, 30(2008):1728-1759.
- Xiong, G., S. Okuma and H. Fujita, (2004), "Multi-Agent Based Experiments on Uniform Price and Pay-As-Bid Electricity Auction Markets," 2004 International Conference on Electric Utility Deregulation, Restructuring and Power Technologies (DRPT2004), April 2004, Hong Kong.