ABSTRACT

Title of dissertation: ELECTRICITY MARKETS: PRICE RISK, POLICIES, AND POLLUTION

Daniel Patrick Werner, Doctor of Philosophy, 2014

Dissertation directed by: Dr. Sébastien Houde
Department of Agricultural & Resource Economics

The electricity sector is a significant contributor to the economic and environmental health of the United States, with annual revenues well over $300 billion and responsibility for approximately one third of all carbon emissions. The last several decades have brought significant changes to economic and environmental policies applicable to the electricity sector, including market restructuring and a variety of air quality improvement policies. This thesis builds on previous research of these issues through three related essays on energy economics and policy.

The first essay examines the local environmental impacts that can be attributed to renewable portfolio standards. Renewable portfolio standards (RPS) have been adopted by the majority of states in the U.S. to encourage electricity from renewable sources. Previous studies omit an analysis of local and regional pollutants, so this paper provides an empirical investigation into pollutant reductions from RPS while accounting for policy heterogeneity across states. Using a nation-wide panel of pollution monitoring stations in conjunction with local and national economic data, I find that adopting RPS results in significant sulfur dioxide reductions and modest nitrogen oxide reductions. I find no evidence of particulate matter reductions. Lastly, the analysis shows that pollution reductions are driven by groups of states whose neighbors also adopt RPS, which is likely because of pollution spillover effects.

The second essay examines the importance of ramping cost to electricity price volatility. High price volatility has plagued electricity market participants for decades and is increasingly important in the context of growing intermittent renewables. Although electricity market price behavior generally has been well studied in the last decade, the literature is sparse when discussing the importance of generator ramping costs to price volatility. This paper contributes to the literature by first formalizing the intuitive link between ramping costs and price volatility in a multi-period competitive equilibrium. The fundamental result of the model shows how price volatility rises with ramping costs. This notion is tested empirically using a pooled event study regression, a two-stage least squares (2SLS) specification, and a generalized autoregressive conditional heteroskedasticity (GARCH) model. The econometric results all confirm that price volatility is significantly decreased by additional natural gas capacity, which has comparatively low ramping costs. This marks the first rigorous study to quantify the pecuniary externalities within the New England market's generating profile. A simulation also explores how annualized volatility changes over time during a shifting generation profile, noting that natural gas generators can offset the volatility increases from increasing wind generation. Lastly, there is no evidence that natural gas capacity additions reduce the forward premium.

The third essay examines price convergence in the wholesale electricity markets in the context of transaction costs on virtual bids. Virtual bidding has been introduced in most restructured electricity markets in the United States with the intent to manage price risk, increase financial liquidity, and minimize deviations between forward prices and spot prices. Previous literature argues that even without virtual bids, generators can attempt to exploit the forward premium through altering bids related to physical scheduling, which is a costly way to induce price convergence. While previous literature has shown that the introduction of virtual bids does lead towards price convergence, it is also a relatively large market shock that potentially introduces new market participants with different risk preferences. This paper is the first to explicitly test the effect of increasing virtual bid transaction costs on forward price premiums using a natural experiment in a market where virtual bidding is already established. Using high-frequency price data with an event study approach, I find that increasing transaction costs on virtual bids leads to significant increases in forward premiums and significant decreases in the total number of cleared virtual bids. Additionally, my analysis supports recent literature arguing that the day-ahead prices have converged to become an unbiased predictor of real-time prices, which is an important condition for efficient markets. Lastly, I find no evidence that increasing transaction costs on virtual bids translates to increases in intra-day price volatility.

ELECTRICITY MARKETS: PRICE RISK, POLICIES, AND POLLUTION

by

Daniel Patrick Werner

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2014

Advisory Committee:
Professor Sébastien Houde, Chair/Advisor
Professor Anna Alberini
Professor Roberton Williams III
Professor Erik Lichtenberg
Professor Peter Cramton

Acknowledgments

This thesis would not be possible without continual support and guidance from university faculty, friends, and family. First, I would like to thank my primary advisor, Professor Sébastien Houde, for his feedback and encouragement on these research projects. He provided excellent guidance throughout each project by encouraging me to think critically about the research questions, empirical approach, and policy importance of my work. I am also very thankful for Professor Anna Alberini, who proved to be a great advisor during my first year qualifying paper and provided detailed help throughout the thesis as well. My thesis also benefited greatly from the guidance of my entire dissertation committee, including Professors Rob Williams, Peter Cramton, and Erik Lichtenberg. I also owe many thanks to my classmates and fellow graduate students for helping me through coursework and research discussions, including Andrew Brudevold-Newman, Davide Cerruti, Jaclyn Evans, Jeffrey Ferris, Monica Saavoss, and Ziyan Yang. Lastly, I'm forever grateful for the guidance, assistance, and wisdom received from my family. I could not have asked for better parents to support me throughout the academic process, and I could not be happier that Allison joined me on this adventure.

Table of Contents

List of Figures

List of Abbreviations

2SLS	Two-stage least squares
AQS	Air Quality System
ARMA	Autoregressive moving-average
BEA	U.S. Bureau of Economic Analysis
CAIR	Clean Air Interstate Rule
CAISO	California Independent Systems Operator
CO	Carbon monoxide
CO_2	Carbon dioxide
CSAPR	Cross-State Air Pollution Rule
DSIRE	Database of State Incentives for Renewables
EIA	United States Energy Information Agency
EPA	U.S. Environmental Protection Agency
FERC	Federal Energy Regulatory Commission
FTR	Financial transmission rights
GARCH	Generalized autoregressive conditional heteroskedasticity
ISO	Independent system operator
ISO-NE	Independent Systems Operator of New England
LMP	Locational marginal pricing
MISO	Midwest ISO
NEPOOL	New England Power Pool
NO_X	Nitrogen oxide
OLS	Ordinary least squares
PM10	Particulate matter
RECs	Renewable energy certificates
RPS	Renewable portfolio standards
RSG	Revenue Sufficiency Guarantee
RTO	Regional Transmission Organizations
SEMASS	Southeast Massachusetts
SO_2	Sulfur dioxide
VOCs	Volatile organic compounds
WLS	Weighted least squares

Chapter 1: The Local Environmental Impacts of Renewable Portfolio Standards

1.1 Introduction

In the past fifteen years, a variety of climate change legislation has been discussed at the state level in the United States. Much of this discussion revolves around energy policy, since the U.S. Environmental Protection Agency (EPA) estimates that over 60% of greenhouse gas emissions come from the transportation and electricity sectors alone (EPA, 2011*b*). The U.S. Energy Information Administration (EIA) estimates that approximately 70% of domestic electricity still comes from fossil-fuel based generators (EIA, 2012). Most potential solutions to climate change require a shift away from fossil-fuel sources of electricity and towards renewable sources. To facilitate this change, state legislators often cite renewable portfolio standards as a solution, and thirty states have adopted renewable portfolio standards since 1997 (DSIRE, 2013).[1]

Renewable portfolio standards (RPS) require a specified percentage of electric-

[1]The states with renewable portfolio standard requirements are AZ, CA, CO, CT, DC, DE, HI, IA, IL, KS, MA, MD, ME, MI, MN, MO, MT, NC, NH, NJ, NM, NV, NY, OH, OR, PA, RI, TX, WA, and WI. Other states have adopted voluntary renewable portfolio goals, including ID, ND, OK, SD, UT, VT, VA, and WV.

ity consumed within the state to be generated from renewable sources by a specific date. These goals typically ramp up gradually over time and often include penalties for utilities in non-compliance. For example, in 2002 California passed Senate Bill 1078 requiring 14% renewable electricity by 2004. The percentage requirement increases almost every year until a final requirement of 20% by 2020, although this bill was later amended to a final requirement of 33% by 2020. The stated goals of the policy were to "promote stable electricity prices, protect public health, improve environmental quality, stimulate sustainable economic development, create new employment opportunities, and reduce reliance on imported fuels" (S.B.1078, 2002).

However, there is currently a debate over the effectiveness of such policies to achieve their stated goals (Borenstein, 2012, Schmalensee, 2012). Empirical studies provide conflicting evidence that these policies encourage electricity generation within the adopting state (Adelaja and Hailu, 2008, Johnson, 2014, Kneifel, 2008, Menz and Vachon, 2006, Shrimali and Kniefel, 2011, Yin and Powers, 2010). The most rigorous approaches fully account for heterogeneity in RPS policy details across states and find that the policies are effective at encouraging within-state renewable generation (Johnson, 2014, Yin and Powers, 2010), although the policies may not be cost effective (Fell and Linn, 2013, Johnson, 2014, Schmalensee, 2012).

Since these policies are often discussed in the context of climate change, Bushnell, Peterman and Wolfram (2008) argue that a state's RPS is highly symbolic and ineffective pollution legislation because CO_2 emissions are a global issue with pollution spillovers. While regulators have developed and argued for RPS policies in order to reduce greenhouse gases, such policies may also abate local and regional pollutants. Whether RPS should instead be motivated from the standpoint of local pollution is also controversial within the United States. The negative externalities associated with local and regional air pollutants from electricity generation have been addressed by a variety of policies over the last several decades, including the Clean Air Act, Acid Rain Program, NO_X State Implementation Plan, Clean Air Interstate Rule, Ozone Transport Commission NO_X Budget Program, and RPS legislation. Simultaneous policies may have perverse effects when overlapping with other policies, which has been found in other similar contexts (de Gorter and Just, 2008, 2009), but is beyond the scope of this analysis.

Notably absent from all the empirical investigations surrounding RPS is a discussion of local pollution reduction, although local pollutants have been studied in other contexts, such as the Clean Air Act (Auffhammer et al., 2009, Greenstone, 2004) and automobile policy (Cerruti, 2014, Charmley, 2004, Wolff and Perry, 2010). Chen et al. (2009) and Heeter et al. (2014) provide excellent reviews of the simulation studies on RPS costs and benefits, some of which forecast the emissions reductions on a state-level basis (Brattle, 2010, Delmarva, 2012, IPA, 2012, LEI, 2012, NYSERDA, 2013, PUCO, 2013). These forecasts can be informative, but require specific assumptions for projecting emissions that will occur without the policy. An ex-post analysis can complement these studies by empirically providing the true pollution reductions attributable to the renewable portfolio policies. Meanwhile, emission reductions from renewable generation have been empirically studied (Cullen, 2013, Kaffine et al., 2013, Novan, 2011), though a thorough discussion of emissions attributable to specific renewable energy policies is beyond the scope of their analysis. While it is intuitive that RPS legislation will encourage electricity generation and displace fossil fuel emissions, quantifying the exact changes to local air quality is of particular interest to policymakers, economists, and public health advocates discussing the benefits and costs of additional regulation at the state and national level.

In this paper I seek to understand the role of renewable portfolio standards in reducing local and regional pollutants in an effort to provide stakeholders better estimates of the ambient air quality benefits attributable to RPS legislation. More specifically, what are the within-state reductions to ambient concentrations of sulfur dioxide (SO_2), nitrogen oxide (NO_X), and particulate matter (PM10) from RPS legislation? Given heterogeneity in renewable portfolio standards across states, how do these ancillary air quality benefits change with specific policy details?

To study these issues, I use panel data from air quality monitoring stations to estimate the effect on three local and regional pollutants associated with electricity generation, while controlling for various economic and time trends. I find that RPS legislation significantly reduces in-state sulfur dioxide levels, and accounting for heterogenous policy details across states is essential to capture this effect. However, there are only modest reductions in NO_X concentrations, and there is no statistically significant effect of RPS adoption on PM10 levels. The pollution reductions in SO_2 and NO_X only occur in states whose neighbors also adopt RPS legislation, while policies adopted in isolation are ineffective. Similar gains accrue in groups of states which allow interstate trading of renewable energy certificates (RECs), as expected since policies with regional participation minimize pollution spillover effects. These results are robust to falsification tests using pollutants emitted primarily only from automobiles, and to autoregressive moving-average (ARMA) models commonly used with high-frequency data when an adequate control group is not available.

This paper contributes to the existing literature on regulating multiple environmental externalities from the energy sector. In addition to quantifying the local and regional pollution reductions from RPS legislation, I explore how these results change with respect to specific RPS policy details. It adds to the debate on whether renewable portfolio standards are achieving their stated goals, if they should be enacted, and whether they can properly address both regional and global environmental issues (Borenstein, 2012, Bushnell, Peterman and Wolfram, 2008, Johnson, 2014, Schmalensee, 2012). Lastly, the results have important implications for air quality policies adopted in isolation by suggesting that a geographically integrated RPS legislation is required for significant improvements to air quality.

The remainder of this paper proceeds as follows. Section 1.2 provides a background of RPS legislation and the three pollutants studied in this analysis. Section 1.3 describes the econometric model and related methods, while Section 1.4 describes the data. The related results are presented in Section 1.5 and Section 1.6 concludes.

1.2 Background

Sulfur dioxide (SO_2), nitrogen oxide (NO_X), and particulate matter (PM10) constitute three of the EPA's six "criteria pollutants." Research has found that these pollutants are harmful to human and environmental health, causing acid rain and smog in addition to increasing the risk of chronic bronchitis, asthma attacks, heart attack, arrhythmia, and other lung related illnesses. Sulfur dioxide is the most mobile of theses three pollutants and can travel up to several hundred miles. NO_X does not usually travel as far because it is a chemical precursor to ozone and converts over time in the presence of sunlight and volatile organic compounds (VOCs). Meanwhile, PM10 is relatively immobile because of its particle size and can only travel up to 30 miles (EPA, 2013).

The EPA currently provides National Ambient Air Quality Standards to limit airborne concentrations of these pollutants. In addition, SO_2 and NO_X regulations include cap and trade programs on annual emissions through the 1990 Clean Air Act Amendment and the Clean Air Interstate Rule (CAIR). Although the CAIR was replaced by the more stringent Cross-State Air Pollution Rule (CSAPR), it has been delayed in the courts through the period of this analysis. To the extent that these policies are annually binding, overlapping RPS legislation cannot yield additional national reductions in emissions. However, the focus of this paper is within-state pollution concentrations which may still decrease in the presence of RPS. Given the local and regional nature of the pollutants studied, within-state pollution concentrations are still of interest to policy makers and environmental scientists. This is especially true for those in "non-attainment" areas who are required to develop and implement air quality improvement plans because they are in violation of the EPA's National Ambient Air Quality Standards. The subsequent analysis illuminates the extent to which RPS legislation may help achieve these air quality goals.

Fossil-fuel combustion in the electricity sector is the leading source of SO_2 and CO_2 emissions and the second leading anthropogenic source of emissions for both NO_X and PM10. Meanwhile automobiles are the leading source for NO_X emissions, although they are a small percentage of total SO_2 and PM10 emissions (EPA, 2011*a*). Economic simulations of renewable portfolio standards suggest they can induce reductions in carbon emissions (Chen et al., 2009, Palmer and Burtraw, 2005), but there is little discussion of the effect on the local and regional pollutants. Effective RPS legislation should decrease the relative emissions of all pollutants as non-fossil fuel generation is encouraged, but the actual reductions to ambient pollution concentrations depend on which fossil fuel units are displaced (Kaffine et al., 2013). Further, it is not immediately clear whether the pollution reductions will actually occur within the RPS state because of pollution spillover effects and unobservable environmental characteristics.

The details of RPS legislation can vary but the basic principles across policies are the same. A specific percentage of electricity sold within the state must be generated by renewable sources, or the retail utility will face penalties for non-compliance. The RPS requirements can be satisfied in a variety of ways in a restructured electricity market. Generally, states with RPS requirements include a renewable energy certificate (REC) trading program. A REC is issued for each unit of electricity generated by a producer. The generated electricity is sold into the grid via the wholesale market as usual, but the REC can be sold separately. This is an important feature that allows regulators to track which utilities are providing renewable energy, since a kWh from any source is identical. At the end of the compliance period, each utility dispensing electricity to consumers is responsible for holding a number of RECs based on the stringency of the RPS requirement.

The economics of RPS are intuitively similar to the ethanol blend mandate for gasoline in the United States (de Gorter and Just, 2009), as well as a cap-and-trade system. By mandating a specific blend of electricity, the policy effectively acts simultaneously as a subsidy for renewable electricity generation and a tax on fossil-fuel generators. The policy sets apart renewable and fossil-fuel generation as two separate cohorts which compete amongst themselves. By establishing the minimum number of required RECs and allowing trading among participants, the equilibrium price of a REC will reflect the cost premium of renewable sources above fossil-fuel sources (Berry, 2002). The subsidy received by the renewable generator varies according to the price of a REC, which encourages future development of renewable generation facilities. In a long-run equilibrium, the most cost-effective generator of carbon-neutral electricity prevails within the renewable cohort and the most cost-effective fossil-fuel generator prevails in the conventional cohort.

Although RPS effectively introduces taxes and subsidies, it is not necessarily the case that electricity prices will increase as a result. Fischer (2010) notes the net effect of RPS policies on electricity prices depends on the elasticity of supply for both renewable and nonrenewable generation. Depending on the policy stringency and relative elasticities, RPS legislation could increase generation costs and electricity prices or the policy could reduce demand for conventional generation, along with fossil fuel and electricity prices. This also has implications for the quantities of each pollutant that will be reduced from RPS legislation, since the emission reductions from new generation is not immediately clear without rigorous empirical analysis (Kaffine et al., 2013).

1.3 Econometric Model

Using ambient pollution concentrations as the dependent variable, I estimate a reduced form econometric model with two-way fixed effects, which is analogous to a difference-in-differences regression. I assume that the ambient pollution in time period t is explained by economic activity, fuel prices, total in-state electricity generation, unobserved monitoring station fixed effects, various unobserved time fixed effects, and pollution in the previous period. Since the model estimates primarily local pollutants I assume the error structure is spatially correlated within states but not at larger aggregations such as regional census divisions. Thus, the general econometric model is

$$y_{ist} = \beta_{i0} + \beta_1 y_{i,t-1} + \beta_2 RPS_{st} + \beta_3 C_{ist} + \beta_4 M_{st} + \varepsilon_{ist} \tag{1.1}$$

where y_{ist} is the ambient concentration of pollutant y at monitoring station i within state s during each day t. β_{i0} represents a monitoring station fixed effect, RPS measures the presence of an RPS policy, C is a matrix of economic controls, and M

is a matrix of month and year-region fixed effects[2] in addition to a linear time trend. Lastly, ε_{ist} is an error term clustered at the state level. While the autoregressive nature of dynamic panels can lead to biased estimates, the bias is inversely related to the total number of periods, T, when using the within estimator (Anderson and Hsiao, 1981, Arellano and Bond, 1991, Kiviet, 1995, Nerlove, 1971, Nickell, 1981). The panel data used here have a large T and use the within estimator, so the regression results should be consistent even with the lagged dependent variable, $y_{i,t-1}$, as a regressor.

In the most basic specification, RPS is simply a binary dummy variable indicating if the policy is in effect at time t for state s. Thus, β_2 is the coefficient of interest, which captures the overall effect of the policy on local pollution concentrations. Given policy heterogeneity however, this simple specification does not consider the real strength of the policy and is inadequate to capture the marginal effect of the policy. The preferred specifications use a continuous variable to account for policy variations, which is discussed in detail in Section 1.4. Although there may be concerns over the endogeneity of pollution with RPS adoption, both Lyon and Yin (2010) and Chandler (2009) show that the timing of RPS adoption is not actually correlated with pollution concentrations. Schmalensee (2012) argues that the environmental rationale for RPS adoption is unclear, further reducing endogeneity concerns. Lastly, increases to the strength of the policy occur in regularly scheduled intervals over time and it can be credibly argued that the increases are exogenous to the current period pollution. My preferred identification strategy relies

[2] This is an interaction of year fixed effects with the nine U.S. census region fixed effects.

on this variation to make a casual inference between RPS legislation and pollution reductions.

The matrix of controls, C, includes electricity fuel average prices for coal and natural gas. I construct quarterly average coal prices, by census region, to account for variations in coal type and region specific transportation costs. Since the Acid Rain Program of the 1990 Clean Air Act established SO_2 permit trading, I also interact the EPA's annual auction allowance prices with dummy variables for each respective census region to allow for heterogeneous regional effects. Collectively, this should ensure that my results are not biased by coal plants switching from high-sulfur to low-sulfur coal, as these decisions are driven by the costs captured through these variables. Although the cap-and-trade nature of the program prevents additional reductions to aggregate national SO_2 emissions, within-state emissions may still change. Further, the program works through an emissions price for all states, so non-adoption states are still an adequate control group when using the two-way fixed effects regression framework where the timing of RPS adoption is exogenous with respect to local pollution. Lastly, the average monthly wellhead price of natural gas is also included as a control. With the rapid fall in natural gas price in recent years, electricity from natural gas generators has become cheaper than electricity from coal plants,[3] raising the possibility that pollution concentrations decrease because natural gas generators are less polluting per MWh generated when compared to coal

[3]The EIA uses "total system levelized" cost to compare electricity across generator sources because it considers capital, operating, and fuel costs. They estimate $100/MWh for conventional coal and $67/MWh for combined cycle natural gas. For comparison, renewable technologies cost at $87/MWh, $144/MWh, and $90/MWh for onshore wind, solar photovoltaic, and hydroelectric, respectively. (EIA, 2013*a*)

(Cullen and Mansur, 2013, Lafrancois, 2012).

Included within the matrix of controls are variables for state-level personal income, local GDP, total state electricity generation, and corporate average fuel economy standards for passenger cars and light trucks. Real quarterly personal income at the state level captures changes in statewide economic trends across time. To capture local economic activity I include total GDP from the nearest metropolitan area. GDP per capita is not used here because measured pollution is more likely connected to total economic activity, which is highly correlated with population density. Due to data limitations, a continuous measure of local population densities is not included. Finally, I include monthly in-state electricity generation data to control for any pollution changes explained by an increase or decrease from total electricity generation. While some of the determinants of pollution from neighboring states may impact pollution locally, the model attempts to mitigate this concern through various spatial aggregations of the control variables in combination with spatial and temporal fixed effects. This should allow for proper identification of the variables of interest without substantial contamination or omitted variable bias, though this is discussed further with the results.

Lastly, the control matrix includes a variety of state renewable energy incentives, renewable policies, and national air quality policies. These include binary variables equalling one if the pollution monitoring station is present in the applicable state and year. The state policy controls include public benefit funds, property tax incentives, corporate tax incentives, generation subsidies, and renewable project grants. I control for national pollution policy controls using dummy variables for

Table 1.1: Summary of Control Variables

Control Variable	Data Source
Renewable Portfolio Standards	Database of State Incentives for Renewable Energy
Coal price	Form EIA-923, EIA-423, and FERC-423
Natural gas price	EIA-895 and EIA-176
Personal income	Department of Commerce, Bureau of Economic Analysis
Local GDP	Department of Commerce, Bureau of Economic Analysis
Electricity generation	EIA-826, EIA-923, EIA-860, and EIA-861
SO2 emission price	Environmental Protection Agency
OTC NOX Budget Program	Environmental Protection Agency
NOX State Implementation Plan	Environmental Protection Agency
Public Benefit Fund	Database of State Incentives for Renewable Energy
Property Tax Incentives	Database of State Incentives for Renewable Energy
Corporate Tax Incentives	Database of State Incentives for Renewable Energy
Generation subsidies	Database of State Incentives for Renewable Energy
Renewable project grants	Database of State Incentives for Renewable Energy
Clean Air Interstate Rule	Environmental Protection Agency
Diesel fuel sulfur standards	Environmental Protection Agency
Aircraft exhaust emission standards	Environmental Protection Agency
Corporate average fuel economy standards	Department of Transportation

the states and years affected by the Clean Air Interstate Rule, Ozone Transport Commission NO_X Budget Program, the NO_X State Implementation Plan, changes in diesel fuel sulfur standards, and aircraft exhaust emission standards.[4]

1.4 Data

Data on RPS policies and state incentive controls is taken directly from the Database of State Incentives for Renewables & Efficiency (DSIRE), which compiles information on renewable policies directly from state legislation. The project is supported by the U.S. Department of Energy and administered by North Carolina State University's Solar Center in collaboration with the Interstate Renewable En-

[4]Although an extensive review of these policies is beyond the scope of this analysis, interested readers should refer to Charmley (2004) and EPA (2002) for a discussion of automobile standards, emissions, and health impacts. Archived information with additional policy details can be found on the EPA's website: http://www.epa.gov/airmarkets/

ergy Council. The DSIRE includes the dates of each state's nominal requirement, the exemptions, and the carve-outs for specific renewables. As an example of a carve-out, Maryland's final RPS requirement is 20% by 2022 but there are additional requirements "carved out" for 2% solar generation. Further, some states allow RECs trading with other states to meet RPS requirements. The nominal percentage requirements, coverage exemptions, eligibility requirements, carve-outs, and REC trading status all serve as examples of the heterogeneity of RPS policies. As noted by Yin and Powers (2010), the heterogeneity in policy details is important to consider. Their analysis shows this is especially true when considering the nominal requirements because sometimes states which appear to have strong requirements actually have relatively weak or unbinding requirements. Thus, I use a similar measure similar to convert "nominal" RPS requirements to "real" requirements as follows:

$$REAL_{st} = NOM_{st} * COV_s - \frac{GEN_{sT}}{SALES_{sT}}$$

where $REAL_{st}$ is real RPS percentage requirement for state s in year t, NOM_{st} is the nominal RPS percentage requirement, COV_s is the non-exempt percentage covered by the policy, $SALES_{sT}$ is the total electricity sales at legislation date T, and GEN_{sT} is the total in-state renewable generation at legislation date T. This adjustment accounts for a situation in which a nominal increase in RPS strength is non-binding. For example, a state might have its RPS requirement increase from 5% to 7%, but already be producing electricity with 8% renewables at the legislation date several years ago. In this case, the percentage requirement increased, but we would expect no effect on pollution or renewable generation. Data on coverage comes directly from the DSIRE database, while sales and generation data comes from the EIA. Table 1.2 provides a cross-sectional comparison of nominal and real RPS policies for selected states in effect during 2012, including if hydro-electric power or biomass are allowed to satisfy the policy. A lengthy discussion of eligible generation is omitted because each state is different, but an extensive review of hydroelectric and biomass requirements was performed in the data construction.[5] Notably, many states with strong nominal requirements are shown to have much more modest requirement. Of the 48 contiguous states considered in this analysis, 28 have an RPS with an average nominal requirement of 10.4% renewable generation as of 2012. However, when considering exemptions to the policy and previous renewable generation prior to the policy, the real percentage requirement averages only 3.5%. Note that the final total RPS requirement may be much higher, as these policies gain strength and many do not reach their final requirements for another decade.

For the dependent variable, I use an unbalanced panel of pollution data from the EPA's Air Quality System (AQS), spanning from 1997-2012. The data controlling for other national air quality policies described in Section 1.3 also come from the EPA. The AQS database provides data from over one thousand pollution monitoring stations which were collected by local, state, and federal agencies. Depending on the monitoring station or the pollutant measured, pollutant concentrations are reported as either hourly or daily averages. Data are reported to the

[5]Interested readers should refer to the DSIRE database for additional details and legislation links.

Table 1.2: Summary Statistics: 2012 RPS Policies

State	Nominal %	Real %	Hydro	Biomass
ME	35.0	4.1	Yes†	Yes
CA	20.0	7.4	Yes∓	Yes‡
CT	16.0	11.9	Yes†	Yes‡
MA	14.1	6.0	Yes	Yes‡
MT	10.0	6.1	No	Yes‡
NM	10.0	6.7	Yes†	Yes
MD	9.0	3.8	Yes	Yes‡
TX	4.8	2.6	Yes	Yes
AZ	3.5	2.0	No	Yes
NC	3.0	0.1	Yes†	Yes‡
MO	2.0	1.1	No	Yes
OH	1.5	0.1	Yes	Yes
Total Mean	10.4	3.5		

† Currently includes new hydro only

∓ Includes small hydro only (under 30 MW)

‡ Special restrictions apply, refer to DSIRE database for additional details

EPA quarterly, but data are occasionally missing due to incomplete reporting or subsequent correction by the EPA. Low quality data are marked by the EPA using qualifier codes representing various extraordinary circumstances, for example "lab error," "operator error," or "building/site repair." All flagged values are dropped before daily concentration averages are calculated for each monitoring station. The analysis is limited to the 48 contiguous United States, and thus excludes data for Alaska, Hawaii, and Washington D.C.

Data on fuel price and electricity generation come directly from the Energy Information Agency's annual reports. Total in-state net electricity generation[6] is included at the monthly level. Coal prices are included using quarterly regional average data, while national monthly average wellhead natural gas price is included.

[6]EIA defines net generation as the total power generated by a plant, minus the amount of electricity used to run the generation functions.

Table 1.3: Summary Statistics: Pollutant Concentrations

	Mean and Standard Deviation (in parentheses) Number of Observations [in brackets]			
	1997		2012	
	No RPS States	Pre-RPS States	No RPS States	Post-RPS States
Sulfur Dioxide (ppb)	5.405	4.980	1.737	1.495
	(6.187)	(5.818)	(2.717)	(2.516)
	[68K]	[150K]	[32K]	[55K]
Nitrogen Oxides (ppb)	22.97	31.05	8.01	13.37
	(26.70)	(33.65)	(10.25)	(16.65)
	[14K]	[72K]	[17K]	[46K]
PM10 (μg/m^3)	28.32	27.56	20.57	28.69
	(20.25)	(26.90)	(16.33)	(47.45)
	[16K]	[48K]	[23K]	[38K]

These are included instead of wholesale electricity price because of the possibility for price and localized fuel prices to be endogenous to the RPS policy. While fuel price is highly correlated with electricity price, regional average fuel price is more likely exogenous to any single state adoption of RPS legislation. Further, if the fuels are traded on a national market then any local variation is the result of shipping costs and will be picked up by monitoring station fixed effects, if shipping costs remain stable over time. To the extent shipping costs vary across time but are stable within regions, the effect will be captured through the regional average prices. Lastly, the U.S. Bureau of Economic Analysis (BEA) provides state-level quarterly personal income data and annual metropolitan GDP data. I convert all nominal data into real prices using the GDP deflator provided by the BEA.

Table 1.3 provides summary statistics of daily average ambient levels for each pollutant. Since states vary in the timing of RPS adoption, simply showing pre-adoption and post-adoption averages for adoption states could be deceiving because it would include overlapping years. Instead the first two columns show a cross-section of average pollution levels for treatment and control states during 1997, while the last two columns show data from the same states in 2012, after the policies have been in effect. The data clearly show large decreases in ambient pollution levels for SO_2 and NO_X within both the adoption states and the states without RPS. Ambient concentrations of SO_2 decreased by 67.9% (3.67 ppb) in non-adoption states and 70.0% (3.49 ppb) in adoption states between 1997 and 2012. The dispersion of ambient concentrations also decreased heavily during this time period for both groups, as shown by the lower standard deviations in 2012. Ambient concentrations of NO_X decreased by 65.1% (14.96 ppb) in non-adoption states and 56.9% (17.68 ppb) in adoption states between 1997 and 2012. Meanwhile, the respective columns show decreases of 27.4% (7.8 μg/m^3) in ambient PM10 within the non-adoption states, but a slight increase of 4.1% (1.13 μg/m^3) in adoption states. Reviewing these summary statistics across time underscores the importance of controlling for time trends, such as the region-year fixed effects in the econometric regressions described in Section 1.3. In addition, it helps to motivate the regression analysis, as it is not immediately clear what portion of these reductions can be attributed to RPS legislation.

I also perform tests for pre-treatment time trends to ensure that the adoption states do not contain pre-existing trends that could invalidate the difference-in-differences strategy. I regress pollution concentrations on the control variables while interacting time trend variables with a dummy for RPS adoption states. With a p-value of 0.698, I fail to reject the null hypothesis that there are no pre-existing time trends unexplained by the control variables. The same test yields similar results for NO_X and PM10, with p-values of 0.465 and 0.160, respectively. This suggests that both the adoption states and the non-adoption states exhibit the same unobservable trends up to the point of adoption. When removing the control variables for SO_2 and only including time, geographic, and monitoring station fixed effects I still fail to reject that there are no pre-existing time trends with a p-value of 0.123. Using this method to test for separate pre-existing trends in NO_X and PM10 again yields no significant differences across adoption and non-adoption states, with p-values of 0.479 and 0.440, respectively.

The distribution of sulfur dioxide, nitrogen oxide, and PM10 concentrations are shown in Figures 1.1, 1.2, and 1.3, respectively. The "no-adoption states" label refers to the distribution for all years in the sample for states which never adopt RPS. The "pre-adoption RPS states" label shows the distribution of pollution concentrations for adoption states for all the years prior to adoption. The "post-RPS adoption states" label shows the distribution of ambient pollution for all the years after adoption. Each density graph is truncated at an EPA health standard limit,[7] and they show a similar pattern. There do not appear to be large differences in the distributions between adoption states and non-adoption states. For SO_2 and NO_X the post-adoption period appears to shift the distribution slightly towards lower concentrations, in addition to lowering the variance. Within PM10 there is no

[7] EPA limits are 30 ppb annual average for SO_2, 100 ppb hourly average for NO_X, and 50 $\mu g/m^3$ for PM10.

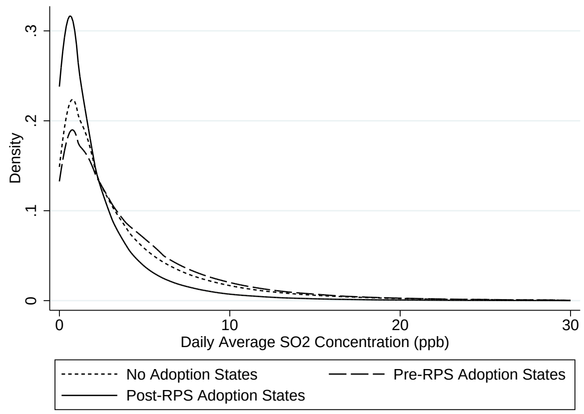

Figure 1.1: Kernel Density of SO_2 Concentrations

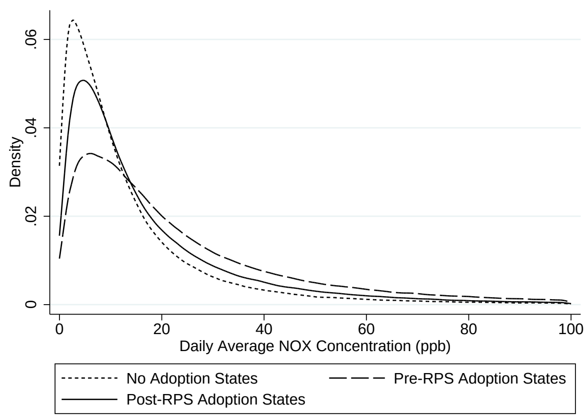

Figure 1.2: Kernel Density of NO_X Concentrations

discernable trend across groups. Further, any differences between the three density groups are negligible for all pollutants before reaching the EPA health standard. This suggests that any pollution benefits associated with RPS standards may not translate into large health benefits. However, such an investigation is left for future research.

Table 1.4 provides summary statistics of the control variables for the entire sample, non-RPS states, and RPS adoption states for the SO_2 analysis. The summary statistics of controls for the analysis of NO_X and PM10 are qualitatively iden-

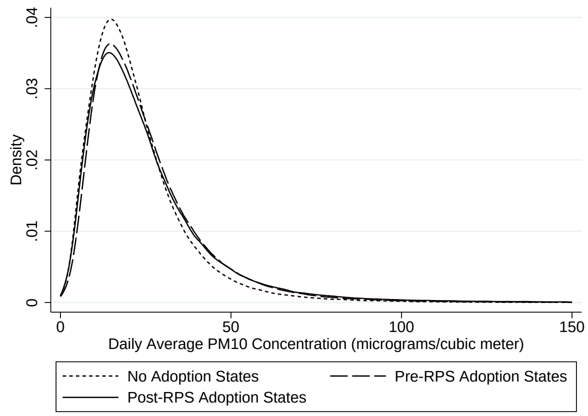

Figure 1.3: Kernel Density of PM10 Concentrations

tical and contain small, insignificant differences due to the unbalanced panel for each pollutant studied. Generally, RPS adoption states appear to have larger incomes, with a larger total GDP in the nearby metro areas. However, these summary level data are skewed by several large states which have RPS policies, notably California, Texas, and New York. Other adoption states have large metropolitan areas as well, for example Illinois and Pennsylvania. However, the large standard deviations imply plenty of overlap in the controls for both groups, so these differences are not a serious concern when coupled with the monitoring station fixed effects that captures the stable effects within heterogeneous urban areas. Finally, Table 1.4 shows little difference in the fossil fuel price data between the two groups.

Table 1.5 provides selected summary statistics organized by the nine US census divisions. Sulfur dioxide concentrations are generally higher in the Northeast and in the Southeast, which is expected because of higher population densities and number of coal plants. There does not seem to be a clear geographical pattern with nitrogen oxides, although concentrations are lowest in the West North Central and East South

Table 1.4: Summary Statistics: Controls

Mean and Standard Deviation (in parentheses)

	Total n = 2,789K	No RPS States n = 1,026K	RPS States n = 1,763K
Metro GDP	113,999	53,978	148,937
(Annual $millions)	(204,097)	(96,154)	(239,172)
Total state income	366,351	190,151	468,914
(Quarterly $millions)	(348,368)	(172,636)	(382,183)
Total state generation	10,465	8,434	11,647
(Monthly GWh)	(7,194)	(4,507)	(8,140)
Coal price	33.64	34.32	33.25
($/ton)	(17.54)	(16.77)	(17.95)
Natural gas price	4.37	4.42	4.34
($/1,000ft^3)	(1.84)	(1.84)	(1.84)

Central census divisions. The geographical discrepancies between regions could be picking up variations in population density, as regions with large urban areas may have larger NO_X concentrations due to higher automobile activity. Meanwhile, PM10 seems to be relatively consistent across regions, although the West South Central census division has the highest mean.

1.5 Results

The regression results for sulfur dioxide, nitrogen oxides, and PM10 are reported in Tables 1.6, 1.7, and 1.8, respectively. The five different specifications in columns (A) through (E) are identical across tables and each table corresponds to a different dependent variable. Every specification controls for the various fossil-fuel prices, economic activity, in-state electricity generation, and fixed effects for month, region-year, and monitoring station. To address the spatial correlation of the error

Table 1.5: Summary Statistics: Census Division

Mean and Standard Deviation (in parentheses)
Number of Observations [in brackets]

	SO_2 (ppb)	NO_X (ppb)	PM10 ($\mu g/m^3$)
New England	3.833	22.73	21.96
	(4.237)	(25.02)	(15.82)
	[173K]	[116K]	[11K]
Middle Atlantic	5.301	28.39	21.90
	(5.687)	(27.77)	(14.32)
	[433K]	[189K]	[132K]
East North Central	4.613	27.87	26.47
	(5.374)	(26.42)	(16.79)
	[535K]	[96K]	[90K]
West North Central	2.187	10.67	24.71
	(4.280)	(13.61)	(20.09)
	[333K]	[99K]	[127K]
South Atlantic	4.232	19.67	23.01
	(5.368)	(25.50)	(13.04)
	[463K]	[96K]	[151K]
East South Central	4.123	7.88	26.06
	(4.713)	(10.94)	(17.47)
	[202K]	[20K]	[67K]
West South Central	2.392	14.6	28.80
	(3.970)	(16.87)	(24.32)
	[215K]	[316K]	[56K]
Mountain	2.579	19.14	28.53
	(4.395)	(28.29)	(29.07)
	[209K]	[132K]	[496K]
Pacific	1.491	26.33	26.58
	(1.871)	(31.28)	(95.78)
	[225K]	[525K]	[173K]

terms from nearby pollution monitoring stations, the standard errors are clustered at the state level for all specifications.[8]

Column (A) of each table provides a basic specification with only a dummy variable equal to one during the periods in which the state had an RPS in effect. Column (B) includes the nominal percentage requirement, so the coefficient captures the average marginal effect of increasing the nominal requirement by one percentage point. However, as discussed in Section 1.4, this is a weak measure of strength for each policy because it does not consider whether such a percentage requirement is beyond the existing renewable generation. Columns (C) through (E) include the real RPS requirement to capture the actual strength of the RPS policy, as described in Section 1.4. The specification of Column (D) is identical to Column (C), except that it replaces the region-year fixed effects with separate fixed effects for both year and region. Lastly, Column (E) removes the policy controls to address concerns that alternative renewable energy and air quality policies are endogenous with RPS legislation due to the possibility that policymakers adopt a bundle of policies simultaneously with RPS.

Table 1.6 gives the regression results using ambient sulfur dioxide concentrations as the dependent variable. Column (A) shows a small, but statistically insignificant, decrease of 0.08 ppb in sulfur dioxide concentrations using the dummy variable for presence of an RPS policy. This is not surprising because the variation of state policy strength implies that a simple policy dummy will be ineffective at

[8]Robustness checks have also used clustering at the county level and the monitoring station level. Both of which yield qualitatively identical results.

capturing the true marginal treatment effect. Column (B) uses the nominal strength of the RPS and shows a very small, but statistically significant, marginal increase from pollution. The coefficient is interpreted as the marginal effect of an additional percentage point strength in nominal requirements, which leads to an SO_2 increase of 0.01 ppb, or 0.2% from the 1997 levels within adoption states. Again, this small positive effect is contrary to the basic intuition that these policies will either decrease pollution or be ineffective. However, my prior expectation is that nominal strength does not capture the true effect of an RPS policy because many state policies appear nominally strong but are actually quite weak.

The preferred specifications for Table 1.6 are provided in Column (C), where RPS shows larger statistically significant reductions when considering the real strength of the policy. In Column (C), a one percentage point increase in the real requirement leads to a significant reduction in SO_2 concentrations by 0.08 ppb, or 1.6% from 1997 adoption state means. Using the 3.5% average real RPS strength in 2012, this means that RPS policies are responsible for approximately 5.4% of the total decline in SO_2 concentrations from 1997 to 2012. The wide variation in marginal effects from the first three specifications provide strong evidence that properly accounting for policy heterogeneity across states is paramount to accurate analysis, and these findings are consistent with the implications of previous studies (Johnson, 2014, Yin and Powers, 2010). In Column (D) the region-year fixed effects are removed, which decreases the marginal effect, although it is not statistically different from Column (C). Removing the policy controls in the specification of Column (E) increases the magnitude of the coefficient to -0.10, but it is not statistically different from either

Table 1.6: Regression Coefficient Results: Sulfur Dioxide
Dependent Variable: SO_2 Concentration (ppb)

	(A)	(B)	(C)	(D)	(E)
Policy	-0.0844				
	(0.0846)				
Nominal		0.0146***			
		(0.0054)			
Real			-0.0773**	-0.0389	-0.1017**
			(0.0292)	(0.0284)	(0.0415)
Economic Controls	Yes	Yes	Yes	Yes	Yes
Policy Controls	Yes	Yes	Yes	Yes	No
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Region-year Fixed Effects	Yes	Yes	Yes	No	Yes
Monitoring Stations	1,039	1,039	1,039	1,039	1,039
Observations	2,789K	2,789K	2,789K	2,789K	2,789K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Standard errors are clustered and reported in parenthesis.

of the prior two columns. Using this coefficient instead of Column (C) implies RPS legislation is responsible for about 7.1% of the total decline in SO_2 concentrations from 1997-2012.

Table 1.7 shows that an RPS policy significantly decreased nitrogen oxide pollution concentrations by an average of 0.85 ppb, as shown in Column (A). This represents a 2.8% total decrease from 1997 levels in adoption states which show a total decrease of 57% from 1997 to 2012. Consistent with the previous story for SO_2, Column (B) shows an insignificant decline in NO_X concentrations from the nominal RPS requirements. The coefficient suggests that an additional percentage point requirement for renewable generation reduces NO_X concentrations by 0.02 ppb, less than 0.1% from their 1997 levels. Consistent with the previous table for SO_2, the real RPS requirement again shows much stronger decrease in ambient NO_X in Columns (C) and (D). Column (C) suggests that an increase in the real

Table 1.7: Regression Coefficient Results: Nitrogen Oxides
Dependent Variable: NO_X Concentration (ppb)

	(A)	(B)	(C)	(D)	(E)
Policy	-0.8540***				
	(0.3102)				
Nominal		-0.0219			
		(0.0247)			
Real			-0.1725	-0.2696***	-0.1822*
			(0.1060)	(0.0890)	(0.1066)
Economic Controls	Yes	Yes	Yes	Yes	Yes
Policy Controls	Yes	Yes	Yes	Yes	No
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Region-year Fixed Effects	Yes	Yes	Yes	No	Yes
Monitoring Stations	654	654	654	654	654
Observations	1,588K	1,588K	1,588K	1,588K	1,588K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Standard errors are clustered and reported in parenthesis.

requirement of one percentage point leads to reductions in NO_X concentrations of 0.17, or 0.6% for each additional percentage point in real strength. Using the 3.5% average real RPS strength in 2012, this means that RPS policies are responsible for approximately 1.9% of the total decline in NO_X concentrations from 1997 to 2012. Column (D) shows a stronger effect when region-year fixed effects are removed, becoming statistically different than zero but not statistically different from Column (C). Removing the policy controls in the specification of Column (E) changes the magnitude of the effect to -0.18, and it remains statistically significant.

The results using PM10 as the dependent variable are shown in Table 1.8. The basic specification in column (A) shows that an RPS policy is associated with an insignificant decrease in average PM10 concentrations. Column (B) shows that the nominal strength leads to insignificant reductions in PM10, while Columns (C) and (D) show a small and insignificant increase in concentrations from the real strength

Table 1.8: Regression Coefficient Results: Particulate Matter
Dependent Variable: PM10 Concentration (μg/m^3)

	(A)	(B)	(C)	(D)	(E)
Policy	-1.2815				
	(0.8958)				
Nominal		-0.0701			
		(0.0996)			
Real			0.1711	0.0457	0.2600
			(0.1864)	(0.1927)	(0.1884)
Economic Controls	Yes	Yes	Yes	Yes	Yes
Policy Controls	Yes	Yes	Yes	Yes	No
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Region-year Fixed Effects	Yes	Yes	Yes	No	Yes
Monitoring Stations	872	872	872	872	872
Observations	1,303K	1,303K	1,303K	1,303K	1,303K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered and reported in parenthesis.

of the RPS percentage requirement. Lastly, Column (E) removes the policy controls which increases the magnitude of the coefficient, with no change in statistical significance. Thus, there is little to no evidence that RPS legislation reduces ambient concentrations of PM10, which is not particularly surprising since PM10 emissions from the energy sector account for a relatively small portion of total ambient PM10.

Overall, the effects of RPS legislation on ambient pollution are expected to vary across pollutants for several reasons. First, the fraction of total emissions attributed to electricity generation varies widely. For example, the smaller percentage reduction for NO_X relative to SO_2 is not surprising since electricity generation is responsible for 85% of total SO_2 emissions but only 30% of total NO_X emissions. While fossil-fuel combustion is one of the largest human contributor to PM10 emissions, it only accounts for about 5% since the vast majority of PM10 concentrations are from dust, agriculture, and fires (EPA, 2011*a*). Further, generators vary in their relative emissions for each pollutant due to the fuel type used, so the effects on ambient pollution depend on which generators are offset by the additional renewable generation attributable to RPS legislation. However, this is the subject of debate in the previous literature because emissions reductions depend on the marginal generator type that is displaced, conditional on the ramping constraints of the generating profile (Callaway and Fowlie, 2009, Kaffine et al., 2013, Novan, 2011).

The results of this analysis compare well to estimates that can be recovered from the previous literature. Although this analysis focuses on ambient concentrations, Kaffine et al. (2013) empirically estimate the emissions reductions from wind generation in Texas. Using their estimates, I calculate an expected weighted ratio of NO_X to SO_2 reductions to be 0.22, after considering relative emissions from electricity and other sectors. Using my estimates, ambient reductions of NO_X to SO_2 are 0.36, when compared on a percentage basis. The two estimates are similar and show larger percentage reductions in SO_2 than NO_X, although my estimated ratio is slightly higher. The differences are arguably due to differences in the generating profile between Texas and the rest of the nation.

It is also worth discussing how these reductions translate to monetary benefits, although a thorough discussion of the related costs are beyond the scope of this analysis. Previous literature generally discusses the benefits from pollution reduction in terms of marginal costs per unit of emission (Cullen, 2013, Kaffine et al., 2013), which is not directly applicable in this analysis because it focuses on ambient pollution concentrations. However, I calculate a back-of-the-envelope approximation using benefit estimates from the EPA (EPA, 2010*a*,*b*), which rely heavily on the methodologies of earlier research (Fann et al., 2009, Laden et al., 2006, Pope III et al., 2002). This analysis suggests that each percentage increase in national RPS strength translates to between \$92.6 and \$228.5 million worth of health and environmental benefits from decreases in local ambient pollution alone. These benefits come through \$62 to \$153 million from SO_2 reductions when using the EPA's monetary benefits range of \$0.76 to \$1.88 billion per ppb reduction in SO_2 nationally. An additional \$30.7 to \$75.4 million comes from NO_X reductions when using the EPA's monetary benefits range of \$0.175 to \$0.430 billion per ppb reduction in NO_X nationally. However, these estimates do not consider the costs associated with RPS legislation. Lastly, these benefits calculations should be used cautiously since they rely on assumptions from previous literature that are regularly debated, such as the value of a statistical life (Viscusi and Aldy, 2003).

In considering policy differences across states, one important factor is the ability to trade RECs with neighboring states. Economic theory suggests this policy component is important to achieve pollution reductions at the least cost. Further, Yin and Powers (2010) note that REC trading leads to lower renewable generation within states that have such policies. To understand the effect of neighboring states and to investigate how REC trading affects regional pollution reduction I take the baseline specification (Column (D) of Tables 1.6, 1.7, and 1.8) and separate out the real RPS strength for two groups of states. Column (A) of Tables 1.9, 1.10, and 1.11 each show a separate treatment for states with adopting neighbor states and those who adopted without neighboring states for SO_2, NO_X, and PM10, respectively. States with adopting neighbors are those enacting RPS in addition to at least half of their bordering states.[9]

The results show that RPS policies in states with adopting neighbors had strong reductions to SO_2 and NO_X relative to states with non-adopting neighbors. The reductions of 0.11 ppb and 0.27 ppb for SO_2 and NO_X represent a marginal decrease of approximately 2.1% and 0.9%, respectively, for each additional percentage point strength in the real RPS requirement. Meanwhile, those states who adopted RPS when their neighbors did not showed no evidence that RPS policies lead to reductions in local pollutants. As with previous regressions, both groups show no statistically significant effect on PM10. This supports the idea that adopting RPS in isolation does little to achieve the stated goals without a regional policy. It also helps explain why the previous tables found that a percentage point increase in real RPS strength could lead to a greater than 1% reduction in SO_2, as pollution reductions can culminate from neighboring states. This also suggests it is possible that the determinants of ambient pollution from a neighboring states could create an omitted variable bias in the selected state. The preferred empirical specification attempts to mitigate this through various spatial aggregations of the control variables in addition to fixed effects at varying spatial and temporal aggregations. Further, the results are qualitatively identical to a robustness check specification that adds regional averages of control variables in an attempt to capture the determinants of pollution from neighboring states.

[9] Interacting the state's RPS measure with a dummy for neighboring state adoption provides a simple and intuitive coefficient to consider attenuation bias due to pollution spillover effects. While the selection of using at least half the bordering states is arbitrary, the analysis is also done using 2/3 of bordering states with qualitatively identical results. Using an average of neighboring states is not done because it essentially reduces to a regional average without sufficient variation or statistical power.

Table 1.9: Alternative Regression Results: Sulfur Dioxide
Dependent Variable: SO_2 Concentration (ppb)

	(A)	(B)	(C)	(D)
Real (non adopting neighbors)	-0.0139			
	(0.0405)			
Real (adopting neighbors	-0.1049***			
	(0.0358)			
Real (non trading)		0.0113		
		(0.0646)		
Real (trading)		-0.1042***		
		(0.0314)		
Real (upwind avg)			-0.0589	
			(0.0797)	
Real (current year)				-0.1191**
				(0.0480)
Future Real (1 year)				0.0236
				(0.0308)
Future Real (2 years)				-0.0391
				(0.0329)
Economic & Policy Controls	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes
Region-year Fixed Effects	No	No	Yes	Yes
Monitoring Stations	1,039	1,039	1,039	780
Observations	2,789K	2,789K	2,789K	2,020K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered and reported in parenthesis.

Table 1.10: Alternative Regression Results: Nitrogen Oxides
Dependent Variable: NO_X Concentration (ppb)

	(A)	(B)	(C)	(D)
Real (non adopting neighbors)	-0.1895			
	(0.1406)			
Real (adopting neighbors	-0.2717***			
	(0.0859)			
Real (non trading)		-0.0038		
		(0.1759)		
Real (trading)		-0.2848***		
		(0.0827)		
Real (upwind avg)			0.0337	
			(0.2168)	
Real (current year)				-0.3960**
				(0.1841)
Future Real (1 year)				0.1419
				(0.1308)
Future Real (2 years)				-0.0808
				(0.1287)
Economic & Policy Controls	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes
Region-year Fixed Effects	No	No	Yes	Yes
Monitoring Stations	654	654	654	491
Observations	1,588K	1,588K	1,588K	1,103K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Standard errors are clustered and reported in parenthesis.

Column (B) of these same tables uses a similar technique to separate the RPS effect among states with a REC trading program and those without a REC trading program. A state is considered to have a trading program if they allow out-of-state RECs to satisfy the requirements. As these are a similar group of states to those which have neighbors adopting RPS policies, I expect a similar result. Larger pollution reductions are expected in states which trade RECs because of spillover effects, even though trading may lower the within-state subsidies for renewable generation. The results for SO_2 and NO_X are very similar to Column (A). States allowing REC trading decreased SO_2 and NO_X by 2.1% and 0.9%, respectively, for each percentage point of real RPS strength. The non-trading states showed no significant reductions in SO_2 and NO_X. Meanwhile, the PM10 results do show weak statistical significance, with trading RPS states reducing PM10 concentrations by 1.3% and non-trading states increasing concentrations by 1.2%. This result should be interpreted with caution and weighed against all previous regressions showing RPS policies had no statistically significant effect on PM10 concentrations.

Column (C) of Tables 1.9, 1.10, and 1.11 addresses the issue of upwind and downwind states when considering pollution and abatement spillovers. Generally, pollution flows east because of the atmospheric flow known as the Westerlies. Thus, the in-state real RPS requirement is replaced with an average of all real RPS state requirements within 200 miles west of the treated state, weighted by electricity generation. More formally,

Table 1.11: Alternative Regression Results: Particulate Matter
Dependent Variable: PM10 Concentration (μg/m^3)

	(A)	(B)	(C)	(D)
Real (non adopting neighbors)	-0.0490			
	(0.3006)			
Real (adopting neighbors	0.1109			
	(0.2431)			
Real (non trading)		0.3276**		
		(0.1569)		
Real (trading)		-0.3730*		
		(0.1995)		
Real (upwind avg)			0.7059**	
			(0.2731)	
Real (current year)				-0.9342**
				(0.4405)
Future Real (1 year)				0.3834
				(0.2310)
Future Real (2 years)				0.2060
				(0.2470)
Economic & Policy Controls	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes
Region-year Fixed Effects	No	No	Yes	Yes
Monitoring Stations	872	872	872	506
Observations	1,303K	1,303K	1,303K	820K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered and reported in parenthesis.

$$REALAVG_{st} = \frac{\sum_{i \in W} GEN_{it} * (NOM_{it} * COV_i - \frac{GEN_{iT}}{SALES_{iT}})}{\sum_{i \in W} GEN_{it}}$$

where $REALAVG_{st}$ is the real upwind weighted average RPS percentage requirement for state s in year t, NOM_{it} is the nominal RPS percentage requirement, COV_i is the non-exempt percentage covered by the policy, $SALES_{iT}$ is the total electricity sales at legislation date T, GEN_{iT} is the total in-state renewable generation at legislation date T, and $i \in W$ represents all states i within 200 miles west of state s, inclusive of s.

Column (C) of Tables 1.9, 1.10, and 1.11 shows that weighted average policies from upwind states have no statistically significant reductions in SO_2, NO_X, or PM10. This suggest the the results are not driven by large movements in pollution from states which are 200 miles west, even though pollution can travel long distances in some cases. However, exercise caution when interpreting the results of this column because the assumption of eastern flowing pollution is a generalization of long run wind patterns that is not always the case locally. Column (C) of Table 1.9 shows the largest decrease from upwind states within SO_2, with a one percentage increase in real upwind average policies decreasing ambient concentrations by 0.06 ppb, or 1.2% of 1997 averages. Although statistically insignificant, this is in line with expectations, as coal-based electricity generation in the midwest has been known to increase acid rain deposits even as far as the northeast. Regarding NO_X, Column (C) of Table 1.10 shows that the weighted average of upwind state policies leads to an insignificant increase of 0.1%. Although different from SO_2 effects, this is still within expectations since NO_X converts to ozone over long distances when exposed to VOCs and sunlight so I expect no effect. Lastly, Column (C) of Table 1.11 shows no evidence that upwind RPS averages decreased PM10 concentrations, and actually shows a small positive correlation.

Finally, Column (D) addresses the issue of anticipatory responses to future binding policies. It is possible that renewable capacity could be installed in the current year in anticipation of a future year's binding policy, even if it has no real strength in the current year. If renewable generators are being constructed due to the anticipated binding strength of the policy in a future year, it could bias my results. Intuitively, this may create a downward bias because pollution reductions occurring in the control period are captured by annual fixed effects instead of the variable for the real policy strength. To test this hypothesis, I include the future real RPS strength in Column (D) of Tables 1.9, 1.10, and 1.11 for SO_2, NO_X, and PM10, respectively. As shown by the tables, the primary results do not appear to be driven by anticipatory responses. The "future real" strength of the policy in either of the next two years does not lead to significant reductions in pollution for any of the pollutants studied. The magnitude of the marginal effect is insignificantly larger for SO_2 and NO_X reductions from the current year real RPS requirement, when compared to the closest specification in Column (C) in Tables 1.6 and 1.7, respectively.

1.5.1 Robustness Checks

Tables 1.9, 1.10, and 1.11 suggest that pollution spillovers from neighboring states may prevent clean identification of the policies because they lack an isolated and comparable control group. Thus, another possible model takes advantage of the high frequency time series data when a good control group is not available, through vector autoregression models more commonly used with high frequency data in finance and macroeconomics. The autoregressive moving-average (ARMA) model of autoregressive order p and moving average order q is

$$y_t = \beta_0 + \sum_{i=1}^{p} \beta_i y_{t-i} + \beta_2 RPS_{it} + \beta_3 C_{it} + \sum_{j=1}^{q} \beta_j \varepsilon_{j-i} + \varepsilon_t$$

where y_t is the ambient pollutant concentration at a monitoring station during time t, RPS represents the real strength of an RPS policy, C is a matrix of economic controls, and ε_t is an error term correlated across time.

The ARMA model of order 1 is calculated for each monitoring station separately and for each of the three pollutants. All three pollutants give results within the expectations of the preferred specification discussed previously. For SO_2, the median of all ARMA results shows a slightly lower marginal effect of -0.03 ppb, or about a 0.6% decrease for each additional percentage point RPS requirement. The effect of RPS legislation on NO_X increases under the ARMA model to -0.55 ppb, about 1.7% of pre RPS levels. Lastly, PM10 again shows little evidence of reductions from RPS policies, with a median insignificant effect of 0.178 μg/m^3.

To further test the robustness of the panel data results, I perform an identical analysis on two control pollutants as a falsification test. While about three-fourths of all sulfur dioxide emissions are from generating electricity from fossil fuel, automobile activity is a significant contributor to emissions of NO_X, PM10, and a host of other pollutants (EPA, 2011*a*). For example, automobiles make up around 62% of NO_X emissions while fuel combustion[10] is responsible for about 30%. Meanwhile, the EPA notes that fuel combustion is responsible for approximately 84% and 5% of SO_2 and PM10 emissions, respectively, while automobiles account for approximately 2% of both SO_2 and PM10 emissions. To ensure that my findings are not the result of changes in unobservable automobile activity that is not captured in the econometric model through time and policy fixed effects, I use carbon monoxide (CO) and volatile organic compounds (VOC) as control pollutants. The EPA notes that automobiles are responsible for approximately 86% of carbon monoxide emissions, while total fuel combustion makes up only 6%. Automobiles are also responsible for about 45% of VOC emissions, while total fuel combustion is responsible for under 4%.

The summary statistics for the control pollutants are shown in Table 1.12, while the regression results are shown in Tables 1.13 and 1.14 for carbon monoxide and VOCs, respectively. The summary statistics of pollution concentrations follow a similar pattern to the original three pollutants of interest. Pollution concentrations seem to decrease in both non-adoption states and adoption states from 1997 to 2012.

The regression results for both CO and VOC concentrations show that the real strength of RPS policies have no statistical significance, consistent with our expectations. For CO, the first two columns actually show small but significant increase in

[10]Total fuel combustion as described by the EPA includes electricity generation, industrial boilers, commercial, and residential activities

Table 1.12: Summary Statistics: Control Pollutants

	Mean and Standard Deviation (in parentheses) Number of Observations [in brackets]			
	1997		2012	
	No RPS States	Pre-RPS States	No RPS States	Post-RPS States
Carbon Monoxide (ppb)	758.2	813.9	312.8	298.9
	(517.4)	(608.9)	(222.8)	(216.5)
	[37K]	[134K]	[13K]	[43K]
VOCs (ppb)	80,584	4,333	31,791	224
	(93,691)	(9,301)	(36,653)	(1,837)
	[1.4K]	[8.5K]	[1.1K]	[5.3K]

Table 1.13: Robustness Check Regression Results: Carbon Monoxide
Dependent Variable: CO Concentration (ppb)

	(A)	(B)	(C)	(D)	(E)
Policy	10.610**				
	(4.984)				
Nominal		1.464**			
		(0.674)			
Real			1.696	0.100	1.702
			(2.432)	(1.904)	(2.445)
Economic Controls	Yes	Yes	Yes	Yes	Yes
Policy Controls	Yes	Yes	Yes	Yes	No
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Region-year Fixed Effects	Yes	Yes	Yes	No	Yes
Monitoring stations	821	821	821	821	821
Observations	2,164K	2,164K	2,164K	2,164K	2,164K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered and reported in parenthesis.

Table 1.14: Robustness Check Regression Results: VOCs
Dependent Variable: VOC Concentration (ppb)

	(A)	(B)	(C)	(D)	(E)
Policy	327.22				
	(881.04)				
Nominal		6.75			
		(27.73)			
Real			19.11	152.56	34.92
			(178.60)	(295.60)	(235.38)
Economic Controls	Yes	Yes	Yes	Yes	Yes
Policy Controls	Yes	Yes	Yes	Yes	No
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Region-year Fixed Effects	Yes	Yes	Yes	No	Yes
Monitoring stations	305	305	305	305	305
Observations	192K	192K	192K	192K	192K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered and reported in parenthesis.

CO from the policy dummy and nominal RPS strength, which further underscores the importance of accounting for RPS heterogeneity. As with the primary regression results in Tables 1.6, 1.7, 1.8, improper consideration of RPS strength can lead to significantly different marginal effects from the policy. Taken together, the results of this falsification test provide strong evidence that the primary results of this paper are not driven by changes in automobile activity.

An additional falsification test is provided by shifting the date of the RPS change backwards by five years, which should show no significant marginal effects. As a result, the five most recent years are dropped from the sample. The results are shown in Table 1.15, with each column providing the regression results for the ambient pollution concentration of the variable listed in the top row. The specification is exactly the same as the baseline specification, so it can be compared to Column (D) of Tables 1.6, 1.7, and 1.8 for SO_2, NO_X, and PM10, respectively. As expected,

Table 1.15: Falsification Test

	SO_2	NO_X	PM10
Real	-0.0135	-0.0863	0.3769
	(0.0403)	(0.1515)	(0.2337)
Economic Controls	Yes	Yes	Yes
Policy Controls	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes
Region-year Fixed Effects	No	No	No
Monitoring stations	601	374	396
Observations	1,399K	748K	552K

Note: ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered and reported in parenthesis.

the results from each column show no significance to the falsified treatment time, with the coefficients much closer to zero for SO_2 and NO_X.

Lastly, a "reasonableness test" is performed using annual state level emission data provided by EPA from 1997-2012. A simple regression uses the natural log of state-level emissions as the dependent variable, and the results are consistent with the previous analysis. SO_2 and NO_X are significantly reduced by the "real RPS" measure, with larger SO_2 reductions when compared to NO_X on a percentage basis. Unfortunately, the equivalent data was not available for PM10, CO2, and VOCs.

1.6 Conclusion

Renewable portfolio standards are often encouraged by lawmakers as a mechanism to achieve better air quality, further investment in renewable energy, enhance diversity in electricity generation, and create jobs. While pollution reduction from RPS is usually discussed in the context of climate change, local pollution benefits have not been fully studied. As a result, rigorous empirical investigations into the effect of RPS on local pollutants have been notably absent in the previous academic literature. Empirical work becomes increasingly difficult given the heterogeneity of RPS policies across states. However, as shown in the analysis above, this is important to consider because some states have strong nominal requirements that may be rendered non-binding of because large pre-existing renewable capacity and utility exemptions. The resulting regression from a simple difference-in-difference model using a dummy treatment variable will produce bias estimates of the true policy effect. A similar notion holds for the nominal RPS strength, which is shown to be a poor measure of the real treatment intensity.

Quantifying the local ambient pollution reductions that can be attributed directly to RPS also adds to the general debate over adopting RPS legislation in the presence of overlapping policies. If local pollutant externalities are already properly priced through other policies such as feed-in-tariffs or pollutant specific cap-and-trade programs, then the additional reductions to local pollutants shown here will contribute to a deadweight loss. However, if the other mechanisms are inefficiently designed and underprice the true marginal cost associated with local pollutants, then these additional pollution reductions are interpreted as ancillary benefits of the RPS that move towards an efficient allocation in a second-best scenario.

Although prior literature has forecast local emission reductions from RPS and empirically analyzed emissions reductions from specific renewable sources, this paper marks the first rigorous approach to empirically quantify the within-state pollution reductions from RPS legislation, ex-post. I interpret the regression results as strong evidence for reductions in SO_2 from RPS adoption, but only if neighboring states also adopt. Of the large reductions in ambient sulfur dioxide concentrations from 1997 to 2012, my analysis suggests that approximately 5.4% can be attributed to RPS policies. However, for states whose neighbors also adopt RPS legislation, the RPS policies could be responsible for up to 7.4% of the ambient reductions in SO_2. This trend is more moderate when looking at NO_X, showing average reductions of 2% of total NO_X reductions from 1997 to 2012 and 3% reductions for states with adopting neighbors. Lastly, the policies seem ineffective at reducing PM10 concentrations, a robust finding across all model specifications. The variations in ambient reductions across pollutants are arguably due to the type of fossil fuel generation that is offset by additional capacity attributable to RPS. The fundamental results of the analysis are shown to be robust to falsification tests using carbon monoxide and VOCs as control pollutants. The results are also robust to monitoring station specific ARMA models, regressions which drop the non-adopting states,[11] and weighted least squares specifications.[12]

This analysis suggests that each percentage increase in national RPS strength translates to between \$92.6 and \$228.5 million worth of health and environmental benefits from decreases in local ambient pollution alone, though the monetary benefit calculation does not consider the costs associated with RPS legislation or the benefits of CO_2 reduction. Further, these monetary benefits should be interpreted

[11] In the regression dropping the non-adopting states, the "control" group is comprised of states whose real RPS requirement does not change in addition to states who eventually adopt RPS but have not yet adopted.

[12] Due to concern over the non-random placement of monitoring stations, the weighted least squares (WLS) specification weights by the inverse number of monitoring stations within a county. This robustness check ensures that the results are not being driven by particular behavior in specific counties with a large number of monitoring stations. The WLS specification yields qualitatively identical results, and small differences are statistically insignificant.

with caution, as they are approximated using benefit calculations from previous researchers, which use assumptions that are the subject of academic debate.

By comparing the nominal percentage requirements to the real percentage requirements it becomes clear that many state policies are merely symbolic, although this may change in the future as they gain strength over time. In fact, some of the nominal percentage requirements have yet to be higher than the status quo at the time of passing legislation. Thus, in many states RPS legislation appears to be more of a political statement than an effective environmental regulation, supporting the argument of Bushnell, Peterman and Wolfram (2008) described in Section 1.1. However, as these relatively weak policies ramp up over the next decade, it is possible that additional environmental benefits will begin to accrue. Further, when policymakers are considering RPS legislation, they must consider the policies of other states in their region as it is clear that pollution can flow across state borders. Local pollution reductions only accrue in states with regional REC trading programs where neighboring states also adopt, and isolated adoptions are shown to be ineffective at reducing ambient concentrations. Overall, this analysis suggests that other policies aimed at addressing the negative externalities associated with local and regional pollution from electricity production may need to be adjusted in the presence of RPS adoptions by states.

Chapter 2: Electricity Market Price Volatility: The Importance of Ramping Costs

2.1 Introduction

Within the past fifteen years, most electricity markets across the United States have restructured to allow competition in the generation of electricity. Electricity price behavior has been a concern in restructuring activities and related wholesale electricity market design (Bask and Widerberg, 2009, Borenstein et al., 2002, Bushnell, Mansur and Saravia, 2008, Chang and Park, 2007, Metaxoglou and Smith, 2007, Wolak and Patrick, 2001). Price volatility has been examined generally (Hadsell et al., 2004, Higgs, 2009, Higgs and Worthington, 2008, Worthington et al., 2005, Zareipour et al., 2007), however there are fewer studies which examine how price volatility is influenced by the generating profile of the market, even though this is increasingly important in the context of rising renewables, aging nuclear facilities, and President Obama's recent Clean Power Plan which may result in additional coal plant closures. High electricity price volatility has plagued wholesale electricity prices since restructuring, creating major implications for risk-averse market participants and system operators tasked with grid reliability. Further, price volatility is a primary input into conventional options pricing models, pushing real costs onto consumers of electricity as power purchasing retailers use costly options to hedge away from price risk. When compared to other energy commodities, intra-day volatility in wholesale electricity markets is many times larger and varies across regions. For example, daily electricity market volatility ranges from 6-28% compared to 1-1.5%, 2-3%, and 3-5% for stock indices, crude oil, and natural gas, respectively (Simonsen, 2005, Zareipour et al., 2007).

With the rise of non-dispatchable renewable generators such as wind and solar, short run volatility has grown increasingly important to Regional Transmission Organizations (RTO) managing the electricity grid (Navid and Rosenwald, 2012). To ensure adequate ability of adjusting generator output, known as ramping ability, grid operators are discussing alterations to the current market design in California (Xu and Tretheway, 2012) and the Midwest (Navid and Rosenwald, 2013). Under the standard market design, ramping ability will be properly priced in a deterministic model because flexible generators that can quickly adjust output will be able to profit from large movements in price. In practice however, actual market conditions often deviate from those previously scheduled by the RTO. It has been shown that the current market design may not properly price short-term ramping ability due to suboptimal dispatching under uncertainty (Angelidi, 2012, Wang and Hobbs, 2014). When designing an electricity market for the generating profile of the future, it is important to thoroughly understand how different generator types affect price volatility. The increased investment in natural gas generators over the last decade and the anticipated rise of renewables motivates the focus of this paper, since the direct effect of natural gas capacity on electricity price volatility has not been well studied.

Much of the variability in electricity prices is driven by the physical characteristics of electricity, notably the requirement to perfectly adjust supply to meet a demand that varies significantly throughout the day and across seasons. The mainstream view is that high price volatility within electricity markets is due to the lack of hourly retail pricing in combination with the lack of cost-effective electricity storage mechanisms. In traditional commodity markets, forward contracts stabilize spot prices because any deviations allow for arbitrage through selling previously stored goods (Kaldor, 1939, Working, 1948). However, current technologies do not allow cost-effective electricity storage on any meaningful scale, rendering traditional forward pricing models inapplicable. Instead, Bessembinder and Lemmon (2002) develop a seminal equilibrium model of forward contracts between risk-averse electricity generators and retailers, within the context of nonstorable commodities. Their work implies a forward contract premium to accompany high expected demand or demand variance. The essentials of their model are empirically supported (Cartea and Villaplana, 2008, Douglas and Popova, 2008, Lazarczyk, 2013, Longstaff and Wang, 2004), though more recently Haugom and Ullrich (2012*b*) find that the forward price has converged to an unbiased predictor of the spot price.

While Bessembinder and Lemmon (2002) capture the essentials behind forward contracts in non-storable commodities, their model ignores the storability of inputs to electricity generation. Intuitively, if inputs can be stored and capacity exists to instantaneously convert these inputs into electricity, then a stabilizing pressure is applied to price during unexpected demand shocks. Storable fossil fuel inputs provide an indirect storage of electricity, assuming sufficient capacity exists with little to no production ramping constraints. Different generator technologies would affect volatility differently, as they vary in their ability to adjust output. Heterogeneity in ramping costs, or costs of adjusting output, allow some generators to flexibly adjust output during periods of higher demand (Reguant, 2014), putting more downward pressure on prices compared to other generators. Thus, there exists a cross-commodity price relationship as pointed out by Routledge et al. (2001) in an extension of their previous work (Routledge et al., 2000).

A similar notion is empirically tested by Douglas and Popova (2008), who find that larger natural gas storage decreases the premium of forward contracts in electricity markets. While they note that the effectiveness of the indirect physical hedge requires availability of transmission and generation capacity, this is absent from their empirical specification. Further, natural gas storage is likely endogenous to electricity price and forward contract premiums, creating bias in their empirical estimates. In a separate analysis across European electricity markets, Huisman and Kilic (2012) attribute differences in risk premiums to be from differences in the storability implicit within the generation profile, a point more explicitly noted previously (Huisman and Kilic, 2010). However, cross-sectional analysis is inadequate to infer causal relationships when the markets also vary widely in observable and unobservable characteristics. My empirical analysis improves on this literature by explicitly addressing endogeneity issues associated with the supply of generator types and electricity price.

In this paper, I seek to understand the role of ramping costs in the price volatility of non-storable and perishable commodities. More specifically, I ask three connected research questions related to natural gas capacity, which has comparatively low ramping costs (Reguant, 2014, Wolak, 2007). First, what is the impact of additional natural gas capacity on electricity price stability and how does this compare to inflexible capacity such as nuclear? Next, how does the forward premium change on price contracts in the presence of additional natural gas capacity? Finally, what is the value of such volatility reductions to power purchasers and how does this change with the rise of intermittent renewable generators?

To explore this topic, a basic theoretical framework is developed to establish the connection between price volatility and generator ramping costs. Under standard economic assumptions, the analytical model clearly suggests that price volatility increases with generator ramping costs. Further, the theoretical model implies a reduced form econometric specification where the intra-day price volatility is a function of natural gas capacity, intra-day demand volatility, daily average demand, and unobservable time trends. To explore these ideas empirically, I use high-frequency price data from the New England Independent Systems Operator for the period 2005-2011. Data on natural gas capacity and nuclear capacity outages are taken from the U.S. Energy Information Agency and the U.S. Nuclear Regulatory Commission, respectively. The task is complicated by endogeneity between price and capacity, since natural gas is the marginal generator in New England, but my empirical strategy explicitly addresses this issue.

The preferred results include a pooled event study regression, which finds strong evidence that natural gas capacity additions reduce price volatility an order of magnitude more than additional nuclear generation capacity. These results are robust to a two-stage least squares (2SLS) specification, as well as a generalized autoregressive conditional heteroskedasticity (GARCH) model. I attribute the differences in volatility reductions between the two generator types to the relatively low ramping costs of natural gas. Lastly, a simulation explores how the volatility impact from natural gas will change over time, in the presence of non-dispatchable renewable generation. In terms of volatility impact, natural gas provides an excellent complement to new wind generation in the New England market.

This research adds to the broad existing literature that discusses electricity market design (Navid and Rosenwald, 2012, Reguant, 2014, Wang and Hobbs, 2014, Wolak and Patrick, 2001), market efficiency (Borenstein et al., 2002, Metaxoglou and Smith, 2007), electricity price behavior (Hadsell et al., 2004, Worthington et al., 2005), and forward premiums on perishable commodities (Bessembinder and Lemmon, 2002, Douglas and Popova, 2008, Haugom and Ullrich, 2012*b*, Longstaff and Wang, 2004). By formalizing the link between ramping costs and price volatility, the model provides a clear theoretical mechanism to explain how ramping costs increase price volatility. Most importantly, this research provides the first rigorous empirical analysis that supports the role of natural gas capacity to reduce price volatility. This research provides concrete evidence for policymakers to consider the pecuniary externalities associated with generation types. This underscores the importance of investments into ramping ability, which adds to the current discussion on market design alterations. While environmental externalities are beyond the scope of this analysis, ramping costs are also important for such researchers to consider because they can fundamentally alter the abatement cost curves, as they may change the dispatch order of generators.

The remainder of this paper proceeds as follows. Section 2.2 discusses ramping costs in more detail and the theoretical framework is established in Section 2.2.1, which formalizes the intuitions described above into a basic analytical model. The econometric strategy to test these relationships is described in Section 2.3. A brief background of the New England electricity market structure is provided in Section 3.2, while the related data is discussed in Section 2.4.1. The econometric results and primary empirical contributions are discussed in Section 2.5. Finally, additional regression analysis studying the impact of natural gas capacity on the forward premium is provided in Section 2.6, while Section 2.7 concludes.

2.2 Ramping Costs

Electricity generation is itself a complex process, made more complicated through the necessity of balancing supply and demand instantaneously to prevent grid failure. In typical fossil-fuel generators, fuel is burned to convert the embedded chemical energy into thermal energy which heats up water into steam. The pressurized steam flows to turn a turbine, which is connected to a generator that converts the mechanical energy into electricity. Nuclear reactors work in a similar way, except the nuclear reaction creates the heat for the steam turbine.

The mechanical complexity inherent to the generation process imposes extra costs to adjusting electrical output from hour to hour, known as ramping costs. Ramping costs appear through fixed investments as well as marginal costs. Within the fixed costs, physical ramping constraints accompany certain technologies and these require higher investments to overcome. For example, the turbine system and related components require special designs and construction materials to be able to rapidly ramp output and to withstand the extra stress of ramping without failure (Tanaka, 2006).

Regarding marginal costs, previous literature notes that ramping output up or down will decrease the fuel efficiency of the unit compared to a constant operating output. Further, ramping output puts additional stress on the generator components, leading to larger replacement costs. More specifically, ramping induces rapid pressurization and decompression which stresses essential pieces such as the rotor, turbine shaft blades, boiler, and turbine chamber (Tanaka, 2006). This thermal stress induces microscopic fractures known in the engineering literature as “fatigue damage”, which is the second leading cause of boiler tube failure (EPRI, 2006).

Engineering studies also note that fatigue damage to the rotor assembly increases non-linearly with ramping speed and can alter the optimal commitment of generating units (Wang and Shahidehpour, 1994, 1995). Regarding the efficient dispatch of generators, Shrestha et al. (2004) note that ramping may be used strategically in deregulated markets. They point out that, in general, generators start up and shut down slowly to avoid any ramping costs and turbine damage. However, during periods of high prices it can be profitable to incur ramping costs if the generator has sufficient capacity. This is consistent with the intuition behind the theoretical and empirical approach in Sections 2.2.1 and 3.3, respectively.

There are also indirect costs associated with ramping ability. The lower ramping costs associated with natural gas generators presumably enhance grid stability and allow reliable grid operation with lower operating reserves to be sufficient. Additionally, if sufficient capacity does not exist with ramping capabilities to accompany demand changes then there is a large risk of system blackouts. These considerations are discussed by Chao (1983), as blackout risk imposes significant economic costs. However, my analysis is concerned primarily with price risk, so changes to the probability of grid failure due to ramping ability is left for future researchers.

Since the focus of this analysis is on natural gas capacity and nuclear capacity, it is worth noting their differences in ramping ability. The marginal operating costs of nuclear generators are estimated to be one fourth of natural gas generator marginal costs (EIA, 2013*b*) so they generally provide the base load of the electricity supply. Further, technical constraints make cost-effective hourly ramping of nuclear generators infeasible. Nuclear generators may take an entire day to start up or shut down during planned outages, although in emergency situations the reactor can shut down very quickly. Meanwhile, natural gas generators are considered more flexible and follow increases in demand throughout the day. This is confirmed by previous literature which finds that natural gas generators have ramping costs an order of magnitude lower than coal (Reguant, 2014, Wolak, 2007).

Wind and solar generators are non-dispatchable technologies without ramping options, and they are ignored in the empirical analysis because they represent an insignificant portion of supply within the ISO-NE. However, their growing presence increases the relevance of the issues studied here because their inherent supply intermittency increases the volatility of residual demand satisfied by dispatchable generators, such as natural gas. This impact is explored using the simulation in Section 2.5.1.

2.2.1 Theoretical Model

Before discussing the empirical approach, this section formalizes the economic intuition into a basic dynamic model where firms generate electricity to maximize daily profits, π, in a competitive wholesale market. Each day a representative firm i chooses the optimal quantity of electricity, q, to produce in hour h, in order to maximize their profits. Assuming a competitive wholesale market, firms are given hourly market clearing electricity prices, p_h. The model uses a simple generalized cost structure similar to the previous literature (Reguant, 2014, Wolak, 2007), and assumes a convex production cost function, $C_i(q_h)$. There is also assumed to be convexities in the ramping cost function, $R_i(\Delta_{i,h})$ where the change in hourly production is denoted as $\Delta_{i,h} = |q_{i,h} - q_{i,h-1}|$. Demand, D, is exogenous because consumers face a regulated retail price that prevents hourly price pressure, as discussed in additional detail in Section 3.2. Adding fixed costs, F, yields the following objective function for production firms:

$$\max_{q_{h,i}} \pi_i = \sum_{h=1}^{24} \delta_h [p_h q_{i,h} - C_i(q_{i,h}) - R_i(\Delta_{i,h})] - F_i \tag{2.1}$$

subject to $\pi \geq 0, q_h \geq 0, D_h = \sum_{i}^{n} q_{i,h}$

where δ_h is the hourly market discount factor and n is the number of firms. The first two constraints represent non-negative production and non-negative daily profits, though hourly profits can be negative. The final constraint is the standard market clearing condition where production equals demand. Solving for the first order conditions yields the standard result of price equal to marginal costs, for each firm i in hour h:

$$p_h = \frac{\partial C_i}{\partial q_{i,h}} + \frac{\partial R_i}{\partial \Delta_{i,h}} \frac{\partial \Delta_{i,h}}{\partial q_{i,h}} + \frac{\delta_{h+1}}{\delta_h} \frac{\partial R_i}{\partial \Delta_{i,h+1}} \frac{\partial \Delta_{i,h+1}}{\partial q_{i,h}} \tag{2.2}$$

Recall that the intra-day variance of p on day t, denoted by σ_t^p, is defined:

$$\sigma_t^p = \frac{1}{24} \sum_{h=1}^{24} (p_{t,h} - \bar{p}_t)^2 \tag{2.3}$$

where $\bar{p}_t$ is the daily average price. Substituting in equation (2.2) to equation (2.3) and simplifying yields the fundamental result of this model:

$$\sigma_t^p = \frac{1}{24} \sum_{h=1}^{24} \left(\frac{\partial C_i}{\partial q_{i,h}} + \frac{\partial R_i}{\partial \Delta_{i,h}} \frac{\partial \Delta_{i,h}}{\partial q_{i,h}} + \frac{\delta_{h+1}}{\delta_h} \frac{\partial R_i}{\partial \Delta_{i,h+1}} \frac{\partial \Delta_{i,h+1}}{\partial q_{i,h}} - \frac{\sum_{j=1}^{24} \left(\frac{\partial C_i}{\partial q_{i,j}} + \frac{\partial R_i}{\partial \Delta_{i,j}} \frac{\partial \Delta_{i,j}}{\partial q_{i,j}} + \frac{\delta_{h+1}}{\delta_h} \frac{\partial R_i}{\partial \Delta_{i,h+1}} \frac{\partial \Delta_{i,h+1}}{\partial q_{i,h}} \right)}{24} \right)^2 \tag{2.4}$$

As is clear from Equation (2.4) above, price variance depends on the marginal costs of production, marginal costs of ramping, and the variance of demand. The intuition behind this result is straightforward, as the intra-day price variance will depend on the convexity of the supply curve and ramping costs. Decreasing the marginal costs will lower price variance because demand intersects a flatter portion of the convex supply curve. Since the point of convexity along the supply curve is dependent on demand, the model also implies a higher variance during periods of higher demand, *ceteris paribus*. Further, demand volatility is fundamentally driving the price volatility so the model suggests that price volatility increases with demand volatility.

To illustrate this point more clearly, consider a basic two period model where demand increases from D_1 to D_2 such that $\Delta = q_2 - q_1 > 0$ is the change in production. This is shown graphically on Figure 2.1. Without ramping costs the supply curve in both periods remains the same, shown as S, and the simple shift from D_1 to D_2 yields the prices equal to marginal production costs, $p_1 = \partial C_1$ and $p_2 = \partial C_2$ for periods 1 and 2, respectively. However, with ramping costs, the equilibrium prices now become $p_1 = \partial C_1 - \partial R$ and $p_2 = \partial C_2 + \partial R$ for periods 1 and 2, respectively.

Intuitively, firms are willing to produce quantities above those at marginal production cost in period 1 in order to have lower ramping costs in period 2. This is shown on Figure 2.1 as a shift from S to S_1 causing a decrease in prices. In period 2 firms produce quantities below marginal production costs because of ramping constraints. This shifts the supply curve to S_2 in Figure 2.1, increasing prices beyond the equilibrium level without ramping costs. Thus, any losses from "over-production" in period 1 are recouped through lower ramping costs in the profit maximizing multi-period equilibrium.

Adding new capacity with lower ramping costs has two effects. First, the

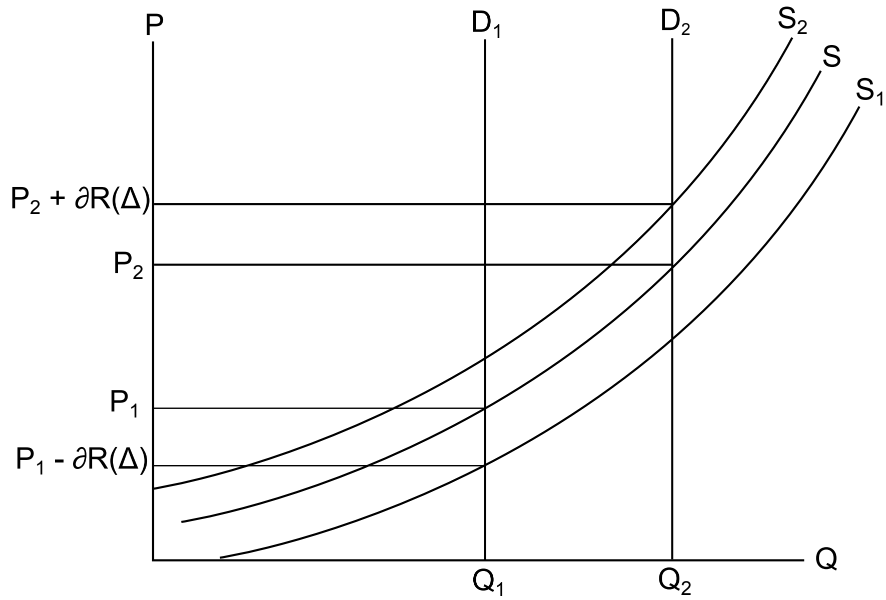

Figure 2.1: Supply and Demand Curves with Ramping Costs

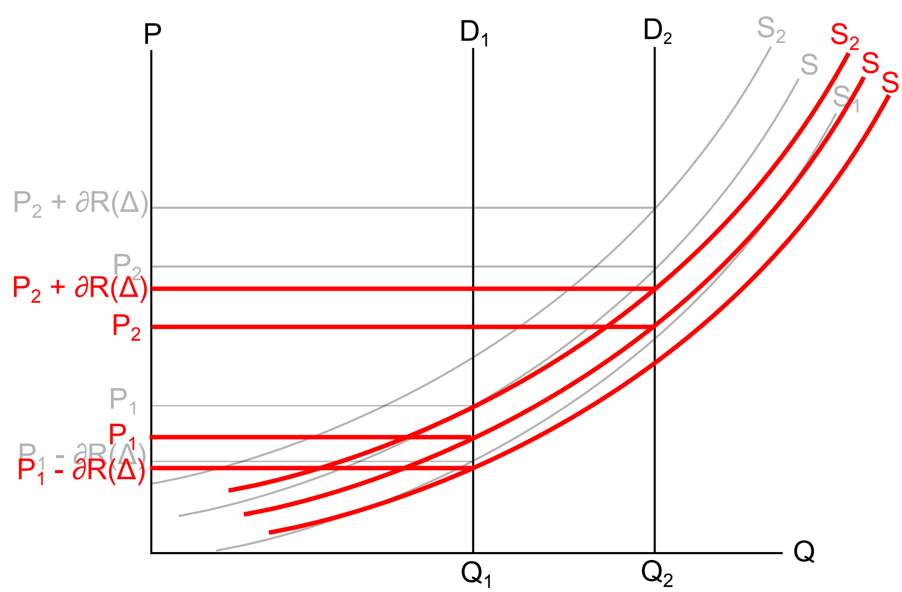

Figure 2.2: Shifting Supply with New Capacity

supply curve shifts outward, which will decrease the difference between p_1 and p_2 because the respective demands now intersect a flatter part of the supply curve. Second, the lower ramping costs squeezes S_1 and S_2 closer to each other, which again decreases the price difference between periods. This is shown graphically in Figure 2.2, where the new equilibrium is shown in red, and the old equilibrium from Figure 2.1 is left in light gray for comparison. Thus, the variance in prices unambiguously decreases from adding new capacity with lower ramping costs and lower marginal production costs.

As discussed in the previous sections, natural gas occupies a critical point along the supply curve where it is the marginal generating unit. Thus, there are two effects from adding new natural gas capacity as captured by the model. First, adding additional new natural gas capacity will lower total marginal costs because the new technologies are assumed to be slightly more efficient than current marginal units. I define this as the "supply shift" effect on volatility. This assumption is validated empirically by the decreasing average heat-rate in natural gas units over the last decade (EIA, 2013*c*). The second effect from adding new natural gas capacity, as captured by the model, is decreasing ramping costs. I define the "ramping effect" as the resulting volatility reduction from a decrease in ramping cost associated with new natural gas capacity. Again, this assumption is justified by empirical analysis (Reguant, 2014, Wolak, 2007), as natural gas units have lower ramping costs than coal-fired power plants. Thus, adding natural gas capacity should unambiguously decrease price volatility, *ceteris paribus*.

Meanwhile, nuclear capacity additions should provide only the supply shift effect because it provides baseload power on the far left portion of the supply curve. As previously noted, nuclear technology has a very low marginal cost and generally operates throughout the day without ramping. Thus, the theoretical model implies that changes in active nuclear capacity should reduce price volatility, but less than the volatility reductions from natural gas. The difference between the volatility reductions from these two generator types is interpreted as the ramping effect. This fundamental result of the model is tested in Section 2.3, and explains how production flexibility stabilizes non-storable commodity prices similar to how storage ability stabilizes traditional commodity prices.[1]

2.3 Econometric Specification

To test the implications and conclusion from the theoretical model in Section 2.2.1, I take advantage of high-frequency wholesale electricity price data at the hourly level. Hourly data are collapsed into daily observations which include intra-day price volatility, intra-day demand volatility, and daily average demand. The theoretical model from Section 2.2.1 implies a reduced form econometric specification where the intra-day price volatility is a function of natural gas capacity, intra-day demand volatility, daily average demand, and unobservable time trends.

[1] In storable commodity markets, production can remain constant at the average demand, since excess supply can be stored and sold in a later period. This means that ramping costs and demand volatility can be pushed to zero because the residual demand across different periods are pushed to their aggregate mean. Thus, prices are stabilized at their marginal production costs.

Thus, the model is:

$$v_t = \beta_0 + \beta_1 NGC_t + \beta_2 S_t + \beta_3 D_t + \beta_4 T_t + \varepsilon_t \tag{2.5}$$

where v_t is the intra-day price volatility (as measured through intra-day standard deviation) on day t, NGC_t is total natural gas capacity, S_t is intra-day demand volatility, D_t is mean demand, T_t is a vector of unobservable time fixed effects, and ε_t is a serially correlated error term such that $\varepsilon_t = \rho\varepsilon_{t-1} + u_t$ where u_t is random noise. The vector of unobservable time fixed effects T_t includes month fixed effects and day-of-week fixed effects to capture additional unobservable seasonality that is not captured by daily demand. It also includes a linear time trend variable, as well as year fixed effects to capture non-linear time trends. Both mean demand and intra-day demand volatility are assumed to be exogenous to price and intra-day price variance because of the focus on the wholesale market. As discussed in Section 3.2, retail residential consumers face no price pressures in the short term from the wholesale market because they are billed on a monthly level using a regulated rate instead of the average wholesale market rate. Instead, the primary drivers of daily demand are weather, season, and hour-of-day.

Due to the stepwise increases in capacity from new additions, the preferred specification is a pooled event study using the model above. In this specification, each natural gas capacity change is accompanied by a separate event window fixed effect in an ordinary least squares (OLS) regression. The event window chosen for this analysis includes one month before and after the capacity change, and assumes the exact date of the capacity change is exogenous within this small window. This arguably alleviates endogeneity concerns surrounding natural gas capacity and price, which arise since natural gas units are usually the marginal generating unit and typically determines the marginal price of electricity in the wholesale market. Thus, it is likely that running a simple OLS regression without the event window fixed effects is inadequate because natural gas capacity is endogenous with electricity price and intra-day price variance.

However, if the assumption that capacity comes online exogenously within the event window is not valid, I also provide an instrumental variables approach using a two-stage least squares (2SLS) regression. In this alternative approach, I instrument for natural gas capacity using a 31-day rolling average of the spark spread, lagged by 24 months. The spark spread is the gross margin between electricity price and the cost of generation using natural gas. More specifically,

$$SS_t = \sum_{i=0}^{30} \frac{1}{31}(p_{t-i} - NGP_{t-i} * HEAT_{t-i}) \tag{2.6}$$

where SS_t is the 31-day rolling average spark spread ($USD/MWh) on day t, p_t is the daily average electricity spot price ($USD/MWh), NGP_t is the natural gas price ($USD/MMBtu), and $HEAT_t$ is the heat rate (MMBtu/MWh) which measures how efficiently a natural gas generator can convert gas into electricity. The spark spread gives a measure of the profitability of generating electricity from natural gas and is highly relevant for investment decisions surrounding natural gas capacity. Further, a lagged spark spread is used as an instrument because it is intuitively correlated with future natural gas capacity, but is exogenous with respect to current prices. While some persistence in the spark spread may cause autocorrelation to remain at short intervals, at longer intervals this is shown to not be the case. Thus, a 24-month lag is used in the model. The long lag is due to a natural gas construction time of 18-36 months and should pass the exclusion restriction which requires the instrument to only influence current electricity prices through natural gas capacity.

Finally, a third specification is provided using a generalized autoregressive conditional heteroskedasticity (GARCH) model (Bollerslev, 1987, Engle, 1982), which is sometimes used in the literature on electricity prices and volatilities (Hadsell, 2007, Hadsell et al., 2004, Worthington et al., 2005). In brief, the conditional intra-day volatility estimated by the GARCH model is

$$p_t = \phi + \varepsilon_t \tag{2.7}$$

$$v_t = \beta_0 + \beta_1 v_{t-1} + \beta_2 \varepsilon_{t-1} + \beta_3 NGC_t + \beta_4 S_t + \beta_5 D_t + \beta_6 T_t \tag{2.8}$$

where v_t is the intra-day price volatility (as measured through intra-day standard deviation) on day t, such that v_{t-1} represents the previous period's volatility forecast. Meanwhile, ε_{t-1} is a lagged error term representing new information about volatility from the previous period. Similar to the prior specification, NGC_t is total natural gas capacity, S_t is intra-day demand volatility, D_t is mean demand, T_t is a vector of unobservable time fixed effects, p_t is electricity price, and ϕ is mean electricity price. The GARCH model also requires the dependent variable to be generated by a stationary process, so an augmented Dickey-Fuller test is performed. I reject the null hypothesis that intra-day price volatility contains a unit root with a z-statistic of -29.00 and I reject that mean price contains a unit root with a z-statistic of -14.48. Thus, the additional requirements to use the GARCH model are satisfied by my price data during my study period.

2.4 New England ISO Market Background

Prior to the 1990s, New England's electricity market was comprised of vertically integrated monopolies that were heavily regulated. Private and municipal utilities managed the region's electricity grid through the New England Power Pool (NEPOOL) created in the early 1970s. However, by 1996 the Federal Energy Regulatory Commission (FERC) issued orders that encouraged wholesale electricity markets for competitive electricity generation. The FERC created general guidelines with a recommended market structure where a non-profit Regional Transmission Organization (RTO) is entrusted to manage the transmission grid and electricity markets. This paved the way for the creation of the Independent Systems Operator of New England (ISO-NE) in 1997 to oversee the market restructuring, ensure grid reliability, and establish competitive markets. (ISO-NE, 2014*a*)

New England's competitive electricity markets were first implemented in 1999 and now cover 14 million people across six states.[2] The wholesale market includes over 500 participants and the ISO-NE coordinates over 8,000 miles of transmission

[2]The New England market includes Maine, Vermont, New Hampshire, Massachusetts, Connecticut, and Rhode Island.

lines (ISO-NE, 2014*c*). After restructuring, consumers can choose between several licensed utilities which are responsible for the retail delivery of electricity. Typically residential consumers pay a constant marginal cost for electricity at a rate fixed for several months and face no hourly price pressure from the wholesale market. Thus, consistent with the prior literature, the rest of this analysis assumes demand to be exogenous to wholesale prices at the hourly level.[3]

Major changes to the wholesale market occurred in 2003 when the ISO-NE adopted the "Standard Market Design" of FERC, which established locational marginal pricing,[4] financial transmission rights,[5] and a dual-settlement market. The dual-settlement market system provides a day-ahead market and a real-time market, which clear separately through two competitive auctions. (ISO-NE, 2014*b*)

In the day-ahead market, participants provide hourly bids for the supply and demand[6] of electricity that will be dispatched the following day. For each hour of scheduled delivery, the bids are due by noon of the prior day. ISO-NE then stacks the bids into hourly aggregate supply and demand curves and schedules electricity

[3]At longer time horizons, changes in wholesale electricity prices are eventually passed on to the consumer but the exogeneity assumption is arguably most appropriate for the frequency of the data used in this analysis.

[4]Locational marginal pricing (LMP) is required for efficient markets because of transmission capacity constraints which impose congestion costs. For each node and load zone in the ISO-NE, supply and demand offers are submitted such that the LMP provides the competitive price inclusive of congestion costs. If congestion and transmission losses are zero, the efficient price is equivalent across all nodes and their zonal aggregates.

[5]Since LMP includes congestion costs paid to the ISO-NE by power purchasers, the suppliers may receive less revenue than the final price that includes congestion costs. Thus, financial transmission rights (FTR) are auctioned to market participants, giving them a share of the real-time congestion payments that are absent from the day-ahead market price. For power purchasers, this acts as a hedge against unexpected higher congestion costs, while it can also provide additional revenue for generators or speculators.

[6]While demand is exogenously determined by retail customers, retail utilities have a choice to buy electricity in the day-ahead market or the real-time market. Any unscheduled electricity demanded in the day-ahead market is required to be purchased in the real-time market.

to be delivered for all bidders below the intersection of supply and demand. While the day-ahead market is purely financial since no electricity is physically delivered, suppliers must deliver the agreed amount of electricity in the corresponding hour of the following day. In the event of equipment malfunction, for example, the supplier cannot deliver the ex-ante scheduled amount of power and they are required to buy the appropriate amount in the real-time market. (ISO-NE, 2014*b*)

After the first round of commitment in the day-ahead market, ISO-NE performs a reliability assessment based on its own demand forecast and a "re-offer" period begins. Supply and demand that has not been previously scheduled is eligible for bidding in this market, which forms the foundation of the real-time market. Throughout the following trading day the ISO-NE physically balances supply and demand through these hourly bids while maintaining grid stability through a sufficient operating reserve of electricity. The real-time market prices are from ex-post settlements based on actual power delivery that may deviate from expected demand. (ISO-NE, 2014*b*)

Although the day-ahead market is purely financial, risk averse market participants may prefer the day-ahead schedule. The day-ahead pricing is typically more stable because it is based on expected outcomes, but real demand variations can be unexpected. To ensure the convergence of day-ahead prices with real-time prices, the ISO-NE also allows "virtual bids", which are purely financial trades in the day-ahead market that must be closed out in the real-time market. Thus, any consistent and profitable arbitrage opportunities between the two markets should be removed in the presence of virtual bidding by risk-neutral participants, leaving only a small risk premium. With risk averse participants, the forward premium should be significantly driven by the variance and skewness of spot market prices (Bessembinder and Lemmon, 2002), as previously discussed.

Overall, the New England market is primarily served by electricity generation from nuclear and natural gas. The total GWh generation by source is provided by the ISO-NE and shown in Table 2.1 for 2005-2011, the entire period studied in this analysis. In 2011, generation from nuclear and natural gas facilities comprised around 67% of total generation, not including the 13% from dual-fuel generators, much of which can be attributed to natural gas as well. Meanwhile, coal, hydro, and aggregate non-hydro renewables[7] each generate close to 6% of the ISO-NE total. Thus, this analysis focuses on the two largest generator types of nuclear and natural gas to understand the role of ramping costs in price volatility. Generally, natural gas generators are the marginal unit throughout most of the year, while new nuclear has been discussed as a hedge against the fossil-fuel price volatility that underlies electricity price risk (Kessides, 2010, Roques et al., 2006).

While the ISO-NE wholesale electricity market generally operates independently, there are also thirteen interconnections that allow for the purchase and sale of electricity to grids in New York and Canada. The annual flows of electricity from 2005-2011 are listed for the ISO-NE in Table 2.2. On average, net imports account for 5.7% of electricity consumed within the ISO-NE. The ISO-NE is a net exporter of electricity to the New York ISO, but a net importer from Quebec. From 2005 to

[7] Within non-hydro renewable generation for 2011, 4.9% of total generation is from wood and refuse, 0.6% from wind, and less than 0.3% from landfill gas or solar.

Table 2.1: New England Generation Profile: Annual GWh from 2005-2011

Source	2011	2010	2009	2008	2007	2006	2005
Total Generation	120,610	126,416	119,437	124,749	130,723	128,050	131,877
	100%	100%	100%	100%	100%	100%	100%
Gas	46,378	42,042	38,163	38,338	39,367	39,425	38,583
	38.45%	33.26%	31.95%	30.73%	30.11%	30.79%	29.26%
Nuclear	34,283	38,364	36,231	35,547	36,972	36,923	34,609
	28.42%	30.35%	30.33%	28.49%	28.28%	28.83%	26.24%
Oil/Gas†	15,925	15,542	12,487	12,721	15,791	13,542	16567
	13.2%	12.29%	10.45%	10.2%	12.08%	10.58%	12.56%
Hydro	8,252	7,227	8,354	8,466	6,385	7,498	6,739
	6.84%	5.72%	6.99%	6.79%	4.88%	5.86%	5.11%
Renewables	7,261	7,686	7,331	7,539	7,818	7,675	7,599
	6.02%	6.08%	6.14%	6.04%	5.98%	5.99%	5.76%
Coal	7,080	14,131	14,558	18,596	19,770	19,375	20,789
	5.87%	11.18%	12.19%	14.91%	15.12%	15.13%	15.76%
Pumped Hydro	1,149	854	1,419	1,623	1,744	1,582	1,339
	0.95%	0.68%	1.19%	1.3%	1.33%	1.24%	1.02%
Oil	282	570	895	1,918	2,877	2,030	5,652
	0.23%	0.45%	0.75%	1.54%	2.2%	1.59%	4.29%

†ISO-NE does not have data splitting generation by fuel in dual-fuel units.

Table 2.2: New England Electricity Flow: Annual GWh from 2005-2011

	2011	2010	2009	2008	2007	2006	2005
Total Demand	129,163	130,773	126,838	131,753	134,466	132,087	136,355
	100%	100%	100%	100%	100%	100%	100%
Total Generation	120,610	126,416	119,437	124,749	130,723	128,050	131,877
	93.38%	96.67%	94.16%	94.68%	97.22%	96.94%	96.72%
Pumped Hydro†	-1,589	-1,183	-1,963	-2,247	-2,403	-2,156	-1,819
	-1.23%	-0.9%	-1.55%	-1.71%	-1.79%	-1.63%	-1.33%
Imports	15,880	12,781	15,226	14,256	12,269	10,762	10,152
Exports	5,738	7,242	5,863	5,005	6,122	4,569	3,855
Net Imports	10,142	5,539	9,363	9,251	6,146	6,193	6,297
	7.85%	4.24%	7.38%	7.02%	4.57%	4.69%	4.62%

†Pumped hydro is a net loss of energy generation but can still occasionally be optimal. Essentially it provides relatively small indirect storage of electricity during low demand periods that is released during peak demand periods.

2011, demand has decreased by 5.3% while total generation has decreased by 8.5%. The difference is made up though additional imports which have generally increased over time.

2.4.1 Data

To test the role of natural gas capacity in the price stability of the wholesale electricity market, I use data from the Independent Systems Operator of New England (ISO-NE). Hourly electricity prices from the real-time ISO-NE market are obtained from March 2005 through June 2011. Throughout the analysis, prices and electricity demand loads are taken from the Southeast Massachusetts (SEMASS) zone, as it is geographically central to the ISO-NE. The data for both price and demand load are collapsed at the daily level to provide intra-day volatility for the 24-hour period.

Although "volatility" is colloquially used to imply "variability," for clarity I define volatility as the standard deviation of the data.[8] More formally:

$$\sigma_t^x = \sqrt{\frac{1}{24}\sum_{h=1}^{24}(x_{t,h} - \mu_t)^2} \tag{2.9}$$

where σ_t^x is intra-day volatility for the variable x on day t, h is the hour of day, and μ is the daily average of x. Thus, throughout the remainder of the analysis I use the terms "volatility" and "standard deviation" interchangeably.

[8] This is also sometimes referred to as "historical volatility" in the finance literature, which is distinct from annualized volatility, implied volatility, variance, and the probability of extreme events.

Monthly summary statistics are shown in Table 2.3 for daily mean price, intra-day price volatility, daily mean demand, and intra-day demand volatility. The summary statistics are consistent with previous expectations about the New England electricity market, with the summer and winter months showing higher intra-day volatilities in addition to higher mean prices, mean-demands, and intra-day demand volatilities. The summary statistics suggest a strong seasonality to all variables of interest, which will be important to capture through month fixed effects.

Figure 2.3 shows a clear relationship between intra-day price volatility and intra-day demand volatility. The graph uses a 60-day smoothing average to show general time trends without the daily statistical noise. The seasonality of intra-day demand volatility comes through very clearly, with a strong peak during the summer months and a second, smaller peak during early winter. An overall linear time trend is less obvious for either price or demand volatility, but there may be a slight decrease in both intra-day volatilities over time. Generally, periods of high demand volatility appear to coincide with high price volatility, a finding consistent with the intuition of the theoretical model in Section 2.2.1.

Figure 2.4 shows a similar trend, again with a clear seasonality for both daily mean demand and intra-day price volatility. The second peak during early winter is more pronounced in the mean demand graph than in the intra-day demand graph, but the two graphs are generally consistent with each other. As implied by the basic and intuitive theoretical model, the temporal patterns of volatility and mean demand are highly correlated.

Data on natural gas generator heat rates and Massachusetts gas price are

Table 2.3: Summary Statistics for ISO-NE (March 2005 through June 2011)

		Mean and Std. Dev. (in parentheses)			
Month	Obs (n)	Daily Mean Price ($USD/MWh)	Intra-Day Price Volatility ($USD/MWh)	Daily Mean Demand (MWh)	Intra-Day Demand Volatility (MWh)
January	186	69.07	18.92	1,791.5	278.4
		(21.42)	(10.35)	(107.0)	(33.6)
February	169	62.99	16.02	1,764.2	254.7
		(16.81)	(8.13)	(108.1)	(33.5)
March	217	55.76	13.30	1,667.7	246.6
		(16.74)	(7.58)	(118.4)	(43.0)
April	210	56.43	12.75	1,532.0	239.7
		(20.71)	(8.25)	(91.3)	(37.3)
May	217	58.83	15.71	1,552.9	267.9
		(23.99)	(11.42)	(113.7)	(47.4)
June	210	58.66	16.65	1,813.6	354.0
		(25.37)	(12.58)	(246.3)	(91.1)
July	186	64.70	18.39	2,100.6	425.8
		(26.90)	(11.84)	(284.2)	(99.7)
August	186	63.83	19.81	2,048.3	409.8
		(27.30)	(30.82)	(284.4)	(100.6)
September	180	58.63	16.11	1,722.6	319.6
		(24.94)	(11.69)	(197.3)	(67.6)
October	186	59.51	15.38	1,586.6	278.8
		(26.68)	(13.38)	(100.3)	(36.6)
November	180	55.85	14.80	1,623.5	284.0
		(15.02)	(8.23)	(86.6)	(32.2)
December	186	71.17	18.49	1,794.0	303.3
		(24.15)	(9.11)	(116.4)	(38.2)
Total Sample	2313	61.11	16.28	1744.7	304.1
		(23.34)	(13.42)	(242.2)	(84.5)

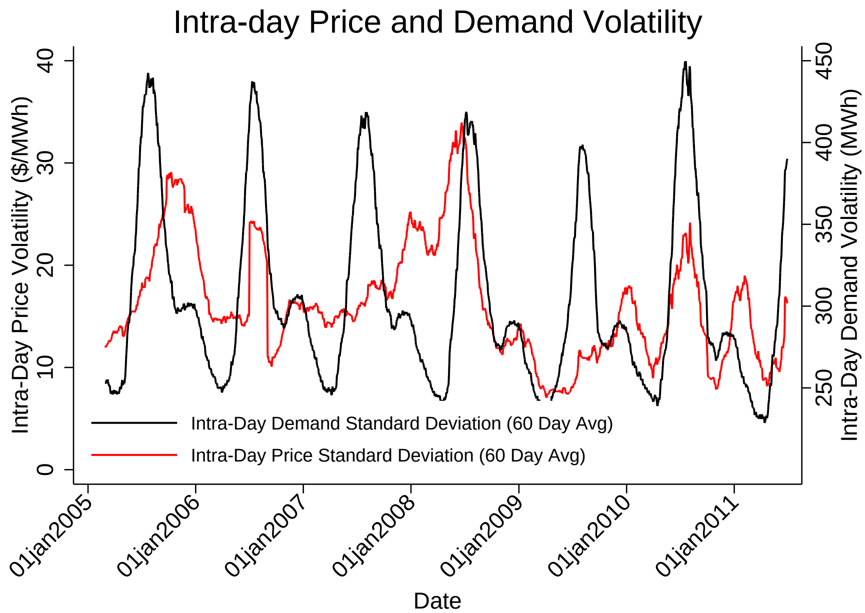

Figure 2.3: Intra-day Price and Demand Volatility

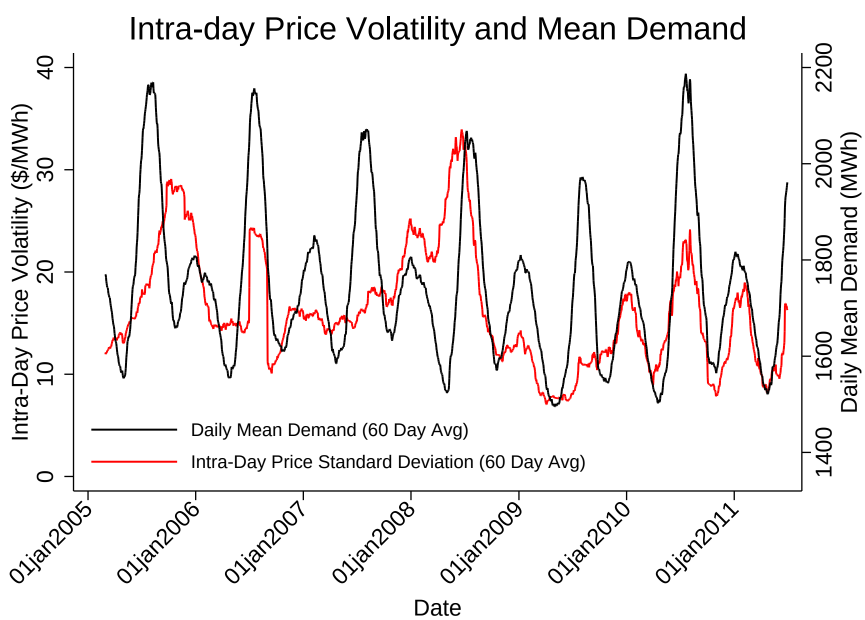

Figure 2.4: Intra-day Price Volatility and Mean Demand

taken directly from the United States Energy Information Agency (EIA). Since heat rate data is provided by the EIA only at annual averages through their "Electric Power Annual Report" (EIA, 2013*c*), a monthly rolling average is constructed which assumes linear technological improvements within the year. The EIA also provides monthly average natural gas prices paid by Massachusetts power plants using data from their "Monthly Cost and Quality of Fuels for Electric Plants Report" (form EIA-423) and "Power Plant Operations Report" (form EIA-923). The monthly data is then used to construct the marginal cost of electricity from natural gas, without considering operational expenses. Finally, a daily spark spread is constructed as the difference between the daily average electricity spot prices within the SEMASS zone and the marginal cost of electricity from natural gas, as described in Section 2.3.

Summary statistics for all variables used to construct the spark spread are shown in Table 2.4. As expected, the average heat rate improves over time from 9,207 Btu/kWh in 2003 to 8,159 Btu/kWh in 2009. Note that the heat rate data covers from March 2003 through June 2009, although the primary period of this analysis is from March 2005 through June 2011. This is because of the 24-month lagged spark spread used as the instrumental variable for natural gas capacity. Thus, the data from March 2003 through February 2005 is only used to calculate the instrumental variable and is not used as the dependent variable in the primary regression results of Section 2.5.

The natural gas price paid by Massachusetts power plants during this period is \$7.9 per thousand cubic feet. This is expected, although it is slightly above the United States average of \$7.19 paid by power plants from March 2003 through June

Table 2.4: Instrumental Variable Construction (March 2003 through June 2009)

Variable	Mean	Std. Dev.	Min	Max
Heat Rate (Btu/kWh)	8548	290.5	8159	9207
MA Gas Price ($USD/1000 ft^3)	7.924	2.378	4.30	14.76
Electricity Cost from Gas ($USD/MWh)	65.62	19.01	34.23	122.39
Daily Average Price ($USD/MWh)	62.00	22.24	22.48	277.80
Spark Spread ($USD/MWh)	-3.62	14.70	-62.07	210.73

2009. After calculating the marginal cost of electricity from the natural gas prices and the EIA average heat rates, the daily average is $65.62 per megawatt-hour. As expected, this is very close to the mean spot price during this period ($62/MWh) because natural gas generators are typically the marginal generator and thus set the electricity price. The difference between these leads to a small average spark spread of -$3.62/MWh.

While a trivial average spark spread is expected it is also important to note the large variation. During the sample period, the daily average spark spread runs from -$62/MWh to $211/MWh. Further, many natural gas generators are "load following units" meaning that they ramp up generation to follow the increased demand during peak hours of the day when prices and demand are highest. The relatively low ramping costs of natural gas units means they can selectively operate during profitable hours. Thus, it is certainly possible to make a profit using natural gas generators even though the small negative daily average spark spread initially suggests otherwise. Further, the 31-day rolling average spark spread that is used as an instrument will smooth away from daily noise and remains a good measure of overall profitability for natural gas units. If the spark spread average remains high for some time, the increased profitability will induce additional entrants to build capacity. Thus, a positive spread should encourage new investment in natural gas capacity.

Data on natural gas capacity is gathered from the EIA's "Annual Electric Generator Report" (form EIA-860). The dataset includes generator level data for power plants in the United States and includes the state of operation, nameplate capacity, date placed in service, and date retired when it applies. Generator level data is collected for all six states within the ISO-NE and changes in natural gas capacity are constructed for 2005-2011 using installation and retirement dates. During this period total natural gas capacity in the EIA database increased by 730.1 MW, which amounts to just over 6% of installed natural gas capacity in 2010 (FERC, 2010). The additions came through nineteen new generators, with an average capacity of 60 MW each. These additions happened through thirteen new power plants, with an average capacity of 87 MW each. Further variations in total capacity come from the nine natural gas generator retirements, with an average capacity of 45 MW each. These capacity reductions happened through the closure of seven power plants, with an average capacity of 58 MW each.

As a visual example of a capacity event, Figure 2.6 graphs intra-day volatility over time, where the x-axis shows the number of days from the day of the capacity addition. The event shown is a 58 MW natural gas generator addition, which was chosen because it represents the average size of a capacity change from 2005-2011. As the graph demonstrates, intra-day volatility is generally noisy, although volatility does appear to decrease in the period following the capacity change. The data from all capacity additions are plotted for the corresponding two weeks before and after

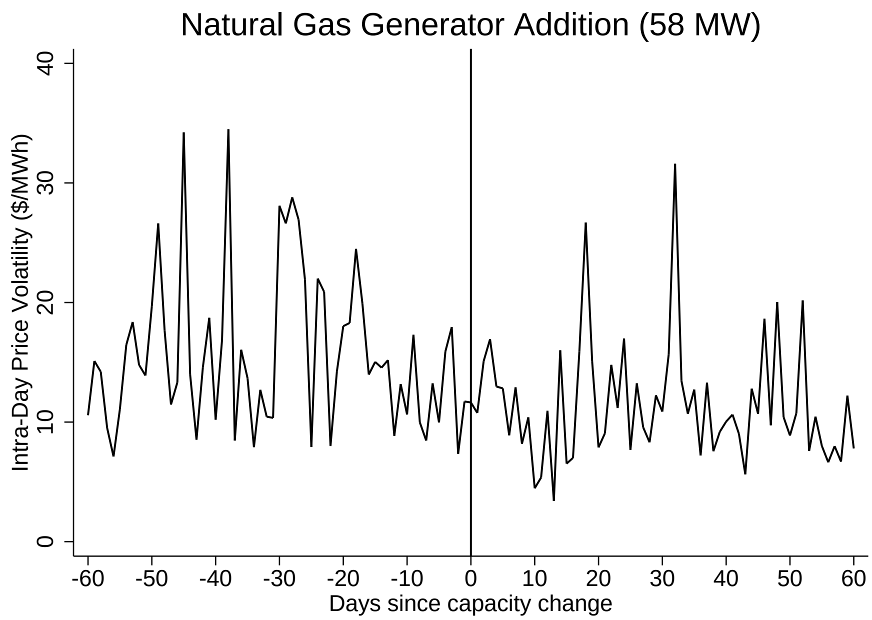

Figure 2.5: Example of Natural Gas Generator Addition

the capacity change in Figure 2.6. Since intra-day volatility is generally noisy, a lowess curve is also fitted to the pre-period and post-period. While the lowess curve omits the control variables from the econometric model, it does show a small and distinct discontinuity. The resulting decrease in intra-day volatility from the capacity addition appears to persist in time with no major change to the lowess curve trend.

While no new nuclear capacity has been installed or retired during the period studied, nuclear capacity occasionally goes offline for both planned and unplanned outages related to refueling, maintenance, and safety. Planned outages are typically scheduled months in advance and occur during regular refueling times. As such, the exact outage date is arguably exogenous with respect to the current intra-day price volatility, but the data is also analyzed using unplanned "forced outages" with no

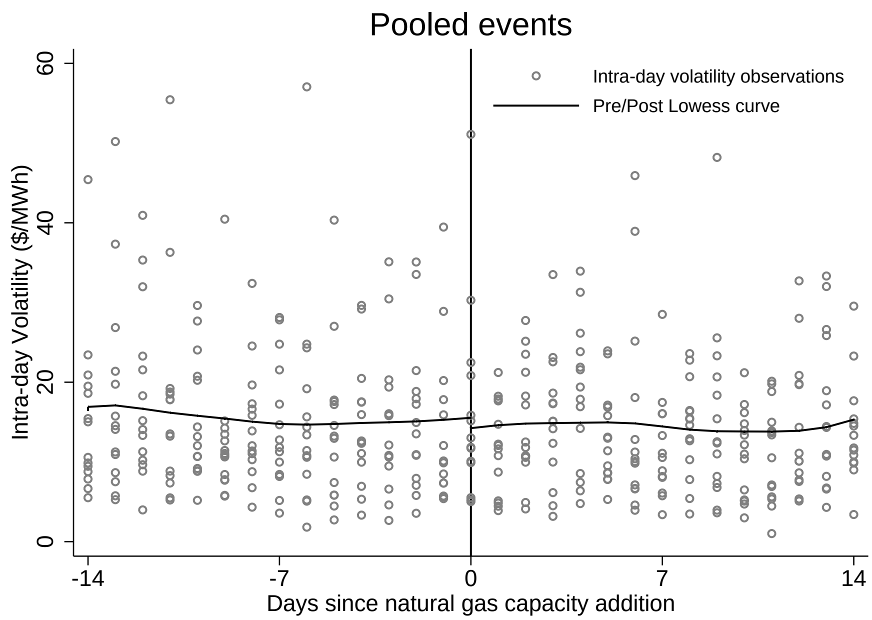

Figure 2.6: Lowess Curve of Natural Gas Generator Additions

change to the results discussed in Section 2.5. Data on nuclear capacity outages within the ISO-NE comes directly from the US Nuclear Regulatory Commission's "Power Reactor Status Report." There are five active nuclear generators within the four nuclear power plants located inside the ISO-NE load area.[9] The generators have an average capacity of 917 MW per generator, for a total installed nuclear capacity of 4,586 MW. During the sample period, the average active installed capacity is 4,217 MW, such that active capacity was below installed capacity for 391 total days, or 17% of the sample. Included among these are 185 days from forced outages, or 8% of the total sample days. Since there are overlapping outages, perhaps a more insightful statistic during the sample period is an average outage time of 21.9 days per nuclear generator per year.

[9]The four power plants are Millstone Nuclear Power Station in Connecticut, Pilgrim Nuclear Generating Station in Massachusetts, Seabrook Nuclear Power Plant in New Hampshire, and Vermont Yankee Nuclear Power Plant in Vermont.

2.5 Results

The regression results show that natural gas capacity significantly decreases intra-day price volatility in the wholesale electricity market, supporting the theoretical model in Section 2.2.1. Table 2.5 provides the coefficients of interest for each of the three primary specifications, with Newey-West standard errors reported when applicable to correct for serial correlation. Column (A) is the preferred specification, using the pooled event study approach. Column (B) gives the second stage results for the two-stage least squares (2SLS) specification. Finally, Column (C) provides the regression results for the generalized autoregressive conditional heteroskedasticity (GARCH) model. As discussed in Section 2.3, each specification controls for intra-day demand volatility, daily demand means, month fixed effects, year fixed effects, day-of-week fixed effects, and linear time trends.

The results across all three specifications continually show natural gas capacity leading to a significant decrease in price volatility. The coefficient in Column (A) suggests that each additional MW of natural gas capacity decreases price volatility by \$0.010/MWh. The average generator addition during my sample period is 60 MW, so the results suggest that a typical generator addition decreases intra-day price volatility by about 4%, or \$0.62/MWh. This volatility decrease is approximately 1% of the mean electricity price during the sample period. The 2SLS results in Column (B) show an increase in the magnitude of natural gas coefficient to -0.028, thought this is not statistically different from the pooled event study regression in Column (A). Using the 2SLS coefficient instead suggests that adding an additional natural

Table 2.5: Regression Results: Natural Gas Capacity
Dependent Variable: Intra-day Price Volatility ($/MWh)

	(A) OLS	(B) 2SLS	(C) GARCH
Natural Gas Capacity (MW)	-0.0103**	-0.0278**	-0.0099***
	(0.0051)	(0.0109)	(0.0025)
Demand Volatility (MWh)	0.0622***	0.0629***	0.0552***
	(0.0071)	(0.0068)	(0.0037)
Demand Mean (MWh)	0.0159**	0.0161***	0.0149***
	(0.0067)	(0.0026)	(0.0015)
Time Fixed Effects	Yes	Yes	Yes
Observations	2313	2313	2313

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively. Newey-West standard errors are reported in parenthesis for Columns (A) and (B) to correct for serial correlation.

gas generator decreases intra-day price volatility by about 10%, or $1.66/MWh. Finally, the GARCH model gives a coefficient of -0.011 in Column (C), insignificantly different from either of the two previous specifications.

All three specifications continually show that intra-day price volatility significantly increases with both intra-demand volatility and mean demand, after considering seasonality and time trends. This is again consistent with the expectations of the theoretical model. Further, the marginal effect of an increase in demand volatility has a much larger effect than an increase in daily mean, as is intuitively expected since intra-day demand volatility is a main driver of the intra-day price volatility. These coefficients are especially interesting because they can be used to recover the volatility impacts from renewable generation. Renewable generation is typically nondispatchable, and thus can be thought of as reductions in the net demand satisfied by conventional generation. To assess the robustness of the 2SLS results in Column (B) of Table 2.5, variations on the spark spread are used and the first stage

Table 2.6: First Stage 2SLS Results
Dependent Variable: Natural Gas Capacity (MW)

	(A)	(B)	(C)
Lagged Spark Spread ($/MWh)	4.526***	4.624***	4.219***
	(0.349)	(0.403)	(0.432)
Time Fixed Effects	Yes	Yes	Yes
Kleibergen-Paap rk-statistic	167.87	131.21	95.14
Observations	2313	2313	2313

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively. Newey-West standard errors are reported in parenthesis to correct for serial correlation.

results are reported in Table 2.6. Columns (A), (B), and (C) coincide with the same columns of the second stage in Table 2.7. As expected, the lagged spark spread is strongly correlated with increases in natural gas capacity. In Column (A), the spark spread used is lagged two years and is a 31-day rolling average as in Column (B) of Table 2.5. I also perform a weak instrument test using the rk-statistic of Kleibergen and Paap (2006) because the F-statistic of Cragg and Donald (1993) is not valid when the standard errors are not *i.i.d.* normal. Previous literature suggests a rule of thumb where there is little concern of a weak instrument with an F-statistic above 8.96 (Stock et al., 2002, Stock and Yogo, 2001). The preferred specification in Column (A) of Table 2.6 shows that the lagged spark spread is arguably a very strong instrument, with a Kleibergen-Paap rk-statistic of 167.87.

The regression results shown in Columns (B) and (C) use a 60-day and 90-day rolling average for the lagged spark spread instead of a 31-day average. The results are not particularly sensitive to the number of days included in the rolling average of the spark spread. The marginal effects are not statistically different from Column (A), although they do increase. The first stage results in Table 2.6 show that the instrument remains strong and yields no significant change in magnitude.

Table 2.7: Second Stage 2SLS Results
Dependent Variable: Intra-day Price Volatility ($/MWh)

	(A)	(B)	(C)
Natural Gas Capacity (MW)	-0.0278**	-0.0380***	-0.0496***
	(0.0109)	(0.0123)	(0.0146)
Demand Volatility (MWh)	0.0629***	0.0630***	0.0632***
	(0.0068)	(0.0069)	(0.0070)
Demand Mean (MWh)	0.0161***	0.0159***	0.0157***
	(0.0026)	(0.0026)	(0.0027)
Time Fixed Effects	Yes	Yes	Yes
Observations	2313	2313	2313

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively. Newey-West standard errors are reported in parenthesis to correct for serial correlation.

As discussed in Section 2.2.1, there are two effects of adding natural gas capacity. First, is the outward shift in the supply curve which should yield a decrease in intra-day price variance because demand intersects on a flatter convexity. The second effect is the decrease in ramping costs which squeezes together the dynamic supply curve shifts, which also yields a decrease in intra-day price variance. The regression above captures both of these effects, but the ramping costs effect is of particular interest to this paper. It is arguably possible to separate out these two effects using capacity changes that only affect volatility through outward supply curve shifts, for example using nuclear outages. Since nuclear power is a low marginal cost provider of baseload power and does not typically ramp production during the day, it seems reasonable to assume that nuclear power outages will only shift the supply curve inward, without changing the intra-day dynamics involved from ramping costs. Thus, running the same specification on nuclear power should show changes in volatility due only to the supply curve shift. As previously noted, no new nuclear capacity has been built during the time period studied but outages do occur for

Table 2.8: Regression Results: Natural Gas and Nuclear Capacity
Dependent Variable: Intra-day Price Volatility ($/MWh)

	OLS (A)	OLS (B)	OLS (C)	2SLS (D)	GARCH (E)
Natural Gas Capacity (MW)			-0.0152**	-0.0347***	-0.0166***
			(0.0059)	(0.0120)	(0.0025)
Nuclear Capacity (MW)	-0.0013***	-0.0014***	-0.0019***	-0.0026***	-0.0019***
	(0.0004)	(0.0005)	(0.0006)	(0.0007)	(0.0002)
Nuclear Capacity X Forced Outage (MW)		0.0001			
		(0.0004)			
Demand Volatility (MWh)	0.0626***	0.0626***	0.0628***	0.0632***	0.0554***
	(0.0072)	(0.0072)	(0.0071)	(0.0068)	(0.0036)
Demand Mean (MWh)	0.0167**	0.0167**	0.0160**	0.0164***	0.0144***
	(0.0071)	(0.0071)	(0.0067)	(0.0026)	(0.0015)
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations	2313	2313	2313	2313	2313

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Newey-West standard errors are reported in parenthesis to correct for serial correlation.

refueling, planned maintenance, and occasional emergency shutdowns. The specifications shown in Table 2.8 use these temporary outages in nuclear capacity to understand the volatility changes from the supply shift.

The OLS results in Column (A) of Table 2.8 show a small but statistically significant decrease to price volatility from nuclear capacity. The marginal effect of an additional MW of nuclear capacity leads to a $0.0013/MWh decrease in intra-day price volatility. Although nuclear outages are generally assumed to be exogenous, Column (B) uses an interaction effect between nuclear capacity and forced outages to ensure that forced outages behave similarly to planned outages. The results show that forced outages have a very small, insignificantly different effect on intra-day price volatility when compared to regular outages.

Columns (C), (D), and (E) of Table 2.8 include natural gas capacity outages in the same regression and can be compared with Columns (A), (B), and (C) of Table 2.5. The pooled event study approach is shown in Column (C) of Table 2.8 while the 2SLS and GARCH models are shown in Columns (D) and (E), respectively. When including the capacities of both nuclear and natural gas power plants, the marginal effect of natural gas and nuclear capacity on price volatility changes insignificantly across specifications.

As discussed above, the discrepancies in the marginal effect between nuclear capacity and natural gas capacity are attributed to ramping costs. The preferred results in Column (C) suggest that adding 60 MW of nuclear capacity decreases intra-day price volatility by 0.7%, or \$0.114/MWh, while adding 60 MW of natural gas capacity decreases intra-day price volatility by 5.6%, or \$0.912/MWh. Thus, empirically it appears that the reduction of volatility from the supply shift effect is actually quite small, although still statistically significant. The bulk of the volatility reduction from adding natural gas generators comes through supply flexibility via decreased ramping costs. The results imply that adding 60 MW of natural gas capacity will decrease intra-day price volatility by 4.9 percentage points, or \$0.798/MWh, more than adding a lower marginal cost inflexible generator. This volatility reduction amounts to approximately 1.3% of the mean electricity price.

The theoretical model also implies the ramping cost effect to be greater during the summer months for two reasons. First, the ramping cost effect is more pronounced because demand intersects a steeper section of the convex supply curve. Second, demand volatility is greater during the summer months which also induces larger dynamic shifting of the supply curves. This notion is tested empirically in

Table 2.9: Regression Results: Natural Gas and Nuclear Capacity
Dependent Variable: Intra-day Price Volatility ($/MWh)

	(A) OLS	(B) 2SLS	(C) GARCH
Natural Gas Capacity (MW)	-0.0105*	-0.0391***	-0.0123***
	(0.0060)	(0.0144)	(0.0025)
Natural Gas Capacity X Summer (MW)	-0.0351***	-0.0225***	-0.0171***
	(0.0050)	(0.0053)	(0.0027)
Nuclear Capacity Capacity (MW)	-0.0019***	-0.0027***	-0.0018***
	(0.0005)	(0.0007)	(0.0002)
Nuclear Capacity Capacity X Summer (MW)	-0.0047*	-0.0049*	0.0009
	(0.0024)	(0.0026)	(0.0012)
Time Fixed Effects	Yes	Yes	Yes
Observations	2313	2313	2313

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Newey-West standard errors are reported in parenthesis to correct for serial correlation.

Table 2.9 using the pooled event study, 2SLS, and GARCH models in Columns (A), (B), and (C), respectively.

Each column of Table 2.9 uses an interaction effect with a dummy variable equal to one during the summer months of June through August, when demand and intra-day demand volatility are highest. Consistent with our expectations, each regression shows that natural gas capacity provides a significantly larger stabilizing effect during the summer. Meanwhile, the supply shift effect shown by the coefficient on nuclear capacity is also larger during the summer months but it is again an order of magnitude below that of natural gas. This effect is only significant in the first two columns. While the size of the coefficients for natural gas and nuclear does change across specification, they all are consistent with the intuition of the theoretical model which implies larger volatility reductions in the summer months due to the ramping cost effect.

Figure 2.7 shows the monthly marginal effect of additional natural gas capacity on intra-day price volatility. Although segmenting the sample by months limits statistical significance because of a smaller sample size, the relative coefficient magnitudes are revealing. There is a clear intra-year trend, with natural gas providing larger reductions to price volatility during months with larger demand volatility. Consistent with the regression in Table 2.9, the decrease in volatility from natural gas capacity in the summer months is several times greater than the rest of the year.

This has important implications for future price behavior in the presence of non-dispatchable wind or solar generation. While wind generation will reduce the residual demand that is supplied by conventional generators, it also has intermittency concerns that may increase the demand volatility served by conventional generators. Solar has similar concerns but the effect is more ambiguous since production follows demand, with larger output during the summer and daylight hours. Thus, the results in Table 2.9 and Figure 2.7 underscore the importance of pairing increases in intermittent renewable generators with conventional generators that have low ramping costs. The results suggest that the value of price stability from natural gas is increasing with the share of non-dispatchable generators such as wind, an idea explored further in Section 2.5.1.

While this analysis yields strong evidence that the low ramping costs of natural gas generators provide ancillary benefits to intra-day price stability, this may not be the case at longer time intervals. At the intra-day level, natural gas generators arguably are not subject to fossil fuel price volatility since gas prices paid by generators are often negotiated through bilateral contracts and forward financial

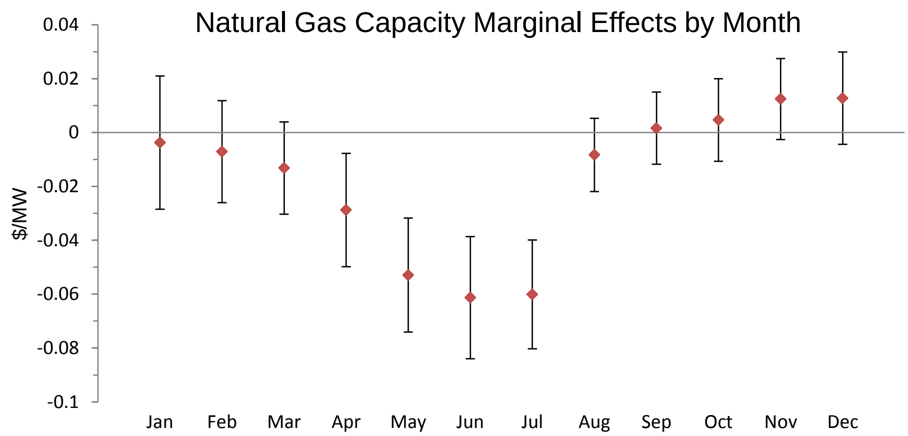

Figure 2.7: Natural Gas Capacity Marginal Effects by Month

markets. However, at longer time horizons natural gas generation is subject to fossil fuel price volatility which means the price stability benefits shown here at the daily level may not translate to volatility reductions at the monthly level.

To investigate this notion, an analysis identical to the baseline specification is performed using the same data at the monthly aggregation. The results of natural gas capacity on intra-month spot price volatility are shown in Table 2.10. The results show that increases in natural gas capacity lead to reductions in intra-month price volatility similar in magnitude to the daily reductions. However, the results are no longer statistically significant. This is arguably the result of both a smaller sample size and larger variations in fossil fuel prices which suppress the price stability benefits shown at the daily level. The results suggest that the ancillary pecuniary benefits from natural gas at the intra-day level do not necessarily generate benefits at longer time horizons. Finally, they emphasize the importance of disaggregated data analysis when investigating electricity markets and related price behavior.

As an additional robustness check to the intra-day price volatility analysis, I

Table 2.10: Regression Results: Intra-month Volatility
Dependent Variable: Intra-month Price Volatility ($/MWh)

	(A) OLS	(B) 2SLS	(C) GARCH
Natural Gas Capacity (MW)	-0.0327	-0.0338	-0.0102
	(0.0213)	(0.0229)	0.0206
Demand Volatility (MWh)	0.1203	0.1517***	0.1048*
	(0.0939)	(0.0563)	(0.0630)
Demand Mean (MWh)	-0.0179	-0.0118	0.0115
	(0.0551)	(0.0256)	(0.0279)
Time Fixed Effects	Yes	Yes	Yes
Observations	76	76	76
Kleibergen-Paap rk-statistic		28.53	

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Newey-West standard errors are reported in parenthesis to correct for serial correlation.

perform the same regression using an alternative measure of volatility that is also used in finance literature focused on electricity prices (Hadsell and Shawky, 2006, Haugom and Ullrich, 2012*a*, Simonsen, 2005, Ullrich, 2012, Zareipour et al., 2007). Here, the historical volatility is defined as the standard deviation of the logarithmic returns:

$$\sigma_t^r = \sqrt{\frac{1}{24}\sum_{h=1}^{24}(r_{t,h} - \overline{r_t})^2} \tag{2.10}$$

where σ_t^r is the intra-day volatility of logarithmic returns on day t and h is the hour of day. Logarithmic returns are defined as

$$r_{t,h} = ln\left(\frac{p_h}{p_{h-1}}\right) \tag{2.11}$$

where p_h is electricity price for hour h on day t.

Over the entire sample period, intra-day standard deviation of logarithmic returns is 0.2047 and the daily mean returns are close to zero, as expected, at -

Table 2.11: Regression Results: Alternative Volatility Measure
Dependent Variable: Standard deviation of logarithmic returns

	(A) OLS	(B) OLS	(C) 2SLS	(D) 2SLS
Natural Gas Capacity (100 MW)	-0.0129***	-0.0126***	-0.0226**	-0.0212**
	(0.0047)	(0.0047)	(0.0094)	(0.0097)
Nuclear Capacity (100 MW)		0.0005		0.0007
		(0.0007)		(0.0008)
Time Fixed Effects	Yes	Yes	Yes	Yes
Observations	2,313	2,313	2,313	2,313
Kleibergen-Paap rk-statistic			150.52	140.52

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Newey-West standard errors are reported in parenthesis to correct for serial correlation.

0.0040. While the intra-day volatility of returns is quite high, it is consistent with the range found in the previous literature which use this measure (Zareipour et al., 2007).

The regression results for both the pooled event study and the 2SLS specification are shown in Table 2.11 all models include mean daily demand, intra-day demand volatility, month fixed effects, day of the week effects, and a linear time trend. The results are similar to the primary results, with natural gas capacity significantly reducing price volatility. Again, the coefficient on natural gas is an order of magnitude above the coefficient for nuclear capacity for the specifications in Columns (B) and (D). Across all specifications the coefficient for natural gas capacity are not significantly different from each other, while the coefficient on nuclear capacity is insignificantly different from zero. A 60 MW natural gas generator addition will decrease volatility by 0.008, or approximately 3.9%. This is only slightly below the results of previous tables when using the traditional definition of volatility which estimates the typical natural gas generator will reduce intra-day volatility by

5.6%.

2.5.1 Volatility Impact of Renewables

To explore how the volatility value of additional natural gas capacity explores over time within the ISO-NE, I provide a simple simulation which is calibrated using the coefficients from the econometric model results. Internal reports by the ISO-NE note that over two-dozen oil and coal fired generators may be retired within the next decade. If these aging generators do retire, ISO-NE notes that over 6,000 MW of new capacity will need to be produced. A significant percentage of the replacement capacity will come from natural gas generators, and up to 40% of proposed projects are from wind generation (ISO-NE, 2013). Using these retirements to motivate the context, the simulation estimates how the volatility value of additional natural gas capacity changes over time.

First, the impact on price volatility from wind generation is recovered from the marginal effects in Section 2.5. Wind generation is a non-dispatchable resource, so additional wind capacity can be modeled as a decrease in the residual electricity demand that is supplied by dispatchable generators. Further, the intermittency issues over wind can be thought of as increases the intra-day volatility in this residual demand. Thus, the theoretical effect of wind on price volatility is ambiguous because of these two competing effects. Since the coefficients for demand and intra-day demand volatility are recovered from exogenous changes in demand, they can be interpreted as the true coefficients from a supply increase of wind without traditional endogeneity concerns between supply and price. A similar method can be used to recover the price volatility impact from solar or other non-dispatchable renewable generators, but it is not done here due to New England's limited solar potential and lack of a comparable region to calibrate the model.[10]

To calibrate the simulation to incorporate the effects of wind generation on price volatility, actual hourly wind generation data is taken from the California Independent Systems Operator (CAISO), within the NP15 zone. This zone covers northern California, which the National Renewable Energy Laboratory estimates to have similar wind potential as the ISO-NE region (NREL, 2014) and the simulation assumes to have the same ratio of wind generation to volatility. Actual hourly wind generation data from ISO-NE is not available for use, but CAISO's NP15 zone data is preferred regardless because the market wind penetration is one percentage point larger than that of ISO-NE and will more accurately reflect the wind volatility under the growth described in the ISO-NE simulation.

The CAISO hourly 2012 wind production data show that an average 0.47 MWh intra-day volatility increase accompanies every MWh decrease in the daily mean residual demand due to wind. This ratio is used in combination with the demand and demand volatility marginal effects to calculate a net increase in price volatility of approximately 3% from 60 MWh of wind energy production. This more accurately reflects the impact of wind generation on price because it is calibrated using wind production, instead of wind capacity, so it already incorporates non-

[10]Recovering true estimates of solar photovoltaic capacity on price volatility will also require use of hourly coefficients, since nondispatchable solar production generally follows the demand load changes throughout the day and smooths the net demand served by traditional generation.

dispatchability concerns such as intermittency and curtailment.

The 3% marginal increase in price volatility due to wind production can be directly compared to the 3.9% decrease in price volatility estimated from the equivalent natural gas production capacity. Their similar magnitudes emphasize their complimentary nature, as the flexible natural gas generation offsets the entire volatility increase from wind power. The ISO-NE envisions a long term future electricity mix of 42% wind and 52% natural gas (ISO-NE, 2013), which is relatively close to the volatility neutral growth of 54.4% wind and 45.6% natural gas calculated by the simulation.

The simulation results are given in Figure 2.8, which shows how the volatility value of a typical natural gas generator changes over the next ten years under different generator replacement scenarios that cover the 6,000 MW expected need. The model assumes the aging facilities are phased out linearly and thus replaced at a rate of 600 MW per year. Mean daily demand and intra-day demand volatility are assumed to be constant over time, except as altered through additions of wind capacity.

The four scenarios shown graphically in Figure 2.8 include assumptions for low natural gas replacement, high natural gas replacement, volatility neutral replacement, and the ISO-NE envisioned scenario. The low natural gas replacement scenario assumes the replacement generators come from 20% wind, 0% natural gas, and 80% other, where "other" is assumed to be a volatility neutral generating source. The high natural gas replacement scenario assumes the replacement generators are 20% wind and 80% natural gas. As described above, the volatility neutral scenario

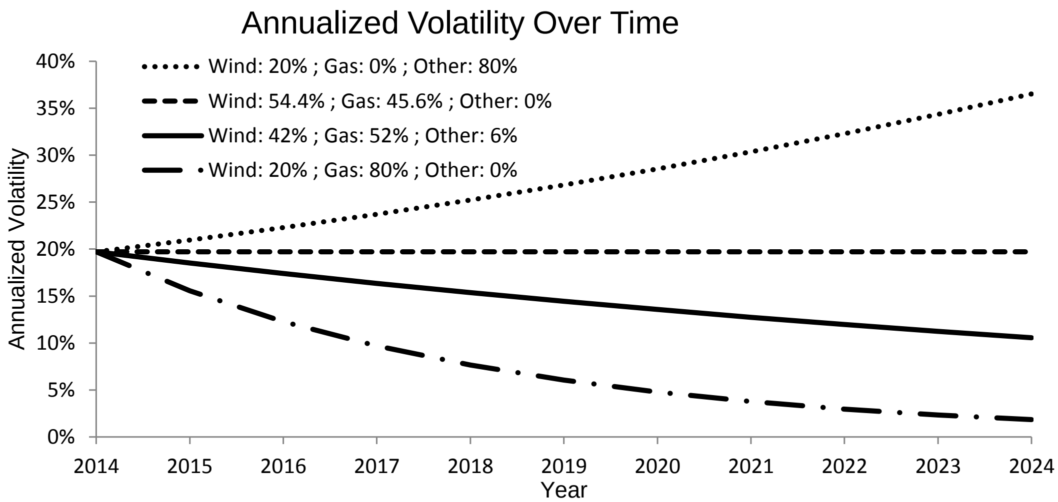

Figure 2.8: Annualized Volatility Over Time

is calculated as replacement generators coming from 54.4% wind and 45.6% natural gas. The fourth scenario uses a replacement rate of 42% wind and 52% natural gas, since this is the ISO-NE envisioned future generation mix.

The simulated annualized volatility is plotted over time in Figure 2.8 in the presence of a shifting generating mix. The low natural gas scenario shows that the marginal value from a natural gas generator increases over time, approximately doubling the annualized volatility percentage. This is because the growth of wind increases price volatility in the future, giving a larger value to volatility reductions from natural gas. This is the opposite of the high natural gas scenario, which shows annualized volatility decaying over time as the growing share of natural gas decreases the additional need for price stability.

The ISO-NE envisioned future shows a gradual decline in the value of volatility reductions from natural gas generators. This is due to the growth of natural gas outpacing the growth of wind, which nets to a slow dampening of price volatility. With the decreasing price volatility comes a decrease in the marginal value of volatility reductions from natural gas, as natural gas generators already represent a large share of the market.

2.6 Natural Gas Capacity and the Forward Premium

Since adding flexible production capacity affects volatility in a similar fashion to electricity storage, there could be implications for the forward premium as well. Douglas and Popova (2008) argue that larger natural gas storage reserves lead to smaller forward premiums, as it is a form of indirect storage. As discussed in Section 2.1, their intuition is largely correct but their econometric model ignores the endogeneity concerns that can bias their results. In this section, I extend their regression analysis with a more rigorous empirical specification that specifically examines the effect of natural gas capacity on the forward premium.

Before starting the regression analysis, recall that the ex-ante forward premium is the difference between the day-ahead price and the expected spot price:

$$PREM_t = FP_t - \mathbb{E}[SP_t] = FP_t - SP_t + u_t \tag{2.12}$$

where $PREM_t$ is the forward premium at time t, FP_t is the forward price, $\mathbb{E}[SP_t]$ is the expected spot price which is assumed equal to the actual spot price plus a random error term, u_t.

The seminal model by Bessembinder and Lemmon (2002) yields a testable hypothesis that the risk premium should be increasing with skewness of the price distribution and decreasing with the variance of the distribution, when generators and retailers are risk averse. Since empirical investigations in the last decade have found mixed evidence in support of this notion (Douglas and Popova, 2008, Haugom and Ullrich, 2012*b*, Longstaff and Wang, 2004), it is worth exploring more in depth here.

The essential intuition is that the risk premium on forward contracts may be lower in markets with lower ramping costs. This is because stored natural gas is equivalent to indirect storage of electricity. Lower ramping costs within new natural gas capacity should imply a greater ability to immediately convert the stored input into electricity. This increases the effectiveness of the indirect physical hedge which reduces the forward premium because the additional ramping ability potentially translates to lower price risk in the spot market.

This notion is tested empirically using a reduced form econometric specification that follows from the previous empirical literature (Douglas and Popova, 2008, Longstaff and Wang, 2004):

$$PREM_t = \beta_0 + \beta_1 NGC_t + \beta_2 VAR_{t-1} + \beta_3 SKEW_{t-1} + \beta_4 T_t + \varepsilon_t \tag{2.13}$$

where $PREM_t$ is the average hourly forward premium on day t, NGC_t is total natural gas capacity, VAR is variance of real-time price, $SKEW$ is the skewness of real-time price, and ε_t is a serially correlated error term such that $\varepsilon_t = \rho\varepsilon_{t-1} + u_t$ where u_t is random noise. More specifically, a 7-day average of intra-day price variance is used for VAR as it arguably represents the best indication of ex-ante

Table 2.12: Regression Results: Natural Gas & the Forward Premium
Dependent Variable: Daily Mean Forward Premium

	(A) OLS	(B) 2SLS	(C) GARCH
Natural Gas Capacity (MW)	-0.0041	-0.0125	-0.0066**
	(0.0047)	(0.0099)	(0.0030)
Variance ($/MWh)	0.0002	0.0001	0.0003
	(0.0002)	(0.0002)	(0.0007)
Skewness ($/MWh)	0.7232	0.6428	0.7548**
	(0.4984)	(0.5682)	(0.3769)
Demand (MWh)	-0.0001	-0.0012	0.0026
	(0.0032)	(0.0023)	(0.0016)
Year Fixed Effects	Yes	Yes	Yes
Observations	2313	2313	2313
Kleibergen-Paap rk-statistic		157.23	

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Newey-West standard errors are reported in parenthesis to correct for serial correlation.

variance expectations. Similarly, a 7-day average of intra-day price skewness is used for $SKEW$. As with previous specifications, T represents a matrix of controls for time trends and includes fixed effects for month, year, and day of the week.

The regression results are presented in Table 2.12. As with previous analysis, Columns (A), (B), and (C) correspond to the pooled event study, 2SLS, and GARCH model, respectively. Across all specifications, natural gas capacity shows a small decrease in the forward premium, although this is only statistically significant in the GARCH model. The coefficients for mean demand and variance are small and insignificantly different than zero in all specifications. The marginal effect of skewness is positive and insignificantly larger in magnitude to the results in Longstaff and Wang (2004), although the difference from zero is only statistically significant in the GARCH model.

Overall, the results are not considered supportive of the Bessembinder and

Lemmon (2002) model and supporting literature (Douglas and Popova, 2008, Longstaff and Wang, 2004). The Bessembinder and Lemmon (2002) model suggests that the forward premium should increase with skewness and decrease with variance, but the effects shown in Table 2.12 are not significant. Instead, these coefficients are more supportive of recent literature by Haugom and Ullrich (2012*b*), who argue that the forward price has converged to be an unbiased predictor of the spot price in the PJM market. This appears consistent with my analysis of the ISO-NE market, which shows an average forward premium of \$0.61, about 1% of the mean electricity price during my sample period.

To ensure that my results are not the result of using aggregated daily forward premiums, I also perform the regression analysis by hour. First, hourly forward premium means with 95% confidence intervals are provided in Figure 2.9 and show premiums statistically different from zero in approximately half of the hours. The regression results again control for the past week's spot price variance, spot price skewness, and mean demand. The marginal effects for natural gas capacity are shown in Figure 2.10. The hourly marginal effects are shown for variance and skewness in Figures 2.11 and 2.12, respectively.

The hourly marginal effects of national gas capacity on the forward premium shown in Figure 2.10 are consistent with the daily forward premium results of Table 2.12. Additional natural gas capacity routinely leads to small, insignificant reductions in the forward premium. The hourly marginal coefficients for variance and skewness are also generally insignificant but do not follow as clear of a trend. The coefficient signs and magnitudes change often, again supporting the work of

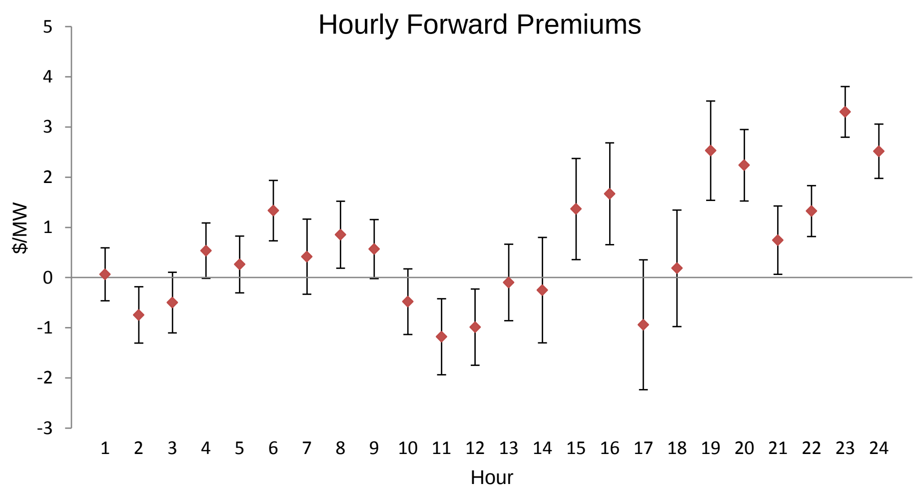

Figure 2.9: Hourly Forward Premiums

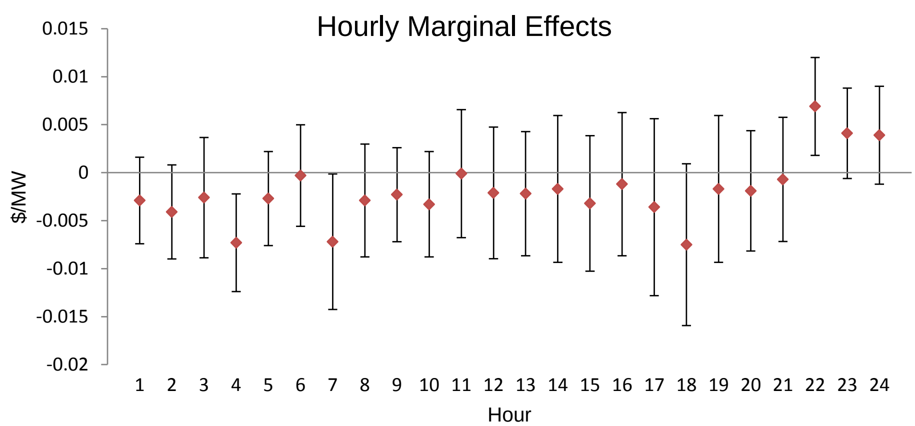

Figure 2.10: Hourly Marginal Effect of Natural Gas Capacity on Forward Premium

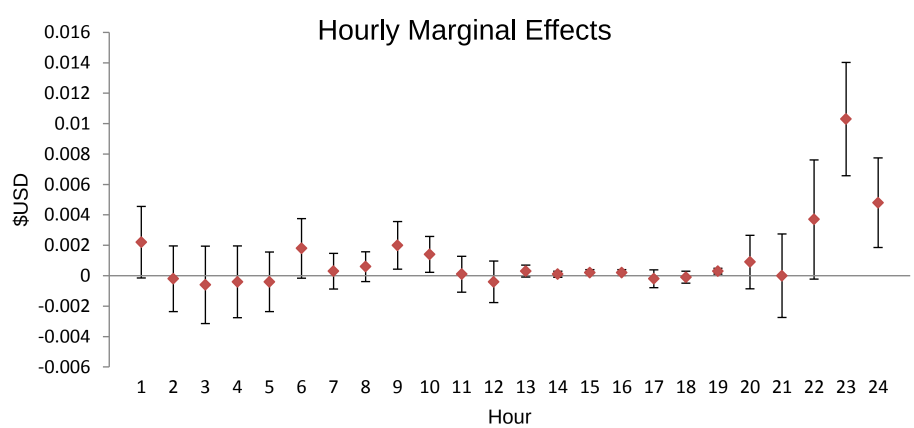

Figure 2.11: Hourly Marginal Effect of Variance on Forward Premium

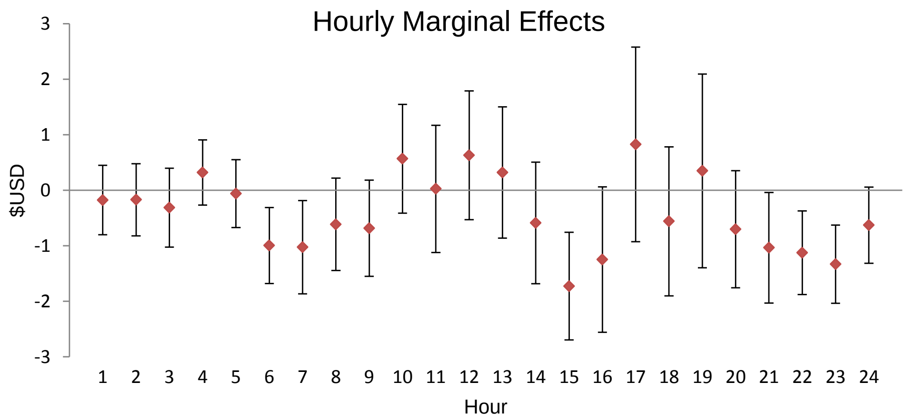

Figure 2.12: Hourly Marginal Effect of Skewness on Forward Premium

Haugom and Ullrich (2012*b*) instead of the earlier literature.

A persistent forward premium implies that there is some risk premium in buying forward price contracts, such that the risk aversion of power purchasers dominates that of electricity generators. Meanwhile, in an efficient market with risk-neutral traders the forward premium should converge to zero (Jha and Wolak, 2013). While there appears little support for the Bessembinder and Lemmon (2002) model with small daily average forward premiums, it is not clear that these results should be interpreted as evidence that the ISO-NE market is operating efficiently in the presence of sufficient risk neutral traders. Although the premiums are small relative to well known inefficient electricity markets Metaxoglou and Smith (2007), statistically significant hourly premiums remain for half the day with primarily positive forward premiums. While it does seem that additional natural gas capacity may slightly decrease the forward premium, this effect is not significant.

2.7 Conclusions

The indirect impacts of additional natural gas capacity on wholesale electricity market price behavior have not been fully analyzed in the previous literature. While natural gas capacity has obvious effects on the mean price of electricity, there is minimal discussion on the implications for price volatility. The ramping ability of natural gas plants is particularly important since there may not yet an efficient market for ramping ability within the FERC's "Standard Market Design" (Angelidi, 2012, Stoft, 2002, Wang and Hobbs, 2014). My analysis provides several contributions to the existing literature on electricity markets, as it describes and quantifies the additional benefits from adding flexible generation capacity. First, it formalizes the intuitive link between natural gas capacity and price volatility due to ramping costs. Second, it implements a rigorous empirical analysis which provides supporting evidence to the theoretical model. Finally, it builds on previous literature connecting natural gas markets and the forward premium in electricity markets, while adding to the debate over the Bessembinder and Lemmon (2002) model.

In this paper I develop a basic theoretical model which details the importance of ramping costs on electricity market price volatility. In the absence of cost-effective storage, ramping costs are a major contributor to price volatility in the electricity market. The model shows that adding generation capacity with lower ramping costs and lower marginal costs will unambiguously decrease intra-day price volatility under standard economic assumptions. Further, the implications of the model easily generalize to all non-storable, or perishable, commodities where there are marginal costs of adjusting output. In brief, flexible production can serve a similar role to storage in ensuring price stability.

A reduced form econometric specification is inferred from the equilibrium conditions of the model and the empirical evidence supports the theory. More specifically, I find that a typical natural gas generator will reduce price volatility by approximately 5.6% in the wholesale market. These results are obtained using a high frequency data within a pooled event study regression analysis, and are robust to a two-stage least squares model and a generalized autoregressive conditional heteroskedasticity (GARCH) model.

The results also show that the effect is larger during the summer months, when intra-day demand is highest and the ramping ability is most important. This suggests the important role that natural gas can play in the future, since expected growth in renewable generation is non-dispatchable and will result in larger residual demand volatilities.

The marginal effects from the econometric results are used to calibrate a simulation exploring how the annualized volatility will change over time in the ISO-NE market as aging generators are replaced. The simulation results show that the annualized volatiity increases dramatically in the presence of large wind growth. Meanwhile, in the ISO-NE envisioned scenario, the annualized volatility falls slowly over the next decade as natural gas growth mitigates the price volatility increases from wind production. This underscores the importance of natural gas as a compliment to non-dispatchable renewable generation because the low ramping costs of natural gas translate to corresponding price stability benefits.

The effect of natural gas capacity on the forward premium is also investigated, showing insignificant decreases to the forward premium from additional natural gas. The results are similar when using daily mean data and disaggregated hourly forward premium data. This supports recent literature suggesting the futures market is operating efficiently in the presence of sufficient risk neutral traders, resulting in the forward price converging to an unbiased predictor of the spot price. However, this result should be interpreted with caution since statistically significant forward premiums remain at the hourly level.

Taken together, the results of this analysis point electricity market regulators towards specific policies. First, market design and policies should acknowledge that there are additional benefits around adding capacity that has both low ramping costs and low marginal costs, such as natural gas generators. This is increasingly important when considering the future growth of non-dispatchable generators such as wind and solar. Volatility-neutral resource planning suggests pairing increases in wind capacity with roughly equal increases in natural gas capacity. Secondly, since the benefits around ramping costs may not be properly priced under the current design of most electricity markets, incentives may be desirable to ensure such benefits are internalized into long-run capital investment decisions. This can be done through an additional market for ramping services, as several transmission organizations have begun to create. In the meantime, construction subsidies may also be offered to ensure additional investment in flexible generators. Incentive-based support mechanisms should remain in place until cost-effective storage reduces ramping issues to irrelevance.

Chapter 3: Price Convergence, Virtual Bids, and Transaction Costs in Electricity Markets

3.1 Introduction

Restructured electricity markets within the United States use a dual-settlement system which yields a day-ahead price, or forward price, as well as a real-time price, or spot price. The forward premium in these markets is important to market efficiency, because accurate price signals are required to properly schedule generators in advance. Since many generators incur start-up costs and require sufficient time to ramp up operations, inaccurate forward prices can result in suboptimal generator dispatch by the independent system operator (ISO) tasked with ensuring grid reliability and effective markets. Thus, the forward premium, which is simply the discrepancy between the day-ahead price and the real-time price, has been a subject of interest to market regulators, transmission operators, and general market participants for over a decade.

In many markets in the United States, virtual bidding has been introduced with the intent to manage price risk, increase financial liquidity, and minimize deviations between forward prices and spot prices. Virtual bids are financial trades which are made in the day-ahead market and closed out in the real-time market, such that no physical commodity has been exchanged. Thus, in a market with risk-neutral participants any forward premium should be arbitraged away, although transaction costs may inhibit price convergence as well (Jha and Wolak, 2013). Persistent deviations beyond a small risk premium have been considered evidence of an inefficient market by previous researchers (Borenstein et al., 2008, Hadsell, 2008, Metaxoglou and Smith, 2007). However, in an efficient market with risk-averse participants, a forward premium may persist due to the positive skewness in electricity prices and a desire for power purchasers to hedge away from high spot price volatility (Bessembinder and Lemmon, 2002, Longstaff and Wang, 2004).

The behavior of the forward premium in electricity markets has been well studied (Bessembinder and Lemmon, 2002, Bowden et al., 2009, Bunn and Chen, 2013, Cartea and Villaplana, 2008, Douglas and Popova, 2008, Haugom and Ullrich, 2012*b*, Longstaff and Wang, 2004), but there are fewer studies on the impact of virtual bids or the role of transaction costs. A report by the New England ISO finds suggestive evidence that the introduction of virtual bids led to price convergence (ISO-NE, 2004). Further, the report notes that a tariff implemented on virtual transactions lowered the volume of virtual transactions, and coincided with a divergence in price, as expected. However, their econometric model ignores the seasonality of the risk premium found by other researchers (Bowden et al., 2009, Bunn and Chen, 2013, Cartea and Villaplana, 2008) and does not control for demand which has been shown to cause a higher risk premium (Longstaff and Wang, 2004).

Jha and Wolak (2013) argue that even without virtual bidding, generators can attempt to exploit the forward premium through altering bids related to physical scheduling. However, this is costly to the generator and imposes additional scheduling risk for the ISO. Thus, they argue that introducing virtual bids essentially lowers arbitrage transaction costs in the California ISO, which should lead to price convergence. Their empirical findings support this notion and are consistent with the official report by the system operator (CAISO, 2012). Within the New York market, Saravia (2003) provides similar findings for the forward premium and Hadsell (2007) shows that the introduction of virtual bidding decreased spot price volatility.

While previous literature provides some evidence that the lower transaction costs of virtual bids lead to price convergence, the introduction of virtual bids is a large market shock introducing many new participants that could alter the risk premium regardless of transaction costs. To more directly test the notion that virtual bid transaction costs are responsible for increasing the forward premium, I take advantage of a natural experiment in the Midwest ISO (MISO) market where virtual bidding was already in practice. In the beginning of the MISO in 2005, virtual bids were not subject to the Revenue Sufficiency Guarantee (RSG) costs, which are essentially a small fee to power producers that ensures generators committed by the MISO are guaranteed cost recovery of startup costs. After a legal debate, the Federal Energy Regulatory Commission (FERC) issued orders requiring RSG fees to be applied to virtual bids. This ruling provides an exogenous increase to virtual bid transaction costs, which allows us to explicitly test the relationship between price convergence and transaction costs.

The following analysis uses the FERC ruling to add to the described literature on price convergence, virtual bids, and transaction costs in electricity markets. I ask three related research questions. First, what is the effect of increasing virtual bid transaction costs on price convergence and intra-day price volatility? More specifically, did imposing RSG costs on virtual bids change the volume of such bids and lead to a larger forward premium? Lastly, does empirically examining the MISO market forward premium support the notion of sufficient risk neutral participants, or does the risk-averse participant model of Bessembinder and Lemmon (2002) apply?

To answer these questions and empirically add to the literature on electricity market design, I use high frequency price data from the MISO day-ahead and real time markets during 2005-2008. To examine virtual bid behavior, I use aggregate level data on cleared virtual bids which also come from the MISO archive. An event study approach is used with a variety of econometric specifications that control for time trends, weather, and past spot price behavior. My analysis shows that the FERC ruling, which increased transaction costs on virtual bids, led to a significant increase in the forward premium within the MISO market. This coincides with a significant drop in the quantity of cleared virtual bids, which is expected when transaction costs are increased. Finally, the analysis coincides with recent literature by Haugom and Ullrich (2012*b*), who find little support for the Bessembinder and Lemmon (2002) model. This is arguably because there is still a sufficient number of risk neutral participants to prevent arbitrage opportunities using past price information alone, even after virtual bids are subject to RSG costs.

The remainder of this paper proceeds as follows. Section 3.2 gives a brief background of the MISO electricity market structure, while the econometric strategy is described in Section 3.3. The related data are analyzed in Section 3.4, with the regression results presented in Section 3.5. The price volatility impacts are analyzed and discussed in Section 3.6, while Section 3.7 concludes.

3.2 Midwest ISO Market Background

Historically, residential electricity in the United States was served through regulated monopolies, but in 1996 the Federal Energy Regulatory Commission (FERC) encouraged the creation of wholesale electricity markets. FERC Orders 888 and 889 established the standard market rules under which the competitive electricity markets are to operate, including the creation of a non-profit Regional Transmission Organization (RTO) to manage the transmission grid and associated electricity markets. By the end of 2001, the Midwest Independent System Operator (MISO)[1] was the first approved RTO and began coordinating the electricity grid in the midwest states. Finally, the MISO established a competitive generating market which began trading on April 1, 2005.

Following the FERC standard market design, the MISO wholesale market includes a dual-settlement system. The dual-settlement system includes a separate competitive auction for both a day-ahead market and a real-time market. The day-ahead market schedules electricity to be delivered in each hour of the following day based on expected demand, so no energy is physically delivered. The real-time market allows adjustments to the delivery schedules, through balancing supply

[1] Due the subsequent growth the MISO coverage area, it is now known as the Midcontinent Independent System Operator, but retains the same acronym.

and demand in each hour. Residential demand faces no immediate price pressure from these wholesale markets, so demand is generally considered exogenous in the subsequent analysis. However, retail utilities have a choice to buy electricity in the day-ahead market or the real-time market, but any unscheduled electricity is required to be purchased in the real-time market.

Given the inflexibility of traditional generators and necessity to perfect balance supply and demand for each hour of the day, the real-time market prices can be especially volatile due to relatively minor deviations from expectations. Thus, many risk averse power purchasers in the wholesale market may prefer the more stable day-ahead market since the price received from their customers is regulated to only change every several months. To ensure that the day-ahead prices and the real-time prices converge, MISO allowed "virtual bids" from the beginning of the competitive markets. Virtual bids are purely financial positions in the day-ahead market that must be closed out in the real-time market. For example, a virtual supply bid in the day-ahead market would increase demand in the real-time market without ever physically changing the final electricity delivered.

Proponents of virtual bids argue they benefit individual market participants as a hedge for supply and demand uncertainty, although there is also some evidence of market manipulation using virtual bids due to market power effects (Birge et al., 2014, Ledgerwood and Pfeifenberger, 2013). Virtual bids also can enhance market liquidity and ensure price convergence of the day-ahead and real-time markets. Allowing virtual supply and demand bids will, in the presence of sufficiently risk neutral traders, prevent any persistently profitable arbitrage opportunities. If anything, only a small risk-premium should remain within the forward premium. However, with only risk averse participants, the forward premium should depend on the variance and skewness of spot market prices, as noted by Bessembinder and Lemmon (2002).

Although virtual bids have been allowed since the beginning, their treatment has changed slightly. In the initial MISO market setup, the rules were unclear if virtual bids should be subject to the Revenue Sufficiency Guarantee (RSG) costs. The RSG is a small fee to power producers that ensures generators committed by the MISO are guaranteed cost recovery of startup costs. For system reliability reasons, the MISO sometimes commits additional generators to be ready for production beyond those that cleared in the day-ahead market. In October of 2005, market rules were clarified that virtual bids should not be subject to the RSG, but this ruling was struck down by the Federal Energy Regulatory Commission (FERC) in April of 2006. This provides a natural experiment with a distinct exogenous shock that increases the transaction costs of virtual bids. Thus, the following analysis builds on the previous literature by providing a more valid and explicit link between virtual bid transaction costs and the forward premium.

3.3 Econometric Specification

Before exploring the data and starting the regression analysis, recall that the ex-ante forward premium is the difference between the day-ahead price and the expected spot price:

$$PREM_t = FP_t - \mathbb{E}[SP_t] = FP_t - SP_t + u_t \tag{3.1}$$

where $PREM_t$ is the forward premium at time t, FP_t is the forward price for hour t, $\mathbb{E}[SP_t]$ is the expected spot price which is assumed equal to the actual spot price plus a random error term, u_t.

In a competitive market with adequate financial instruments and risk neutral participants, the forward premium should converge to zero. In the presence of transaction costs, the forward premium should instead converge to such costs. Thus, the expectation is that the FERC ruling should create a divergence in prices, increasing the forward price premium. This notion is tested empirically using a reduced form event study discontinuity approach that builds from the previous literature (Bessembinder and Lemmon, 2002, Bunn and Chen, 2013, Douglas and Popova, 2008, Haugom and Ullrich, 2012*b*, Longstaff and Wang, 2004):

$$PREM_t = \beta_0 + \beta_1 RSG_t + \beta_2 VAR_{t-1} + \beta_3 SKEW_{t-1} + \beta_4 X_t + \varepsilon_t \tag{3.2}$$

where $PREM_t$ is the average hourly forward premium on day t, RSG_t is a dummy variable indicating the period after the FERC ruling, VAR is variance of real-time price, $SKEW$ is the skewness of real-time price, X represents a matrix of controls, and ε_t is a serially correlated error term such that $\varepsilon_t = \rho\varepsilon_{t-1} + u_t$ where u_t is random noise. More specifically, a 7-day average of real-time price variance is used for

VAR as it arguably represents the best indication of ex-ante variance expectations. Similarly, a 7-day average of real-time price skewness is used for $SKEW$. Included in the matrix of controls, X, are fixed effects for hour, month, year, and day of the week. Several specifications also include temperature controls, since temperature is exogenously responsible for much of electricity demand. Heating degree days and cooling degree days are calculated as the difference from 65°F, as is common when controlling for temperature in the electricity demand literature. The regressions also control for squared terms to account for nonlinear relationships between temperature and electricity demand.

The appropriate event window length is debatable, so several specifications are included with varying window lengths. A longer time horizon allows more accurate estimation of hourly and seasonal effects, though it opens up the possibility of unobservable changes in the market. Thus, the assumptions required to ensure casual inference are generally easier to justify at shorter time horizons.

3.4 Data

Day-ahead and real-time wholesale electricity price data used in the analysis come directly from MISO archive, where the Minnesota Hub prices are used because they are geographically central. A forward premium is calculated as the simple difference between these two prices for all hours from April 2005 through December 2012. Monthly summary statistics of the forward premium for the whole sample period are provided in Table 3.1. The forward premium is higher during the summer and winter months which carry a premium significantly different than zero. This is consistent with previous literature suggesting that the forward premium increases with demand due to high electricity demands in combination with the convexity of the supply function (Bessembinder and Lemmon, 2002, Bowden et al., 2009, Bunn and Chen, 2013, Cartea and Villaplana, 2008, Longstaff and Wang, 2004). A sharp convexity in the supply function implies that symmetric demand shocks create a skewed real-time price distribution which requires a higher risk premium in the forward market. The summary statistics provide mild support for this notion, since months with a higher real-time price skewness generally correspond to higher forward premiums, though August is a notable exception.

As previously discussed, in the initial market design it was unclear if RSG fees would be assessed on virtual bids although they generally were not collected. However, market rules were clarified by early October giving an unambiguous period of lower transaction costs on virtual bids. Thus, Table 3.2 provides summary statistics comparing the 202 days that RSG costs were not assessed with the 202 days after this ruling was struck down by the FERC. As expected by economic theory, the period in which virtual bids were not assessed RSG fees shows a mean forward premium insignificantly different from zero. Meanwhile, in the 202 days after the RSG costs were assigned to virtual bids the forward premium increases significantly to \$2.38/MWh. Unlike the pre-RSG period, the forward premium is also significantly different from zero, with a p-value of 0.0077. In addition, the standard deviation of the hourly forward premium increases from 29.3 to 35.9, and a variance ratio test shows this to be significant at the 1% level. Figure 3.1 shows the distribution

Table 3.1: Summary Statistics for MISO (April 2005 through December 2012)

		Mean, Std. Dev. (in parentheses), and Skewness [in brackets]		
Month	Obs (n)	Day-ahead Price ($USD/MWh)	Real-Time Price ($USD/MWh)	Forward Premium ($USD/MWh)
January	5208	38.7	36.48	2.22***
		(22.33)	(28.17)	(21.51)
		[1.52]	[2.93]	
February	4752	41.73	40.58	1.15***
		(29.57)	(40.01)	(28.9)
		[2.44]	[3.67]	
March	5208	34.91	35.43	-0.52
		(22.86)	(34.59)	(25.33)
		[1.52]	[3.47]	
April	5760	34.51	34.84	-0.34
		(22.47)	(29.32)	(21.04)
		[1.36]	[2.13]	
May	5952	31	30.52	0.48
		(19.11)	(30.66)	(23.2)
		[1.40]	[3.17]	
June	5736	33.24	31.11	2.13***
		(27.67)	(40.9)	(31.23)
		[1.73]	[2.97]	
July	5952	43.97	41.44	2.53***
		(38.51)	(50.83)	(38.8)
		[2.37]	[2.61]	
August	5952	37.22	34.54	2.68***
		(29.74)	(40.26)	(30.81)
		[1.27]	[-1.19]	
September	5760	29.08	29.72	-0.64*
		(21.68)	(32.69)	(26.09)
		[1.60]	[2.24]	
October	5952	34.43	33.98	0.45
		(22.58)	(32.67)	(25.07)
		[1.36]	[2.39]	
November	5760	34.95	33.91	1.04***
		(24.09)	(32.57)	(23.87)
		[1.58]	[2.64]	
December	5952	43.39	41.11	2.28***
		(33.87)	(41.7)	(29.79)
		[2.75]	[3.38]	

Note: ***, **, & * denote statistically significant from 0 at the 1%, 5%, and 10% levels respectively.

Table 3.2: Summary Statistics for MISO (10/6/2005 - 11/14/2006)

		Day-ahead Price ($USD/MWh)		Real-time Price ($USD/MWh)		Forward Premium ($USD/MWh)	
Period	Obs (n)	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Pre-RSG	4848	43.95	34.94	43.98	41.44	-0.027	29.271
Post-RSG	4848	47.34	34.99	44.97	44.40	2.375	35.932
Difference		3.39***	0.05	0.99	2.96***	2.402***	6.660***

Note: ***, **, & * denote statistically significant from 0 at the 1%, 5%, and 10% levels respectively.

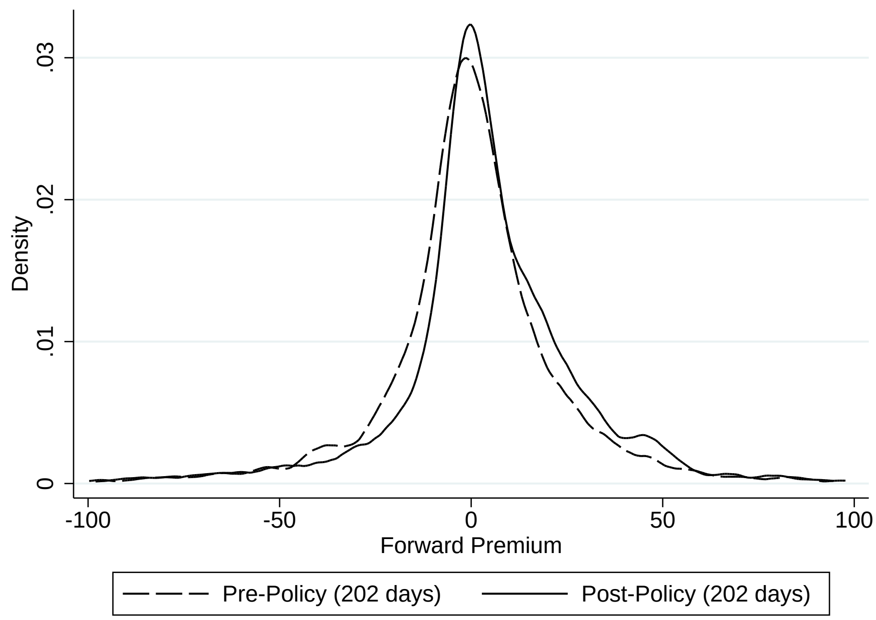

Figure 3.1: Pre/Post Distribution of Forward Premium

of the forward premium in the two periods, showing a positive shift in the post-RSG period. Figure 3.2 gives the disaggregated hourly premiums for the entire 404 day period, showing insignificant premiums for most hours of the day. However, there does seem to be larger and statistically significant premiums during the early evening, which also corresponds to a higher demand periods.

Lastly, given the time-series nature of the data, an Augmented Dickey-Fuller test is performed to ensure the price variables of interest were generated from a

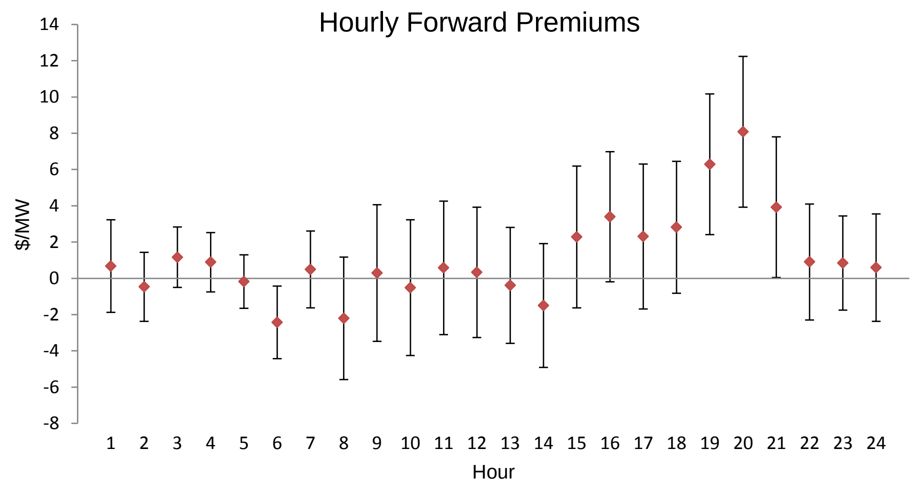

Figure 3.2: Hourly Forward Premium 10/6/2005 - 11/14/2006

stationary process. I reject the null hypothesis that the price variables contain a unit root, with a Dickey-Fuller test statistic of -47.3 and -113.3 for the day-ahead price and real-time price, respectively. Similarly, a I reject the null hypothesis that the risk premium contains a unit root with a Dickey-Fuller test statistic of -154.7.

Figure 3.3 shows the forward premium data marked from the beginning of the market through November 2006. The two vertical black lines mark changes in the treatment of RSG costs for virtual bids, with the first line clarifying that virtual bids are not subject to RSG related fees, although the MISO did not collect RSG fees on virtual bids prior to this clarification. However, since the forward premium in other jurisdictions have been shown to converge to zero as the market matures over several months (Haugom and Ullrich, 2012*b*), the six month period prior to the RSG clarification ruling is excluded in the subsequent regression analysis. It is noteworthy from Figure 3.3 that the weekly forward premium averages following the FERC ruling generally increase though it is not clear how much of this could be due to seasonality from summer months. The regression analysis in the subsequent

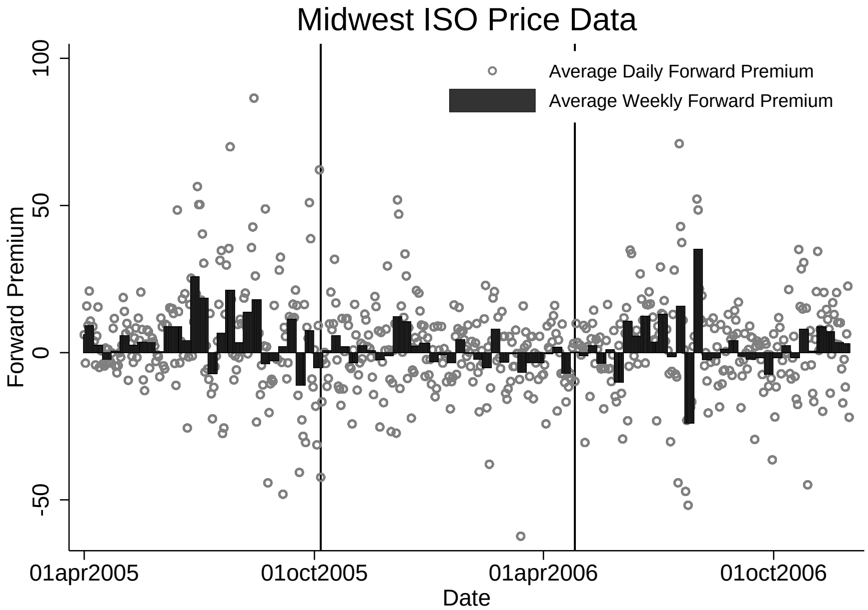

Figure 3.3: Monthly and Daily Forward Premium Averages

section will identify the effect more explicitly, including a variety of specifications that control for seasonal characteristics.

The MISO archive also includes hourly bid data for all cleared virtual bids. Figure 3.4 provides a scatter plot of total daily virtual bids, which are the sum of the hourly cleared bids. The plot also includes a lowess curve fitted for the pre-RSG and post-RSG period. Although the MW of cleared virtual bids seems to be generally declining in the six months prior to the FERC ruling, there is a sharp discontinuity after virtual bids become subject to the RSG fees. Comparing the virtual bid means from two weeks after the FERC ruling to the two weeks prior shows a significant decrease of approximately 30%. This discontinuity is explored further in the regression analysis in the subsequent section, while controlling for a variety of time trends and seasonal effects.

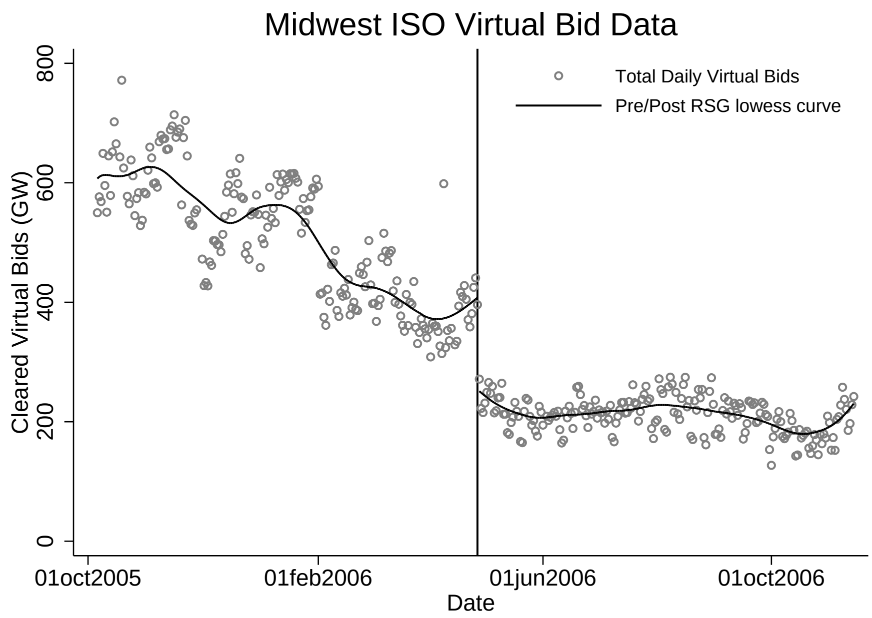

Figure 3.4: Daily Virtual Bids and Lowess Curve

3.5 Regression Results

The regression results show that the FERC ruling subjecting virtual bids to RSG fees results in a significant increase to the forward premium. The regression results are reported in Table 3.3, using a variety of event windows and controls. All columns include controls for time trends and temperature, as discussed in Section 3.3. The row label "RSG" represents a binary variable equal to one in the period following the FERC ruling. Meanwhile, as discussed in Section 3.3, skewness and variance refer to a rolling 7-day averages of the second and third moments of the spot price distribution.

Columns (A) and (B) use the smallest event window, examining only the two weeks before and after the FERC ruling. Column (A) omits controls for skewness and variance, while Column (B) includes them. Both show increases in the forward premium after virtual bids are subjected to the RSG fees. When omitting the spot price distribution controls, Column (A) finds a large and statistically significant effect, while Column (B) finds a modest, insignificant increase when skewness and variance are included. However, the two coefficients are not statistically different from each other, due to the large Newey-West standard errors from the small sample size.

Columns (C) and (D) use the preferred event window length of just over one year, which uses more data for a better estimation of the controls variables. The event window of 404 days is used because there are 202 days between when the FERC ruling and when the MISO clarified the initial ruling that RSG fees would not be assessed to virtual bids. Both columns show a statistically significant increase of the forward premium after the FERC ruling, with no statistical difference between them. The preferred specification in Column (D) includes controls for the spot price distribution and shows the forward premium increased by \$4.6/MWh, or about 10% of the electricity price, after the FERC ruling.

Columns (E) and (F) expand the event window length to the entire sample, from April 2005 through December 2012. Column (E) omits price distribution controls, while Column (F) includes them, and both show the forward premium increasing significantly after virtual bids are subjected to RSG fees. When compared to the preferred specification in Column (D), there is an insignificant change in the RSG coefficient of interest, though the coefficients for skewness and variance do change significantly.

Table 3.3: Regression Results
Dependent Variable: Forward Premium ($/MWh)

	(A)	(B)	(C)	(D)	(E)	(F)
RSG (binary)	5.292***	0.958	3.398*	4.615**	5.024***	4.250***
	(1.996)	(3.791)	(1.856)	(1.887)	(0.861)	(0.852)
Skewness		1.929		0.003		0.257***
		(1.308)		(0.295)		(0.065)
Variance		-0.008		0.004***		0.001***
		(0.006)		(0.001)		(0.0002)
Event Window (days)	28	28	404	404	2,824	2,824
Observations	696	696	9,696	9,696	67,776	67,776
Time controls	Yes	Yes	Yes	Yes	Yes	Yes
Temperature Controls	Yes	Yes	Yes	Yes	Yes	Yes

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively.
Newey-West standard errors are reported in parenthesis to correct for serial correlation.

Across specifications, the coefficients on skewness and variance are shown to be highly unstable. Although the forward premium does appear to increase with spot price skewness, this is only statistically significant in Column (F) where the entire sample is included. Further, the coefficient on variance is positive, small, and significant in this specification as well. Overall, the results are not strong evidence for the Bessembinder and Lemmon (2002) model and endorsing literature (Lazarczyk, 2013, Longstaff and Wang, 2004). The Bessembinder and Lemmon (2002) model suggests that the forward premium should increase with skewness and decrease with variance in the presence of risk averse participants. Instead, these coefficients are more supportive of recent literature by Haugom and Ullrich (2012*b*), who argue that the forward price has converged to be an unbiased predictor of the spot price in the PJM market. While enacting RSG fees are shown to increase the forward premium, this analysis provides suggestive evidence that there is still a sufficient number of risk-neutral virtual bidders to render the Bessembinder and Lemmon (2002) model

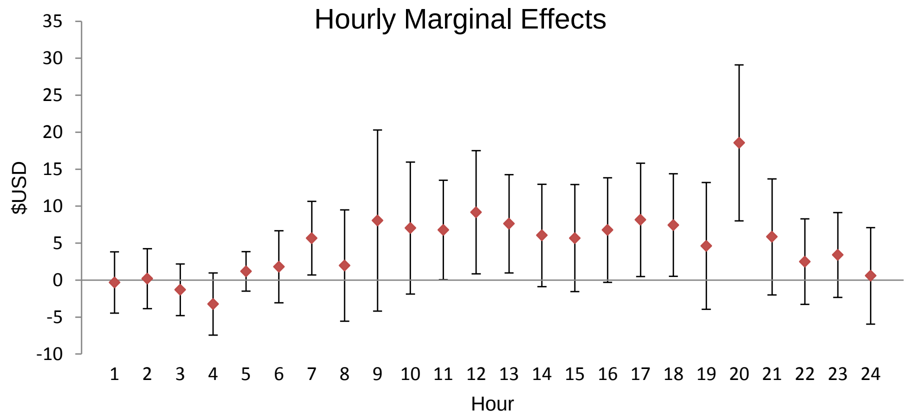

Figure 3.5: Hourly RSG Coefficients

inapplicable.

To explore whether the effect of RSG fees on the forward premium is driven by specific hours, a separate regression is run for each hour. The coefficient results and 95% confidence intervals for the RSG dummy, skewness, and variance are shown in Figures 3.5, 3.6, and 3.7, respectively. Assessing RSG fees to virtual bids appears to have a larger effect during the high demand hours during the day, while having an insignificant effect during the late night and early morning hours when electricity demand is lowest. This supports previous literature arguing that high demand hours have a naturally higher forward premium because greater convexities in the supply function translate to more demand risk for retailers (Longstaff and Wang, 2004). As the transaction costs on risk neutral players increase, forward price behavior is expected to revert towards the theoretical expectations of risk averse players with higher premiums during high demand times.

Overall the hourly coefficients on skewness and variance are contrary to analysis in Douglas and Popova (2008) and generally provide mixed support for the Bessembinder and Lemmon (2002) model. Spot price skewness is shown to increase

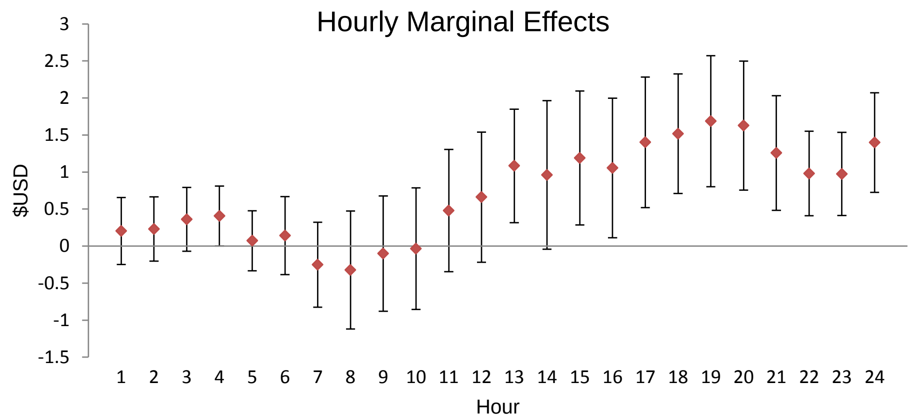

Figure 3.6: Hourly Skewness Coefficients

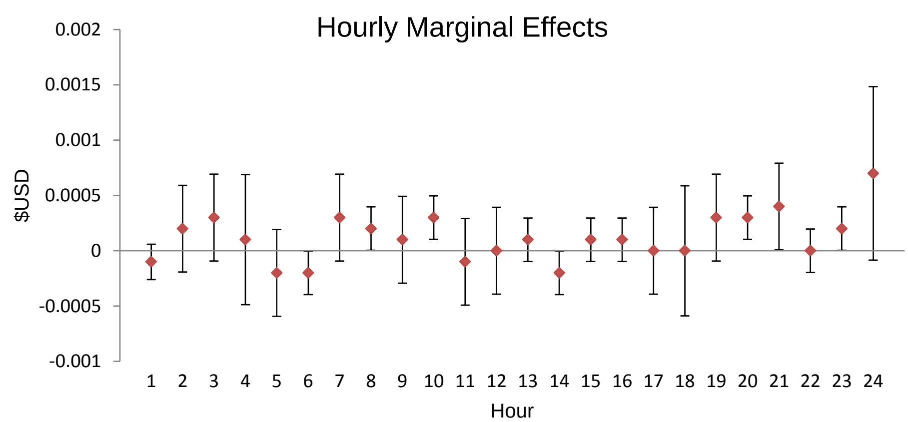

Figure 3.7: Hourly Variance Coefficients

the forward premium during high demand hours, as their model predicts. However, the coefficient for spot price variance provides little support, generally remaining statistically insignificant. The coefficient fluctuates around zero throughout the day with no clear pattern, lending additional support to the findings of Haugom and Ullrich (2012*b*).

To directly test the hypothesis of Haugom and Ullrich (2012*b*) that the forward price has converged with the spot price, a similar analysis is provided on the MISO market. As noted in their analysis, the unbiased forward rate hypothesis provides a simple test of the spot price forecast by the forward price. More specifically,

$$SP_t = \alpha + \beta FP_t + \varepsilon_t \tag{3.3}$$

where FP_t is the forward price for time t in day-ahead market, SP_t is the spot price for time t in the real-time market, and ε_t is a serially correlated error term. Thus, if the forward price is an unbiased predictor of the spot price, α would be insignificantly different than zero and β would be insignificantly different than one.

Following Haugom and Ullrich (2012*b*), the unbiased forward rate hypothesis is tested on mean daily prices and a rolling regression explores how these coefficients change over time. A window size of 365 days is used and the coefficients are shown over time in Figures 3.8 and 3.9 for α and β, respectively. The x-axis shows the ending date of the event rolling regression window, such that the first graphed coefficient covers the sample period from April 1, 2005 through April 1, 2006.

The results show that early in the MISO market history, α and β may have

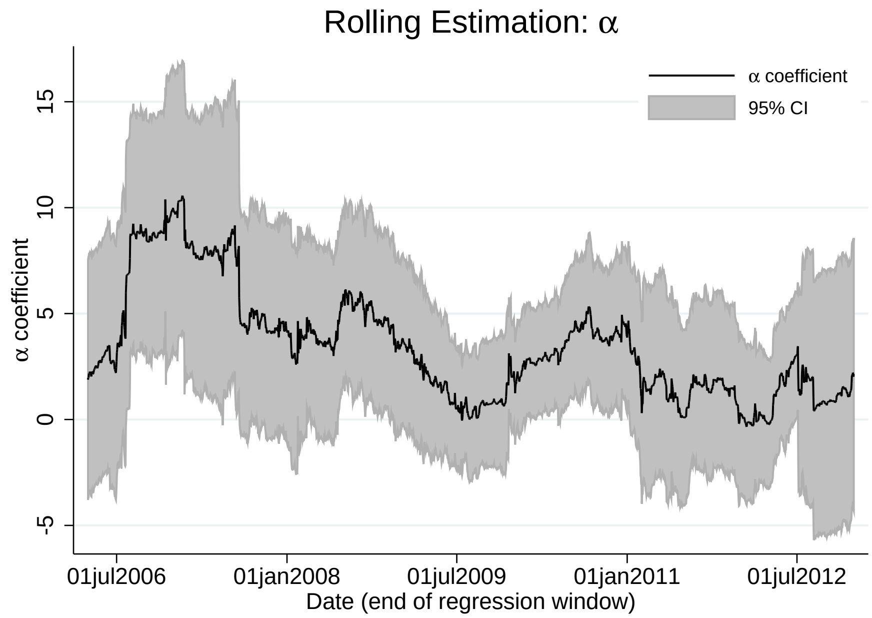

Figure 3.8: α Coefficient Over Time

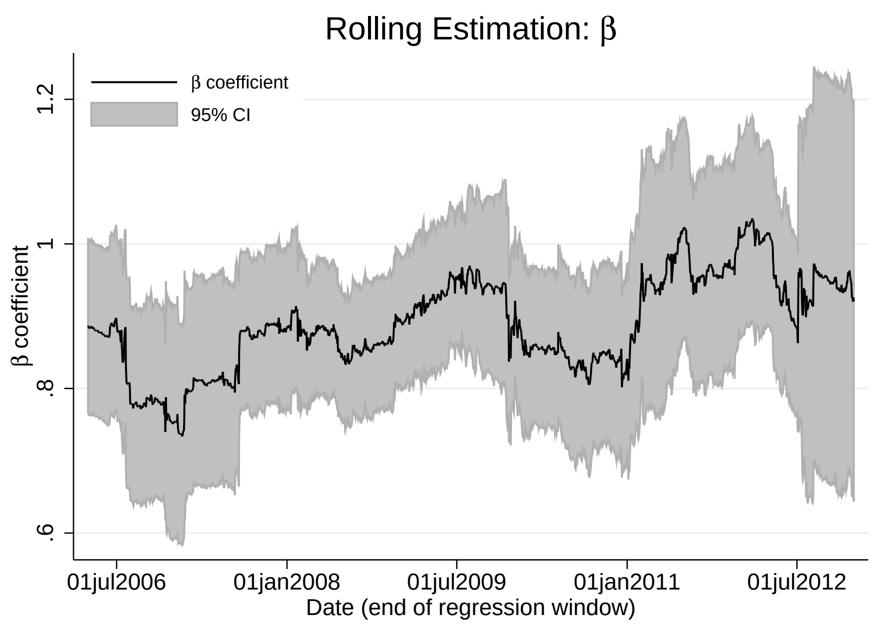

Figure 3.9: β Coefficient Over Time

been statistically different than zero and one, respectively. However, as with the PJM market, the forward price seems to converge to an unbiased predictor of the spot market price as the market matures. As the later years clearly show, there is little evidence of arbitrage opportunities using past price information alone.

3.6 Daily Price Volatility

While the analysis in the previous section argues that the FERC ruling increased the forward premium because of the increased transaction costs on virtual bids, this section analyzes the ruling's impact on price volatility. Previous research argues that the initial introduction of virtual bidding has led to greater price stability (Hadsell, 2007), though this has not been directly analyzed in the context of changing transaction costs on virtual bids. To explore this subject, I first define daily volatility as the standard deviation of the real time price data within the day.[2] More formally:

$$\sigma_t^x = \sqrt{\frac{1}{24}\sum_{h=1}^{24}(x_{t,h} - \mu_t)^2} \tag{3.4}$$

where σ_t^x is intra-day volatility for the variable x on day t, h is the hour of day, and μ is the daily average of x. Thus, throughout the remainder of the analysis I use the terms "volatility" and "standard deviation" interchangeably.

Summary statistics for intra-day real-time price volatility are shown in Table 3.4 for the entire sample. Intra-day price volatility during the entire sample is

[2] This is also sometimes referred to as "historical volatility" in the finance literature, which is distinct from annualized volatility, implied volatility, variance, and the probability of extreme events.

Table 3.4: Summary Statistics for Intra-Day Price Volatility
(April 2005 through December 2012, in $USD/MWh)

	Obs (n)	Mean	Std Dev	Min	Max
January	217	20.12	12.52	2.84	90.41
February	198	23.52	20.73	1.67	132.36
March	217	21.27	17.15	1.77	136.04
April	240	21.19	12.2	2.78	84.69
May	248	21.78	15.21	3.4	102.24
June	239	27.68	22.48	2.89	147.28
July	248	34.25	28.74	3.75	214.4
August	248	28.61	21.1	2.97	159.71
September	240	23.84	17.07	3.17	99.22
October	248	24.11	15.54	2.93	85.15
November	240	23.55	15.13	3.68	91.87
December	248	26.65	20.52	1.4	116.25
Total	2831	24.83	19.16	1.4	214.4

$24.8/MWh, or approximately 70% of the mean real-time prices during this time. However, price volatility itself varies widely, as shown by the minimum of approximately 4% of mean price and a maximum of 608% of mean price. As expected, intra-day price volatility is higher during the summer months, likely because of higher total demand and intra-day demand volatility which are fundamentally driving the intra-day price volatility.

To test if the FERC ruling had any significant effect on intra-day price volatility, a basic multivariate OLS regression model is used with a Newey-West standard error correction because of serial correlation issues. The regression results are reported in Table 3.5 with intra-day price volatility as the dependent variable and the various controls documented in each of the respective columns. The time controls include fixed effects for the day of the week and month, while the temperature controls include daily mean and intra-day volatility for both heating and cooling degree

Table 3.5: Regression Results
Dependent Variable: Intra-day Price Volatility ($/MWh)

	(A)	(B)	(C)	(D)
RSG (binary)	1.688	-2.207	1.507	-3.256
	(2.320)	(2.790)	(2.655)	(2.794)
Event Window (days)	404	404	404	404
Observations	404	404	404	404
Time controls	No	Yes	No	Yes
Temperature Controls	No	No	Yes	Yes

Note: ***, **, & * denote statistical significance at the 1%, 5%, and 10% levels respectively. Newey-West standard errors are reported in parenthesis to correct for serial correlation.

days. For all specifications, the preferred event window of 202 days before and after the FERC ruling is used.

Column (A) of Table 3.5 shows the regression results without any controls, which amounts to a mean comparison. The results show a statistically insignificant increase in volatility of $1.69/MWh or 5.5% of the pre-ruling mean. Column (B) adds time controls and shows the RSG coefficient decreasing to $-2.21/MWh, while Column (C) uses temperature controls to get an RSG coefficient of $1.51/MWh. Lastly, Column (D) includes all controls and again shows no evidence that the higher transaction costs on virtual bids leads to significant changes in intra-day price volatility. Although the sign and magnitude changes across specifications, the columns are not statistically different from each other or from zero. Taken together, this analysis provides no evidence that the increases in forward premiums from the FERC ruling translate to changes in price volatility.

3.7 Conclusions

Price convergence between day-ahead markets and real-time markets has been an important metric for well-functioning electricity markets in the last decade. The forward premium between day-ahead and real-time prices has been noted as a measure of market efficiency by both market regulators and previous researchers (Haugom and Ullrich, 2012*b*, Metaxoglou and Smith, 2007). Since a larger share of electricity is scheduled in the day-ahead markets, inefficient dispatch scheduling can occur if forward prices do not properly predict spot price behavior. This is not easy to correct in the real-time market because engineering limitations create excess start-up costs and production adjustment costs. Thus, many restructured electricity markets have introduced virtual bidding as a simple way to ensure price convergence.

Previous researchers note that some price convergence can be obtained through bid alteration by physical generators, but this is costly and adds unnecessary risk of electricity grid failures. Virtual bids provide a less costly way to achieve price convergence and have been shown to decrease forward premiums in several different electricity markets (Jha and Wolak, 2013). However, a major regulatory change allowing virtual bids can also introduce new market participants with different risk preferences that can alter the forward premium even without reductions in transaction costs. Thus, this paper provides the first to explicitly address price convergence in relation to transaction costs using a natural experiment in the MISO. After litigation, the FERC required revenue sufficiency guarantee (RSG) fees to be applicable to virtual bids, which results in an increase of transaction costs.

Using an event study approach that builds on regression specifications from previous finance literature (Bessembinder and Lemmon, 2002, Haugom and Ullrich, 2012*b*, Lazarczyk, 2013, Longstaff and Wang, 2004), my analysis shows that increasing transaction costs on virtual bids do increase forward premiums and reduce the total number of cleared virtual bids. Although the increase in the forward premium is larger than the actual transaction cost, this may be due to market power issues. Previous research suggests that virtual bids have been used to manipulate prices (Birge et al., 2014, Ledgerwood and Pfeifenberger, 2013), although a thorough investigation of the market power issues surrounding transaction costs. virtual bidding, and the forward premium is left to future researchers. Further, the model supports recent work by Haugom and Ullrich (2012*b*) which argues that the day-ahead price has converged to be an unbiased predictor of the spot price. This notion is also supported when tested directly using a their same methodology in a separate market from their original analysis. A persistent forward premium cannot be fully explained by the past price distribution as hypothesized by Bessembinder and Lemmon (2002) and supported by Longstaff and Wang (2004). These conclusions are supported using daily average premiums as well as a disaggregated hourly analysis. This may be due to a sufficient number of risk-neutral market participants, although statistically significant hourly forward premiums due exist at some hours. Lastly, although the introduction of virtual bids has been shown to provide price stability in other markets (Hadsell, 2007), the transaction cost increases on virtual bids are shown not to increase intra-day price volatility in this case.

Taken together, this research supports the role of virtual bids to aid in price convergence. Electricity market designers interested in minimizing the forward premium to ensure well-functioning wholesale electricity markets should continue efforts to reduce transaction costs on virtual bids and conventional bids, as suggested by traditional economic theory. For historically regulated electricity markets considering restructuring, dual-settle markets which allow virtual bidding should be encouraged to ensure efficient generator dispatch scheduling within competitive wholesale electricity markets. However, if price volatility is the primary concern of market designers, marginal reductions to virtual bid transaction costs may not be sufficient to reduce price risk through price stabilization.

Bibliography

Adelaja, S. and Hailu, Y. (2008), 'Effects of renewable portfolio standards and other state policies on wind industry development in the U.S.', Working Paper.

Anderson, T. and Hsiao, C. (1981), 'Estimation of dynamic models with error components', *Journal of the American Statistical Association* **76**(375), 598–606.

Angelidi, G. (2012), 'Integrated day-ahead market: Draft technical description', California Independent System Operator Technical Report.

Arellano, M. and Bond, S. (1991), 'Some tests of specification for panel data: Monte carlo evidence and an application to employment equations', *Review of Economic Studies* **58**(2), 277–297.

Auffhammer, M., Bento, A. and Lowe, S. (2009), 'Measuring the effects of the Clean Air Act Amendments on ambient PM10 concentrations: The critical importance of a spatially disaggregated analysis', *Journal of Environmental Economics and Management* **58**(1), 15–26.

Bask, M. and Widerberg, A. (2009), 'Market structure and the stability and volatility of electricity prices', *Energy Economics* **31**(2), 278–288.

Berry, D. (2002), 'The market for tradable renewable energy credits', *Ecological Economics* **42**(3), 369–379.

Bessembinder, H. and Lemmon, M. L. (2002), 'Equilibrium pricing and optimal hedging in electricity forward markets', *The Journal of Finance* **57**(3), 1347–1382.

Birge, J., Hortaçsu, A., Mercadal, I. and Pavlin, M. (2014), 'Virtual bidders and financial transmission rights in the MISO electricity market', Working paper.

Bollerslev, T. (1987), 'A conditionally heteroskedastic time series model for speculative prices and rates of return', *The Review of Economics and Statistics* pp. 542–547.

Borenstein, S. (2012), 'The private and public economics of renewable electricity generation', *The Journal of Economic Perspectives* **26**(1), 67–92.

Borenstein, S., Bushnell, J., Knittel, C. R. and Wolfram, C. (2008), 'Inefficiencies and market power in financial arbitrage: A study of California's electricity markets', *The Journal of Industrial Economics* **56**(2), 347–378.

Borenstein, S., Bushnell, J. and Wolak, F. (2002), 'Measuring market inefficiencies in California's restructured wholesale electricity market', *The American Economic Review* **92**(5), 1376–1405.

Bowden, N., Hu, S. and Payne, J. (2009), 'Day-ahead premiums on the Midwest ISO', *The Electricity Journal* **22**(2), 64–73.

Brattle (2010), 'Integrated resource plan for connecticut', The Brattle Group Consulting Report.

Bunn, D. W. and Chen, D. (2013), 'The forward premium in electricity futures', *Journal of Empirical Finance* **23**.

Bushnell, J., Mansur, E. and Saravia, C. (2008), 'Vertical arrangements, market structure, and competition: An analysis of restructured US electricity markets', *The American Economic Review* **98**(1), 237–266.

Bushnell, J., Peterman, C. and Wolfram, C. (2008), 'Local solutions to global problems: Policy choice and regulatory jurisdiction', *Review of Environmental Economics and Policy* **2**(2), 175–193.

CAISO (2012), '2011 annual report on market issues and performance', California ISO Market Monitoring Department.

Callaway, D. and Fowlie, M. (2009), 'Greenhouse gas emissions reductions from wind energy: Location, location, location?', Working Paper.

Cartea, A. and Villaplana, P. (2008), 'Spot price modeling and the valuation of electricity forward contracts: The role of demand and capacity', *Journal of Banking & Finance* **32**(12), 2502–2519.

Cerruti, D. (2014), 'Short and long run effects of vehicle charges: Evidence from Milan, Italy', Working paper.

Chandler, J. (2009), 'Trendy solutions: why do states adopt sustainable energy portfolio standards?', *Energy Policy* **37**(8), 3274–3281.

Chang, Y. and Park, C. (2007), 'Electricity market structure, electricity price, and its volatility', *Economics Letters* **95**(2), 192–197.

Chao, H. (1983), 'Peak load pricing and capacity planning with demand and supply uncertainty', *Bell Journal of Economics* **14**(1), 179–190.

Charmley, W. (2004), 'The federal government's role in reducing heavy duty diesel emissions', SAE Technical Paper 2004-01-2708.

Chen, C., Wiser, R., Mills, A. and Bolinger, M. (2009), 'Weighing the costs and benefits of state renewables portfolio standards in the United States: A comparative analysis of state-level policy impact projections.', *Renewable and Sustainable Energy Reviews* **13**(3), 552–566.

Cragg, J. and Donald, S. (1993), 'Testing identifiability and specification in instrumental variable models', *Econometric Theory* **9**(2), 222–222.

Cullen, J. (2013), 'Measuring the environmental benefits of wind-generated electricity', *American Economic Journal: Economic Policy* **5**(4), 107–133.

Cullen, J. A. and Mansur, E. T. (2013), 'Will carbon prices reduce emissions in the US electricity industry? Evidence from the shale gas experience', Working paper.

de Gorter, H. and Just, D. R. (2008), 'The economics of the US ethanol import tariff with a blend mandate and tax credit', *Journal of Agricultural & Food Industrial Organization* **6**(2).

de Gorter, H. and Just, D. R. (2009), 'The economics of a blend mandate for biofuels', *American Journal of Agricultural Economics* **91**(3), 738–750.

Delmarva (2012), '2012 integrated resource plan', Delmarva Power & Light Company.

Douglas, S. and Popova, J. (2008), 'Storage and the electricity forward premium', *Energy Economics* **30**(4), 1712–1727.

DSIRE (2013), 'Database of state incentives for renewables & efficiency', Department of Energy, administered by North Carolina State University's Solar Center with the Interstate Renewable Energy Council, Inc.

EIA (2012), 'Annual Energy Review 2011', Energy Information Agency Report, U.S. Department of Energy.

EIA (2013*a*), 'Annual Energy Outlook 2013', Energy Information Agency Report, U.S. Department of Energy.

EIA (2013*b*), 'Annual Energy Outlook 2013', Energy Information Agency Report, U.S. Department of Energy.

EIA (2013*c*), 'Electric power annual: 2001-2011', Energy Information Agency Report, U.S. Department of Energy.

Engle, R. F. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation', *Econometrica* pp. 987–1007.

EPA (2002), 'Health assessment document for diesel engine exhaust', U.S. Environmental Protection Agency Office of Transportation and Air Quality.

EPA (2010*a*), 'Final regulatory impact analysis (RIA) for the NO_2 national ambient air quality standards', U.S. Environmental Protection Agency Office of Air Quality Planning and Standards Health and Environmental Impact Division.

EPA (2010*b*), 'Final regulatory impact analysis (RIA) for the SO_2 national ambient air quality standards', U.S. Environmental Protection Agency Office of Air Quality Planning and Standards Health and Environmental Impact Division.

EPA (2011*a*), 'Air emissions sources', United States Environmental Protection Agency. `<http://www.epa.gov/air/emissions/index.htm>`.

EPA (2011*b*), 'Inventory of U.S. greenhouse gas emissions and sinks: 1990-2011', United States Environmental Protection Agency Report.

EPA (2013), 'What are the six common air pollutants?', United States Environmental Protection Agency. `<http://www.epa.gov/air/urbanair/>`.

EPRI (2006), 'Evaluation of thermal-, creep- and corrosion-fatigue of heat recovery steam generator pressure parts', Electric Power Research Institute, Technical Report 1010440.

Fann, N., Fulcher, C. M. and Hubbell, B. J. (2009), 'The influence of location, source, and emission type in estimates of the human health benefits of reducing a ton of air pollution', *Air Quality, Atmosphere & Health* **2**(3), 169–176.

Fell, H. and Linn, J. (2013), 'Renewable electricity policies, heterogeneity, and cost effectiveness', *Journal of Environmental Economics and Management* **66**(3), 688–707.

FERC (2010), '2010 iso/rto metrics report, appedenx f', `<http://www.ferc.gov/industries/electric/indus-act/rto/metrics/iso-ne-rto-metrics.pdf>`.

Fischer, C. (2010), 'Renewable portfolio standards: When do they lower energy prices?', *Energy Journal* **31**(1), 101.

Greenstone, M. (2004), 'Did the Clean Air Act cause the remarkable decline in sulfur dioxide concentrations?', *Journal of Environmental Economics and Management* **47**(3), 585–611.

Hadsell, L. (2007), 'The impact of virtual bidding on price volatility in New York's wholesale electricity market', *Economics Letters* **95**(1), 66–72.

Hadsell, L. (2008), 'Day-ahead premiums on the New England ISO', *The Electricity Journal* **21**(4), 51–57.

Hadsell, L., Marathe, A. and Shawky, H. A. (2004), 'Estimating the volatility of wholesale electricity spot prices in the US', *The Energy Journal* **25**(4), 23–40.

Hadsell, L. and Shawky, H. (2006), 'Electricity price volatility and the marginal cost of congestion: An empirical study of peak hours on the NYISO market, 2001-2004', *The Energy Journal* **27**(2), 157–180.

Haugom, E. and Ullrich, C. (2012*a*), 'Forecasting spot price volatility using the short-term forward curve', *Energy Economics* **34**(6), 1826–1833.

Haugom, E. and Ullrich, C. (2012*b*), 'Market efficiency and risk premia in short-term forward prices', *Energy Economics* **34**(6), 1931–1941.

Heeter, J., Barbose, G., Bird, L., Weaver, S., Flores-Espino, F., Kuskova-Burns, K. and Wiser, R. (2014), 'A survey of state-level cost and benefit estimates of renewable portfolio standards', National Renewable Energy Lab Technical Report NREL/TP-6A20-61042, LBNL-6589E.

Higgs, H. (2009), 'Modelling price and volatility inter-relationships in the australian wholesale spot electricity markets', *Energy Economics* **31**(5), 748–756.

Higgs, H. and Worthington, A. (2008), 'Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The Australian wholesale spot electricity market', *Energy Economics* **30**(6), 3172–3185.

Huisman, R. and Kilic, M. (2010), 'Is power production flexibility a substitute for storability? Evidence from electricity futures prices', Tinbergen Institute Discussion Paper.

Huisman, R. and Kilic, M. (2012), 'Electricity futures prices: Indirect storability, expectations, and risk premiums', *Energy Economics* **34**(4), 892–898.

IPA (2012), '2012 power procurement plan', Illinois Power Agency.

ISO-NE (2004), 'Impact of virtual transactions on New Englands energy market', ISO New England Market Monitoring Department.

ISO-NE (2013), '2013 Regional electricity outlook', ISO New England Annual Report.

ISO-NE (2014*a*), 'ISO New England - History', `<http://www.iso-ne.com/aboutiso/co_profile/history/>`.

ISO-NE (2014*b*), 'ISO New England - How our markets work', `<http://www.iso-ne.com/nwsiss/grid_mkts/how_mkts_wrk/index.html>`.

ISO-NE (2014*c*), 'ISO New England - Key facts', `<http://www.iso-ne.com/nwsiss/grid_mkts/key_facts/index.html>`.

Jha, A. and Wolak, F. (2013), 'Testing for market efficiency with transactions costs: An application to convergence bidding in wholesale electricity markets', Working Paper.

Johnson, E. (2014), 'The cost of carbon dioxide abatement from state renewable portfolio standards', *Resource and Energy Economics* .

Kaffine, D. T., McBee, B. J. and Lieskovsky, J. (2013), 'Emissions savings from wind power generation in Texas', *The Energy Journal* **34**(1), 155–175.

Kaldor, N. (1939), 'Speculation and economic stability', *Review of Economic Studies* **7**(1), 1–27.

Kessides, I. N. (2010), 'Nuclear power: Understanding the economic risks and uncertainties', *Energy Policy* **38**(8), 3849–3864.

Kiviet, J. F. (1995), 'On bias, inconsistency, and efficiency of various estimators in dynamic panel data models', *Journal of Econometrics* **68**(1), 53–78.

Kleibergen, F. and Paap, R. (2006), 'Generalized reduced rank tests using the singular value decomposition', *Journal of Econometrics* **133**(1), 97–126.

Kneifel, J. (2008), 'Effects of state government policies on electricity capacity from non-hydropower renewable sources', Working paper.

Laden, F., Schwartz, J., Speizer, F. E. and Dockery, D. W. (2006), 'Reduction in fine particulate air pollution and mortality: Extended follow-up of the Harvard six cities study', *American Journal of Respiratory and Critical Care Medicine* **173**(6), 667–672.

Lafrancois, B. A. (2012), 'A lot left over: Reducing CO_2 emissions in the United States electric power sector through the use of natural gas', *Energy Policy* **50**, 428–435.

Lazarczyk, E. (2013), 'Market specific news and its impact on electricity prices forward premia', Research Institute of Industrial Economics, IFN Working Paper No. 953.

Ledgerwood, S. D. and Pfeifenberger, J. P. (2013), 'Using virtual bids to manipulate the value of financial transmission rights', *The Electricity Journal* **26**(9), 9–25.

LEI (2012), 'Mpuc rps report 2011 - review of rps requirements and compliance in maine', London Economics International LLC Consulting Report for the Maine Public Utilities Commission.

Longstaff, F. and Wang, A. (2004), 'Electricity forward prices: A high-frequency empirical analysis', *The Journal of Finance* **59**(4), 1877–1900.

Lyon, T. and Yin, H. (2010), 'Why do states adopt renewable portfolio standards?: An empirical investigation', *The Energy Journal* **31**(2), 1140–1149.

Menz, F. and Vachon, S. (2006), 'The effectiveness of different policy regimes for promoting wind power: Experiences from the states', *Energy Policy* **34**(14), 1786–1796.

Metaxoglou, K. and Smith, A. (2007), 'Efficiency of the California electricity reserves market', *Journal of Applied Econometrics* **22**(6), 1127–1144.

Navid, N. and Rosenwald, G. (2012), 'Market solutions for managing ramp flexibility with high penetration of renewable resource', *IEEE Transactions on Sustainable Energy* **3**(4), 784–790.

Navid, N. and Rosenwald, G. (2013), 'Ramp capability for load following in the MISO markets', Midwest Independent System Operator Technical Report.

Nerlove, M. (1971), 'Further evidence on the estimation of dynamic economic relations from a time series of cross sections', *Econometrica* **39**(2), 359–382.

Nickell, S. (1981), 'Biases in dynamic models with fixed effects', *Econometrica* **49**(2), 1417–1426.

Novan, K. M. (2011), 'Valuing the wind: Renewable energy policies and air pollution avoided.', UC Center for Energy and Environmental Economics Working Paper Series E^3 WP-027.

NREL (2014), 'NREL: Dyanmic maps, GIS data, and analysis tools - wind maps', `<http://www.nrel.gov/gis/wind.html>`.

NYSERDA (2013), 'Nyserda renewable portfolio standard main tier 2013 program review - volume 2', New York State Energy Research and Development Authority.

Palmer, K. and Burtraw, D. (2005), 'Cost-effectiveness of renewable electricity policies', *Energy Economics* **27**(6), 873–894.

Pope III, C. A., Burnett, R. T., Thun, M. J., Calle, E. E., Krewski, D., Ito, K. and Thurston, G. D. (2002), 'Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution', *Journal of the American Medical Association* **287**(9), 1132–1141.

PUCO (2013), 'Renewable resources and wholesale price suppression', The Public Utilities Commission of Ohio.

Reguant, M. (2014), 'Complementary bidding mechanisms and startup costs in electricity markets', *Review of Economic Studies* **Forthcoming**.

Roques, F. A., Nuttall, W. J., Newberry, D. M., de Neufville, R. and Connors, S. (2006), 'Nuclear power: A hedge against uncertain gas and carbon prices?', *Energy Journal* **27**(4).

Routledge, B., Seppi, D. and Spatt, C. (2000), 'Equilibrium forward curves for commodities', *The Journal of Finance* **55**(3), 1297–1338.

Routledge, B., Seppi, D. and Spatt, C. (2001), 'The "spark spread:"An equilibrium model of cross-commodity price relationships in electricity', Carnegie Mellon University, Tepper School of Business. Paper 413.

Saravia, C. (2003), 'Speculative trading and market performance: The effect of arbitrageurs on efficiency and market power in the New York electricity market', University of California Energy Institute, UC Berkeley: Center for the Study of Energy Markets, Working Paper 121.

S.B.1078 (2002), 'Senate bill no. 1078, chapter 516', California Legislative Counsel's Digest. <http://www.leginfo.ca.gov/pub/01-02/bill/sen/sb_1051-1100/sb_1078_bill_20020912_chaptered.pdf>.

Schmalensee, R. (2012), 'Evaluating policies to increase electricity generation from renewable energy', *Review of Environmental Economics and Policy* **6**(1), 45–64.

Shrestha, G. B., Song, K. and Goel, L. (2004), 'Strategic self-dispatch considering ramping costs in deregulated power markets', *IEEE Transactions on Power Systems* **19**(3), 1575–1581.

Shrimali, G. and Kniefel, J. (2011), 'Are government policies effective in promoting deployment of renewable electricity resources?', *Energy Policy* **39**(9), 4726–4741.

Simonsen, I. (2005), 'Volatility of power markets', *Physica A* **355**(1), 10–20.

Stock, J., Wright, J. and Yogo, M. (2002), 'A survey of weak instruments and weak identification in generalized method of moments', *Journal of Business & Economic Statistics* **20**(4), 518–529.

Stock, J. and Yogo, M. (2001), 'Testing for weak instruments in linear IV regression', Working paper, Harvard University.

Stoft, S. (2002), *Power System Economics: Designing Markets for Electricity*, Wiley, New York.

Tanaka, M. (2006), 'Real-time pricing with ramping costs: A new approach to managing a steep change in electricity demand', *Energy Policy* **34**(18), 3634–3643.

Ullrich, C. (2012), 'Realized volatility and price spikes in electricity markets: The importance of observation frequency', *Energy Economics* **34**(6), 1809–1818.

Viscusi, W. K. and Aldy, J. E. (2003), 'The value of a statistical life: a critical review of market estimates throughout the world', *Journal of Risk and Uncertainty* **27**(1), 5–76.

Wang, B. and Hobbs, B. (2014), 'A flexible ramping product: Can it help real-time dispatch markets approach the stochastic dispatch ideal?', *Electric Power Systems Research* **109**, 128–140.

Wang, C. and Shahidehpour, S. M. (1994), 'Ramp-rate limits in unit commitment and economic dispatch incorporating rotor fatigue effect', *IEEE Transactions on Power Systems* **9**(3), 1539–1545.

Wang, C. and Shahidehpour, S. M. (1995), 'Optimal generation scheduling with ramping costs', *IEEE Transactions on Power Systems* **10**(1), 11–17.

Wolak, F. (2007), 'Quantifying the supply-side benefits from forward contracting in wholesale electricity markets', *Journal of Applied Econometrics* **22**(7), 1179–1209.

Wolak, F. and Patrick, R. (2001), 'The impact of market rules and market structure on the price determination process in the England and Wales electricity market', NBER Working Paper 8248.

Wolff, H. and Perry, L. (2010), 'Policy monitor trends in clean air legislation in europe: Particulate matter and low emission zones', *Review of Environmental Economics and Policy* **4**(2), 293–308.

Working, H. (1948), 'Theory of the inverse carrying charge in futures markets', *Journal of Farm Economics* **30**(1), 1–28.

Worthington, A., Kay-Spratley, A. and Higgs, H. (2005), 'Transmission of prices and price volatility in Australian electricity spot markets: A multivariate GARCH analysis', *Energy Economics* **27**(2), 337–350.

Xu, L. and Tretheway, D. (2012), 'Flexible ramping products: Second revised draft final proposal', California Independent System Operator Technical Report.

Yin, H. and Powers, N. (2010), 'Do state renewable portfolio standards promote in-state renewable generation?', *Energy Policy* **38**(2), 1140–1149.

Zareipour, H., Bhattacharya, K. and Canizares, C. (2007), 'Electricity market price volatility: The case of Ontario', *Energy Policy* **35**(9), 4739–4748.