Essays on Energy Economics and Market Structure: Investigating Catalysts for Entry and Exit

A Dissertation Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Rebecca Jade Davis

May 2018

© by Rebecca Jade Davis, 2018
All Rights Reserved.

This dissertation and the countless hours spent leading up to and through graduate school are dedicated to my parents. Without their support and endless love, I would not have achieved this accomplishment.

Acknowledgments

I would like to thank the many people who have helped me through the course of this journey. First, I want to thank my family, especially my parents - Brenda, Frank, and Ronnie - and sister, Kelly, for their love, encouragement, and support of an ambitious woman that never took "no" for an answer. They have shaped me into the person I am, and I am forever indebted to them. I express a heartfelt thank you to Daniel, who has been my steadfast companion these past few years. I am very grateful for his unending encouragement and support. You are the light in my darkest days.

I would like to convey my deepest gratitude to my advisor, Charles Sims, for all of his support, guidance, and encouragement during my academic career at the University of Tennessee. I appreciate the countless hours and energy he invested in my graduate education. I would not be where I am without his unfaltering belief in me. Scott Holladay, Matthew Murray, and Burton English are also owed a sincere thank you for their mentorship, assistance, and valuable contributions to this work. Thanks also to Jean Gauger and Don Bruce who have provided me much needed advice and guided me on this path. Thanks to Scott Bellamy, who instilled in me a passion for economics and pushed me to achieve more during my undergraduate studies. I want to extend my appreciation to Sherri Pinkston, the wonderful administrative assistant who always made sure I was on the right track.

Thank you to the friends I have made along the way. Without their companionship, my time completing my dissertation would have been unbearable. I would especially like to thank Beth Glenn, Jill Welch, Matt McMahon, Julie Butler, Clay McManus, Ben Crocker, and Jing Li for being confidants and constant sources of inspiration. I would like to thank my fellow graduate students and the entire faculty at the University of Tennessee for their valuable feedback, which has markedly improved this dissertation.

Lastly, I'm thankful for the financial support for parts of this research that was funded by the Alfred P. Sloan Foundation Pre-Doctoral Fellowship Program on Energy Economics, awarded through the National Bureau of Economic Research.

Abstract

Broadly defined, this dissertation focused on the roles uncertainty, irreversibility, and environmental regulation played in firm entry and exit decisions in energy markets. A prominent innovation of this work was the integration of real options theory and applied econometric techniques. Uncertain energy prices combined with irreversible sunk costs of entry and exit create an economic benefit in delaying entry or exit, known as an option value. An option value approach to understanding firm decisions presents a more robust framework for capturing the impact of uncertainty and irreversibility on energy market structure. Chapters 1 and 2 utilized this modeling approach, while Chapter 3 used applied econometrics.

As a result, this dissertation can be split into two major themes. The first is the role of uncertainty and irreversibility on market entry. Chapter 1 investigated whether a small firm, drilling less than 8 natural gas wells, would have entered the natural gas market during the hydraulic fracturing boom of the 2000s, given historic natural gas prices and reserves. Chapter 1 also examined if technological advances in hydraulic fracturing, horizontal drilling, and 3-D seismic imaging, land lease speculation, or a regime change in natural gas demand drove small firm entry. Each hypothesis was tested using a real options model of market entry and data on natural gas turnover.

The second theme involves the role of uncertainty, irreversibility, and environmental regulation on energy market exit. Chapter 2 investigated this theme with both an optimal stopping model and an empirical analysis of the drivers of coal-fired electricity generator retirement. Chapter 2 developed and implemented a real options model of coal-fired generator retirement to back out the net retirement costs implied by the observed timing of nearly 200 retirement decisions between 2009 and 2015. Propensity score matching was

utilized to assign a retirement cost for the remaining active fleet of coal generators, and the empirical analysis explored the impact of retirement costs on coal generator retirement decisions. Chapter 3 continued work on the second theme by employing a difference-in-differences identification strategy to determine the impact of mercury regulation on coal-fired electricity generator retirements.

Table of Contents

1 Frack to the Future: What Enticed Small Firm Entry During the Hydraulic Fracturing Boom? 1

1.1 Introduction . . . 1
1.2 Entry and Exit in the Natural Gas Market . . . 6
1.3 Data and Parameter Estimation . . . 12
1.3.1 Wellhead Prices & Proved Reserves . . . 12
1.3.2 Sunk Costs . . . 13
1.3.3 Average Production Costs . . . 15
1.4 Results . . . 15
1.4.1 The Benchmark: Market Entry with Pre-boom Expectations . . . 16
1.4.2 Hypothesis 1: Technological Advances . . . 19
1.4.3 Hypothesis 2: Land Leases . . . 20
1.4.4 Hypothesis 3: Prices . . . 22
1.5 Conclusion . . . 24

2 Drivers of Coal Generator Retirements and their Impact on the Shifting Electricity Generation Portfolio in the U.S. 26

2.1 Introduction . . . 26
2.2 Coal-Fired Generator Retirement Model . . . 29
2.2.1 Price Uncertainty . . . 30
2.2.2 The Optimal Timing of Retirement . . . 31
2.2.3 Data and Parameter Estimation . . . 34
2.3 Retirement Cost Analysis . . . 42

2.4 Empirical Analysis . . . 46
2.4.1 Data . . . 47
2.4.2 Estimate Retirement Costs for Active Plants . . . 49
2.4.3 Analysis of Estimated Retirement Costs . . . 50
2.5 Conclusion . . . 56

3 Abate or Exit? The Impact of Mercury Regulation on Coal Generator Retirements 58
3.1 Introduction . . . 58
3.2 Background on Mercury Regulation in the U.S. . . . 62
3.2.1 Federal Regulation History . . . 63
3.2.2 Current Federal Regulation . . . 64
3.2.3 State Regulation Overview . . . 66
3.3 Data and Methods . . . 66
3.3.1 Data . . . 66
3.3.2 Empirical Strategy . . . 70
3.3.3 Descriptive Statistics . . . 72
3.4 Results . . . 72
3.5 Extensions . . . 76
3.5.1 Event Study Analysis . . . 76
3.6 Conclusion . . . 81

4 Conclusions 83

Bibliography 85

Appendices 100
A Appendix 1 . . . 101
A.1 Unit Root Tests . . . 101
A.2 Geometric Mean Reversion Parameter Estimation . . . 101
B Appendix 2 . . . 105
B.1 Unit Root Tests . . . 105

B.2 Geometric Mean Reversion Parameter Estimation 105

Vita **109**

List of Tables

1.1 Natural Gas Entry Parameter Values: Pre-Boom Model, Up to 2000 18
1.2 Pre-Boom Benchmark Model - Minimum Firm Size Entering the Natural Gas Market, 2005-2012 19
1.3 Hypotheses Results, 2005-2008 21

2.1 Average Coal-Fired Generator Parameters by Market Type 37
2.2 Parameter Values for Two Randomly Selected Generators 44
2.3 Data Dictionary 48
2.4 Estimates of the Probability of Retirement 52

3.1 Data Dictionary 69
3.2 Summary statistics for generator outcome, generator controls, plant controls, and state controls in December 2016 73
3.3 Estimates of the Probability of Retirement 77
3.4 Parametric Event Study Estimates 80

A.1 Unit Root Test for U.S. Natural Gas Wellhead Price Data, Monthly January 1976-December 1999 102
A.2 Geometric Mean Reversion Estimates for U.S. Natural Gas Wellhead Price Data 104
A.3 Geometric Mean Reversion Estimates for U.S. Natural Gas Recoverable Reserves per Well Data 104
B.1 Unit Root Test for Delivered Coal Prices at a Random Plant 106
B.2 Unit Root Test for Wholesale Electricity Prices at a Random Generator . . 106

B.3 Geometric Mean Reversion Estimates for Delivered Coal Prices at a Random Plant 108
B.4 Geometric Mean Reversion Estimates for Wholesale Electricity Prices at a Random Generator 108

List of Figures

1.1 U.S. Natural Gas Wellhead Prices and Wet After Lease Separation Proved Reserves per number of gas and gas condensate wells for AR, KY, LA, OH, OK, PA, TX, and WV: 1989-2012. 14

1.2 U.S. Natural Gas Wet After Lease Separation Proved Reserves per number of gas and gas condensate wells for AR, KY, LA, OH, OK, PA, TX, and WV and Nominal Cost per Natural Gas Well Drilled: 1989-1999. 16

1.3 Results with baseline expectations: 2007 - The solid line is the entry threshold of a firm drilling 29 wells. The dotted line is the entry threshold of a small firm drilling 8 wells. 17

1.4 Natural Gas Wellhead Price Densities: The pre-boom density and during-the-boom density have different means and standard deviations. The pre-boom density uses EIA data on monthly wellhead prices from 1976-1999, while the during-the-boom density uses data from 2000-2012. 23

2.1 Retired vs. Operable Coal Generators 35

2.2 Expected and Stochastic Paths for Coal and Electricity Prices. The expected and stochastic paths are simulated using the average $\mathrm{d}P_E$ (Panel A) parameters and $\mathrm{d}P_C$ (Panel B) for all generators. 39

2.3 Coal Generator Technological Efficiency by Operational Year. 41

2.4 Retirement Thresholds for Two Random Generators with and without Electricity and Coal Price Volatility 44

2.5 Retirement Cost Distribution by Market Type 46

2.6 Retirement Cost Density 50

2.7 Marginal Effect of Retirement Costs 54
2.8 Predicted Retirement Density 55

3.1 Number of coal generator retirements between 1993 and 2016. Data are from the Energy Information Administration Form 860 for 2016. A: EPA deemed it "appropriate and necessary" to regulate coal generators, December 2000. B: Final CAMR issued, March 2005. C: D.C. Circuit vacated CAMR, February 2008. D: EPA proposes MATS, March 2011. E: Final MATS issued, December 2011. F: First MATS deadline, April 2015. G: Final MATS deadline, April 2016. 65
3.2 The Marginal Effect of Age on Coal Generator Retirement 75
3.3 Flexible Event Study Estimates 79

A.1 U.S. Natural Gas Proved Reserves per Well, 1989-1999 102

Chapter 1

Frack to the Future: What Enticed Small Firm Entry During the Hydraulic Fracturing Boom?

1.1 Introduction

U.S. natural gas gross withdrawals reached a new high in 2015 at 2.8 trillion cubic feet for the month of March [119]. Production has been on a sharp upward trend since 2005 for two reasons. First, technological advances in horizontal drilling, 3-D seismic imaging, and hydraulic fracturing (fracking) have made it highly profitable for firms to produce large quantities of shale gas [131]. A second, and often, overlooked source of increased production is an increase in the number of active firms in the market.[1] [132] found that between 2000 and 2008 the number of firms actively drilling for natural gas in shale plays increased from 15 to 244. While the natural gas market has historically been dominated by large firms, these new firms were small drilling between one and five wells. Similar changes in market structure have been associated with decreased incumbent firm profits and value [58, 6], overcapitalization [110], productivity growth [62, 42], and increased market responsiveness

[1] A firm is considered active if it drills a well in a particular year; a firm becomes inactive by either exiting the market or mothballing operations. A project is mothballed when it is put into a state of temporary suspension, allowing it to be reactivated in the future at a sunk cost much less than the original upfront sunk costs.

to exogenous shocks [2]. There are also questions about the persistence of these changes in market structure due to the possibility that these small firms are engaging in hit-and-run entry [9]. Identifying the persistence and implications of the recent changes in market structure requires an understanding of the incentives for firms to enter and exit the domestic natural gas market. This study utilizes real options theory and data on natural gas firm turnover to test explanations for the change in market structure that accompanied the recent natural gas fracking boom.

There are three potential explanations for recent changes in the domestic natural gas market. The first is that technological advances in natural gas extraction changed potential entrants' expectations of future resource availability and profits. High rates of entry have been linked with innovation in many industries [44]. In the natural gas industry, two technological innovations increased the expected value and lowered the variance of future reserves. First, advances in horizontal drilling and fracking reduced the costs of production, making previously unrecoverable reserves in shale plays recoverable. Small firms entering the market in the 2000s would have expected recoverable reserves to be larger than had been suggested by historic data. Second, the development of 3-D seismic imaging provided firms with a more accurate estimation of reserves under a piece of land [131]. This made future profits more predictable by reducing the variability in natural gas produced per well.

The second potential explanation concerns land lease speculation. For a firm to enter the natural gas market, it must secure land leases, mineral rights, and drilling permits. A long history of work published in the industrial organization literature illustrates the important role sunk costs play as a barrier to market entry. Small firms that entered the natural gas market in the 2000s were quite possibly engaged in lease speculating. Because land leases can be resold, small firms that entered the natural gas market in the 2000s may have been acquiring mineral rights with the expectation of reselling the lease at a higher price. [132] noted that leases typically require the firm to drill at least one well within a specific time frame and postulated that speculation could be the reason they observed a large number of firms drilling only a single well in their data. Lease appreciation would erode barriers to market entry, such as sunk costs [7, 88, 8, 100, 22], thus making market entry more appealing.

The third potential explanation is regime change in natural gas demand. The 2000s were marked by the highest and most volatile natural gas prices in history. There are numerous reasons potential entrants may have viewed these higher and more volatile prices as a sign of a regime change in natural gas demand. Large energy consumers began substituting coal for natural gas partly in response to expected changes in environmental policy. For example, a widely-adopted strategy for many electric utilities in response to the Clean Power Plan was to retire coal-fired generation capacity and invest in natural gas generation [89]. This increase in natural gas generation market share may have been viewed by potential entrants as a permanent shift to a new regime characterized by higher but more volatile natural gas prices.

Static two-stage oligopoly models are typically used to study market entry in the industrial organization literature [14]. These models are ill-suited for our analysis for two reasons. First, they typically assume incumbents and new entrants are described by a single representative firm. This runs counter to empirical evidence suggesting that new entrants in the natural gas market are much smaller than incumbent firms. Second, static models struggle to capture the effect of firm expectations on market entry decisions [51]. Changes in expected future profits are at the heart of all three of the potential explanations for natural gas market structure changes.

An alternative framework for studying the effect of these explanations on market entry that captures firm heterogeneity and expectations of future profits is real options theory [33, 91, 35, 40]. Real options theory views market entry as an uncertain investment and treats the decision to enter as an investment option. There is a real economic benefit (an option value) for potential entrants in the natural gas market to delaying sunk entry costs in order to observe the evolution of recoverable reserves and prices. A standard Marshallian investment rule would ignore the option value and suggest entry as soon as expected discounted profits marginally exceed the fixed costs required to enter the market. Market entry based on such a strategy would occur too soon and would fail to maximize the value of the firm.

However, accounting for the effect of lease speculation on market entry requires that we account for the incentives that influence market exit as well since the capital gains from lease appreciation are only felt when the firm resells the lease upon market exit. To address the

lease speculation argument, we adopted a regime switching model [16] where firms chose the timing of entry and exit given uncertainty in future reserves and prices. The method treats entry and exit as a set of compound investment options [45], which allows us to account for the option value associated with entering and exiting the natural gas market. These options are linked since entry triggers the option to exit. This feature of the regime switching model is particularly appealing for our investigation of land lease speculation. The cost of exiting the market in this case would be a payment to the firm for selling the mineral rights at a higher price than when purchased. Because the regime switching model accounts for exit costs and the option to exit, we can investigate the possibility of land lease speculation as a driver of small firm entry during the natural gas boom. Taking a regime switching approach allows us to account for hysteresis in which entry and exit depends not only on current reserves and prices, but also the history of reserves and prices [34].

Previous applications of real options theory to energy resources have focused on the timing of production from a single well or mine [87, 26, 24, 66, 60, 82].[2] [76] noted that the mine operator's decision to extract resources from the mine or shut down is analogous to a firm's decision to enter or exit an industry.

However, the timing of market entry differs from the timing of production for a single resource asset in several ways. First, the reserves associated with a single well or mine are finite and decrease with extraction, whereas the reserves held by a firm may increase due to the acquisition of additional mineral rights as long as the firm remains in the market. Treating market entry as the start of production for a single resource asset effectively predetermines the length of time the firm remains in the market. Second, the size of reserves held by the firm in the future is difficult to predict and evolves stochastically as new mineral rights are obtained by the firm. The reserves associated with a single well or mine are relatively well-known when production is initiated, and any uncertainty is more appropriately modeled as a time-invariant random variable. This suggests firms considering market entry face more sources of uncertainty than firms considering the timing of production for a given resource asset. This distinction is even more pertinent for entry into the natural

[2]Real options literature contains many applications in environmental and natural resource economics. These papers cover a wide range, including real options applications to climate change policy [23, 92, 94, 95, 96], renewable resources [21, 99, 70, 59, 83, 107], and nonrenewable resources [17, 25, 108, 27, 28].

gas market due to the relatively short lifetime of a fracked well.[3] Third, the sunk costs of market entry differ from the sunk costs associated with production.[4] The sunk costs incurred to extract from one well or mine typically derive from purchases of capital equipment for drilling (either rented or purchased rigs) and general construction costs. However, the sunk costs associated with entering the natural gas market include the cost of securing the mineral rights and obtaining multiple drilling permits. These costs can constitute 28.8% of the total drilling costs but are typically omitted from studies that focus on the optimal timing of resource extraction [52].

The current investigation of natural gas market entry forces revealed three key findings. First, the expectation of higher natural gas proved reserves per well and higher natural gas wellhead prices are two likely reasons for the boom in small firms entering the natural gas market. It takes less than a 25% increase in expected reserves per well to entice small firms to enter. In contrast, the long-run mean reserve level increased 39% between 2000 and 2014 [115]. This suggests that technological advances such as horizontal drilling and hydraulic fracturing that increased economically viable natural gas reserves helped change the domestic natural gas market structure. In addition, the average natural gas wellhead price in the 2000s was double that of the previous 25 years. If small firms based their expectation of future prices only on observed prices after the boom began (2000 to 2012), our model suggests they would have entered between 2005 and 2008, which is precisely what the data indicates. Second, land lease speculation could also explain small firm entry during the boom. If these firms expected their land lease to appreciate by 30%, a firm drilling just one well would enter at any wellhead price and reserve quantity. Last, implausibly large reductions in reserve volatility would be needed to explain market entry by smaller firms. This suggests that 3-D seismic imaging is an unlikely driver of the observed changes in the domestic natural gas market structure.[5]

[3]The lifetime of a fracked natural gas well is approximately two years, whereas the lifetime of a copper or coal mine can routinely extend beyond twenty years [73, 111].

[4]Following [10], we define sunk costs as costs that cannot be eliminated even by total cessation of production. In contrast, fixed costs are costs that are not reduced by decreases in output so long as production is not discontinued altogether. Thus, not all sunk costs are fixed and not all fixed costs are sunk.

[5]However, these factors may have driven the fracking boom through increases in average production per well, which our model does not consider.

The following section presents the model, and the data and parameter estimation are described in in Section 3. Section 4 discusses the results, and Section 5 offers conclusions.

1.2 Entry and Exit in the Natural Gas Market

A firm may be in one of two states: out of the market ($S = 1$) or in the market ($S = 2$). Based on its existing status, the firm receives a flow payoff $\pi(P, R, S) = \Big(\Big(P(t) - AC(R(t))\Big)Q(S)\Big)$ where Q represents the total natural gas produced by the firm. $P(t)$ is the natural gas wellhead price, and $R(t)$ is the proved natural gas reserves per well held by the firm. $AC(R)$ is the firm's average cost of extraction with $\frac{\partial AC}{\partial R} < 0$ [90]. The firm produces no natural gas if it is out of the market: $Q(1) = 0$. When the firm is in the market, it produces natural gas $Q(2) > 0$ that generates a flow of profits.

Market entry constitutes a move from S=1 to S=2 and instantaneously triggers the flow payoff $\pi(P, R, 2)$. However, entry requires a firm to pay two costs. The first is sunk costs associated with leasing, acquisition, and permitting. Firms acquire mineral rights by consulting geologists on land choice, hiring landmen to negotiate terms with landowners, investigating the mineral interests, and researching the title to ensure the correct landowner is associated with the lease. If the preliminary title check is completed and approved, the lease is signed, taken to the county courthouse, and filed by paying a fee. A signing bonus is typically paid up front and is an important factor for landowners when negotiating the terms of their mineral rights [52]. Following initial research on the mineral rights, the title is further investigated by a title abstractor to check for errors, heirs, rights-of-way, wills, and unrelated mineral interest conveyances. This is done before any construction can be initiated on the land and adds another cost to the firm. One last inspection of the title is completed, which is known as curative title and development. Any missed research is overseen by legal entities before drilling can commence. The curative team may find, for example, an additional family member who must sign the lease. Work completed by the curative team closes out the labor costs associated with leasing and acquisition. The firm must also obtain various permits for drilling for natural gas before entering the market. Once permitting is complete, the firm may begin construction of wellpads and rigs. These sunk

entry costs, K, are a function of the size of the firm proxied by the quantity of natural gas produced, Q, with $\frac{\partial K}{\partial Q} > 0$ and $\frac{\partial^2 K}{\partial Q^2} < 0$. The second entry cost is the price paid for the lease L. The firm can enter the natural gas market at sunk cost $K(Q) + L$ and receive a flow of profits $\pi(P, R, 2) = \Big(\Big(P(t) - AC(R(t))\Big)Q\Big)$.

If prices or reserves fall, a firm may choose to exit the market. Market exit constitutes a move from $S = 2$ to $S = 1$ and terminates the flow of profits: $\pi(P, R, 1) = 0$. Market exit requires a fee, F, that represents remediation and cleanup costs but also allows the firm to sell the mineral rights it holds at price sL. The lease depreciates in value while the firm is in the market when $s < 1$ and appreciates in value when $s > 1$. If prices or reserves rise in the future, the firm may choose to re-enter the market, which requires the firm to pay $K + L$ but restarts the flow of profits. Because market exit is not permanent, our model accommodates both decommissioning ($F > 0$) and temporarily mothballing drilling operations ($F = 0$) to avoid environmental cleanup [82].

While current prices and reserves are known, future prices and reserves are unknown. For example, proved natural gas reserves per well $R(t)$ for a firm's current lease holdings are known, but future levels of reserves held by that firm are unknown. The volume of natural gas that ultimately will be produced cannot be known ahead of time, and the estimates change as extraction technologies improve, markets evolve, and natural gas is produced. Geometric Brownian motion (GBM) has been used to model stochastically evolving reserves for a single resource asset [90]. For the purposes of this study, the stochastic reserve process must reflect the evolution of a firm's reserve holdings. Firms are continuously acquiring new leases, which allows them to replenish what they extract. Reserves per well are assumed to evolve according to geometric mean reversion (GMR), $\mathrm{d}R = r_R(\bar{R} - R)R\mathrm{d}t + \sigma_R R\mathrm{d}z_R$. Here, r_R is the rate of reversion to the mean reserve level, $\bar{R}$ is the long-run mean reserve level, and σ_R is the standard deviation rate. $\mathrm{d}z_R = \epsilon(t)\sqrt{\mathrm{d}t}$ is the increment of the standard Weiner process, where $\epsilon(t)$ is a standard normal variate. Allowing the stock of reserves held by the firm to revert to a long-run mean assumes that exploration, research and development, and

technology advances permit discovery and recoverability of natural gas that matches the amount of reserve withdrawals over time.[6]

Future wellhead prices are also unknown; wellhead prices evolve randomly around a long-run mean following a similar geometric mean-reverting process $\mathrm{d}P = r_P(\bar{P} - P)P\mathrm{d}t + \sigma_P P\mathrm{d}z_P$. The uncertainty encompassing future wellhead prices is due to fluctuations in the market. Shifts in supply and demand influence market prices, and the price elasticity of both demand and supply determine the degree of the price response. For example, there is a significant lead time required for firms to bring additional natural gas to the market since pipeline capacity expansions are necessary to eliminate transmission issues. The uncertainty is not idiosyncratic to a firm. [93] determines that energy prices, including natural gas, are mean-reverting by testing a century's worth of data. Modeling natural gas wellhead prices as GMR prevents any negative prices. Real options results depend on choosing the correct stochastic process; therefore, a unit root test is used to check the GMR assumptions for natural gas wellhead prices. We reject the null hypothesis that the price process follows geometric Brownian motion (Appendix 1). A lack of data prevents the completion of a unit root test on proved reserves per well. However, visual inspection of the data suggests they do not follow geometric Brownian motion (Appendix 1, Figure A.1).

The decision problem is presented in terms of a risk-neutral firm whose objective is to determine if and at what time to enter t_E and exit t_X the natural gas market to maximize the firm's expected discounted profits. The decisions to enter and exit the market in this discontinuous or threshold control setting are made with the knowledge that all future adjustments will be optimal. Using traditional cost-benefit analysis, the firm would enter the market when the expected net present value of profits equals or exceeds the costs associated with entry. However, since many of the costs of entry are sunk, there is an incentive (an option value) to delay entry longer than suggested by cost-benefit analysis to further observe how profitable it will be to drill a natural gas well.

[6] As a robustness check, we model natural gas proved reserves per well as a declining GBM. Switching from geometric MR to declining GBM does not significantly change the results.

At each instant in time, the firm must determine whether to enter the market or stay out given that all future exit and entry decisions are made optimally. Given the discount rate δ, the optimal entry and exit time satisfies the following:

$$V(P_0, R_0, 1) = \max_{t_E} \mathrm{E}_0 \left[\Big(V\Big(P(t_E), R(t_E), 2\Big) - L - K(Q) \Big) e^{-\delta t_E} \right] \tag{1.1}$$

$$V(P_0, R_0, 2) = \max_{t_X} \mathrm{E}_0 \left[\int_0^{t_X} \pi\Big(P(t), R(t), 2\Big) e^{-\delta t} dt + \Big\{ V(P(t_X), R(t_X), 1) - F + sL \Big\} e^{-\delta t_X} \right] \tag{1.2}$$

subject to dP, dR, $P(0) = P_0$, and $R(0) = R_0$.

[35] showed that the decision to enter the market is made based on a comparison of the optimal value function that arises when continuing with the status quo, $S = 1$, and the present value of profits when the firm enters the market, $S = 2$, minus the sunk costs associated with entry $K(Q)$. $V(P_0, R_0, 1)$ is simply the option value since there is no flow of profits when the firm is not in the market. This option value represents the value of delaying entry into the market to gain more information about the profitability of drilling for natural gas.[7] $V(P_0, R_0, 2)$ is the expected net present value of profits plus the option value associated with exiting the market (net the cost of exit).

[16] established that the optimal switching problem can be rewritten as a set of variational inequalities. Before entering the market, the optimal value function satisfies

$$\begin{aligned} \delta V(P, R, 1) \geq\ & r_P(\bar{P} - P) P \frac{\partial V(P, R, 1)}{\partial P} + r_R(\bar{R} - R) R \frac{\partial V(P, R, 1)}{\partial R} + \\ & \frac{1}{2}\sigma_P^2 P^2 \frac{\partial^2 V(P, R, 1)}{\partial P^2} + \frac{1}{2}\sigma_R^2 R^2 \frac{\partial^2 V(P, R, 1)}{\partial R^2} + \sigma_P \sigma_R \rho P R \frac{\partial^2 V(P, R, 1)}{\partial P \partial R} \end{aligned} \tag{1.3}$$

$$V(P, R, 1) \geq V(P, R, 2) - L - K(Q) \tag{1.4}$$

[7] When the firm enters the market, this option value is terminated, making it an additional opportunity cost of entering. It is this opportunity cost that prompts a more cautious response by the firm in the face of uncertainty.

ρ is the correlation coefficient between the two stochastic processes $P(t)$ and $R(t)$: $\rho = corr(\mathrm{d}z_P, \mathrm{d}z_R)$. In financial terms, the firm holds an asset whose value $V(P, R, 1)$ must be optimally managed (i.e. maximized). The left-hand side of (1.3) is the return the firm would require to delay entering the market over the time interval dt. The right-hand side of (1.3) is the expected return from delaying market entry over the interval dt. This equation acts as an equilibrium condition ensuring a willingness to delay prior to market entry. Equation (1.4) compares the total payoff when the firm is in the market and out of the market and acts as a boundary condition for the entry regime. One of the conditions (1.3) or (1.4) is satisfied at each point in the state space of $P(t)$ and $R(t)$. If (1.3) holds as an equality, it is optimal to delay entering the market (remain in regime 1). However, if (1.4) holds as an equality, it is optimal to enter the market immediately (switch from regime 1 to 2).

When the firm is currently in the market, the optimal value function satisfies

$$\delta V(P,R,2) \geq \pi(P,R,2) + r_P(\bar{P}-P)P\frac{\partial V(P,R,2)}{\partial P} + r_R(\bar{R}-R)R\frac{\partial V(P,R,2)}{\partial R} + \frac{1}{2}\sigma_P^2 P^2 \frac{\partial^2 V(P,R,2)}{\partial P^2} + \frac{1}{2}\sigma_R^2 R^2 \frac{\partial^2 V(P,R,2)}{\partial R^2} + \sigma_P \sigma_R \rho P R \frac{\partial^2 V(P,R,2)}{\partial P \partial R} \quad (1.5)$$

$$V(P,R,2) \geq V(P,R,1) - F + sL \quad (1.6)$$

Equation (1.5) compares the expected and required return from delaying market exit. If equation (1.6) holds as an equality, it is optimal to remain in the market (remain in regime 2). Much like equation (1.4), equation (1.6) acts as a boundary condition for regime 2. If equation (1.6) holds as an equality, it is optimal to exit the market (switch from regime 2 to 1).

$V(P, R, 1)$ includes an option value that delays firm entry. $V(P, R, 2)$ includes an additional option value that delays firm exit. From (1.6), this additional exit option value encourages more immediate market entry compared to the case in which entry is irreversible. This highlights the linked nature of the market entry and exit options. The ability to exit the market makes the firm less cautious about entering the market, but the possibility of

exiting the market also increases the option associated with entering the market. Intuitively, the option to enter a market is more valuable if it is only partially irreversible. The impact of exit on the timing of market entry depends on which of these two effects dominate.

The solution to the variational inequalities in (1.3)-(1.6) can be characterized by an entry curve $P_E(R)$ that separates the state space where market entry should occur conditional on the firm being out of the market at the time. Specifically, the entry curve is the set of points where conditions (1.3) and (1.4) are met. An exit curve $P_X(R)$ divides the state space into regimes where the firm should and should not exit given where the firm is positioned in the market. Based on expectations of future profits, the firm optimally enters the market when an increase in prices or reserves crosses the threshold curve $P_E(R)$. The firm optimally exits the market if a drop in prices or reserves cross the threshold curve $P_X(R)$.

Numerical methods are used to determine $P_E(R)$ and $P_X(R)$ for all levels of Q. One goal of this study was to find the minimum total production level of a firm (Q) that entered the natural gas market during the boom and determine if that production size is consistent with a small firm. Numerical methods are required to approximate the unknown value function due to the multi-dimensional nature of the state space [81]. Using piecewise linear basis functions, we approximated $V(P, R, 1)$ and $V(P, R, 2)$ over a subset of the state space [75]. The approximation procedure solves for the 2 x n x m basis function coefficients, which satisfy (1.3) - (1.6) and relevant boundary conditions at a set of $n = 50$ and $m = 100$ nodal points spread evenly over the two-dimensional state space.[8] The following necessary boundary conditions are true whether $S = 1$ or $S = 2$: $V(0, R, 1) = 0$ and $V(P, 0, 1) = 0$. These ensure that there is no value of entering the market when natural gas wellhead prices are 0 or when proved reserves per well are 0.

[8]Upwind finite difference approximations are used to construct a linear spline, which approximates the unknown value function. We used Matlab, along with the CompEcon Toolbox and the smoothing-Newton root finding method, to solve the resulting complementarity problem. The approximated state space ranges from 0 to 30 in the P dimension with $m = 50$ nodal points and from 0 to $1,270,000$ in the R dimension with $n = 100$ nodal points. Extending the state space in either the $P(t)$ or $R(t)$ dimension or increasing the number of nodal points beyond 50 and 100 does not alter general results.

1.3 Data and Parameter Estimation

Estimates of the parameters included in (1.3)-(1.6) are required to approximate $V(P, R, 1)$ and $V(P, R, 2)$ over a subset of the state space. To investigate the three possible explanations behind the structural change of the natural gas market, we parameterized a pre-boom model where firms form expectations over future prices and reserves based on the natural gas market in the 1970s, 1980s, and 1990s - where prices were less volatile and reserves were less plentiful. Following [55], we also assumed natural gas firms use a 9% discount rate $\delta = 0.09$.[9] All parameter values are for the baseline, pre-boom model (Table 1.1). Key model parameters were adjusted to reflect changing firm expectations and lease speculation behavior.

1.3.1 Wellhead Prices & Proved Reserves

A critical step in solving the regime switching model is defining firm expectations over wellhead prices (P) and reserves per well held by the firm (R). The dynamics of wellhead prices and reserves changed drastically around 2000 (Figure 1.1). U.S. natural gas monthly wellhead price data (in dollars per thousand cubic feet: \$/Mcf) are drawn from the U.S. Energy Information Administration (EIA) from January 1976 to December 2012. We completed a Zivot-Andrews unit root test that allows for a single break in the intercept of wellhead price data to determine whether there was a structural break in the time series due to the boom [135]. In January 2000, there was a structural break in the intercept of wellhead prices with a 1% significance level. This result is consistent with [131] and [132]. Therefore, we used price data from January 1976 to December 1999 to characterize price expectations before the fracking boom. We followed Pachamanova and Fabozzi [86] to estimate r_P, $\bar{P}$, and σ_P (Appendix 1). The rate of reversion for wellhead prices is 2.35%, and the long-run mean is \$2.12 per Mcf. Wellhead price volatility is 6.92%.

Firm expectations of reserves per well (in thousand cubic feet per well: Mcf/well) are calculated using EIA data on annual U.S. natural gas (wet after lease separation)

[9] We account for discount rates above and below the baseline model value of 9% in our sensitivity analysis. Our results do not change significantly.

proved reserves and the number of gas and gas condensate wells from 1989 through 1999.[10] Reserves and well count data are specific to Arkansas, Kentucky, Louisiana, Ohio, Oklahoma, Pennsylvania, Texas, and West Virginia to match [132]. These eight states have the largest shale gas plays - Barnett, Marcellus, Haynesville, Eagle Ford, Woodford, Fayetteville, and Devonian shales in the Appalachian basin - and are responsible for over 73% of total U.S. shale gas production in October 2016 [118]. Again, these parameters are estimated using [86] (Appendix 1). The rate of reversion to the long-run mean for proved reserves per well r_R is 0.00013%. The long-run mean $\bar{R}$ is 343,235 Mcf, and the volatility σ_R is 3.00%.

Prices and reserves may be positively or negatively correlated. Negative correlation could reflect increasing prices due to scarcity. However, positive correlation could also reflect increased exploration as prices rise. To capture the correlation between these two stochastic processes, we used the cross-correlation function of the two time series discussed above to calculate the correlation between wellhead prices and proved reserves per well. The correlation coefficient is -0.55 between changes in their levels.

1.3.2 Sunk Costs

[52] contains detailed information on all sunk costs required to drill a natural gas well. There are no publicly available data on sunk costs associated with leasing, acquisition, and permitting for different sized firms that could be used to specify a relationship between K and Q. However, economies of scale are likely. A firm typically negotiates with the landowner for at least one square mile of land, or 640 acres. A well with a depth of 1,000 feet requires one acre of land [85]. That leaves room for far more than one well to be drilled on the leased land. As the number of wells increases, the sunk costs increase but not at a one for one basis. A concave sunk cost function captures these economies of scale and ensures there is a minimum number of wells drilled to trigger entry.[11] For exposition, we assumed the following functional form: $K = b(Q)^{0.75}$. We used the average production per firm for 2000 provided by [132], with the sunk costs of obtaining the mineral rights and permits for drilling one

[10] A Zivot-Andrews unit root test cannot be completed for natural gas proved reserves per well due to a lack of data.

[11] A convex sunk cost curve leads to a firm entering the market up until a maximum firm size based on wells drilled, which is inconsistent with the current natural gas market.

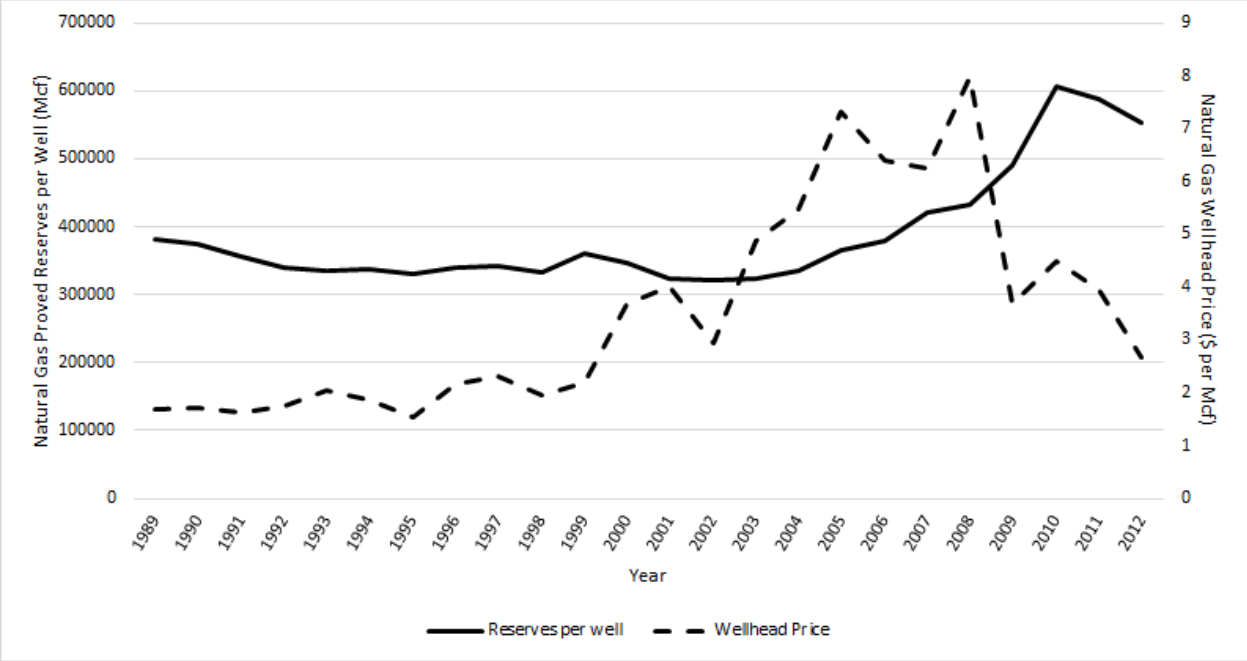

Figure 1.1: U.S. Natural Gas Wellhead Prices and Wet After Lease Separation Proved Reserves per number of gas and gas condensate wells for AR, KY, LA, OH, OK, PA, TX, and WV: 1989-2012.

well from [52], to estimate b. The value for b equals 180.71 and completes the functional relationship between sunk costs and firm production.[12]

The decision to enter the natural gas market requires accounting not only for sunk entry costs, but also the sunk costs associated with exit. Entering a market is inherently tied to the option to subsequently exit that same market. For firms in the natural gas market, the costs of exit are related to site remediation. Included in [52] are estimates for land reclamation costs. Re-contouring portions of the cleared well site and reclamation of temporary roads are part of the remediation costs as well as those associated with restoring topsoil, re-vegetation, and landscaping. Total reclamation fees, F, are on average \$650,000 for one well. Firms that entered the market hold mineral rights and land leases whose value changes over time. Data on how land leases appreciate or depreciate are not publicly available. For the benchmark model, we assumed the lease value just offsets reclamation costs, $sL = F$, so that market exit is costless. For this to be the case, $s = 0.30$, indicating that the lease depreciates by 70% in the benchmark case. We explored alternative values for s in our hypothesis tests.

[12]We varied the exponent between 0 and 1 to account for different degrees of concavity. The results are robust in these changes.

1.3.3 Average Production Costs

To get a measure of average production costs per well in terms of dollars per thousand cubic feet, we collected the annual U.S. nominal cost per natural gas well drilled from the EIA between 1989 and 1999 ($/well). We also collected the annual U.S. natural gas gross withdrawals in thousand cubic feet over the same time frame from the EIA and used the annual number of producing wells for the eight states mentioned above. By dividing gross withdrawals by the number of wells, we obtained a measure of the average production per well (Mcf/well). Taking a ratio of cost per well ($/well) and average production per well (Mcf/well) provides the measure of average production costs per well in terms of dollars per Mcf needed to estimate the $AC(R(t))$ relationship. However, the average cost per well data from EIA includes all costs for drilling and equipping wells for surface-producing facilities - both production and sunk costs.

We removed the portion of average costs not used in the production of the natural gas well by applying criteria explained in [46]. Capital is more sunk when there is a low ability to lease capital, no or limited second-hand markets exist, and if capital depreciates slowly. Using these criteria, we determined which costs included in [52] are sunk. By splitting these costs into production versus sunk costs, we determined the percentage of the costs that are not sunk (57.41%) and multiplied that by the EIA data on nominal costs per natural gas well drilled. This leaves the cost per well in terms of production costs only. Utilizing the apportioned average production cost per well and proved reserves per well data, we found $AC(R)$ is best represented by a power function: $AC = (600,000/R)^{2.634}$ (Figure 1.2).

1.4 Results

According to [132], the natural gas boom was marked by a high volume of small firms entering the market, most intensely from 2005-2008. We focused our attention on the catalysts for small firm entry during the boom. To answer this question, we determined the minimum quantity of natural gas a firm had to produce to enter the market between 1989 and 2012 using pre-boom (prior to 2000) expectations defined by parameters r_P, $\bar{P}$, σ_P, r_R, $\bar{R}$, and σ_R (Table 1.1). These results serve as a pre-boom benchmark; each hypothesis for the increase

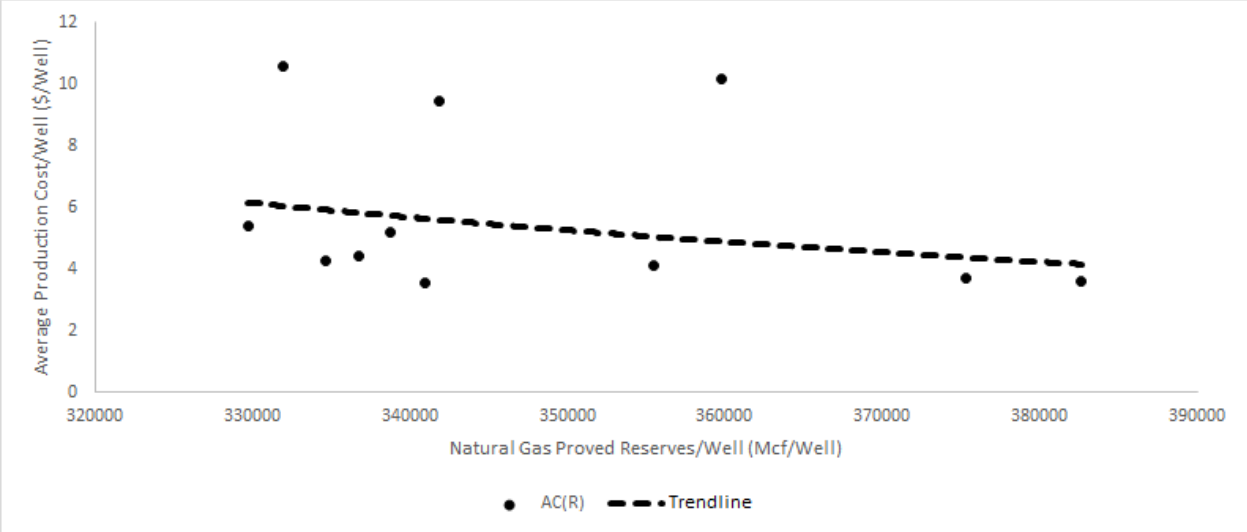

Figure 1.2: U.S. Natural Gas Wet After Lease Separation Proved Reserves per number of gas and gas condensate wells for AR, KY, LA, OH, OK, PA, TX, and WV and Nominal Cost per Natural Gas Well Drilled: 1989-1999.

in small firm entry is tested by altering key model parameters. Significant changes in the minimum natural gas produced signify a driver of small firm entry during the natural gas boom.

1.4.1 The Benchmark: Market Entry with Pre-boom Expectations

Entry thresholds for two total production levels that are the solution to the variational inequalities in (1.3)-(1.6) represent market entry rules for a potential entrant whose expectations of recoverable reserves and wellhead prices are based on historic data (Figure 1.3). The dotted line reflects the entry threshold for a firm producing over 300,000 Mcf per year. To put that into context, we calculated the average rate of production per well from 1989 to 2012 using EIA data on U.S. natural gas gross withdrawals in thousand cubic feet and the number of producing wells for the eight states identified in section 3.[13] A simple ratio of total production over calculations of average production rates provides the minimum number of wells necessary for a firm to enter the market in any given year. The reserves per

[13]Constant average production rates by year are used in the spirit of [5]. They determine that production or flow from wells is constrained due to pressure in underground reservoirs. Thus, the rate of extraction does not play a significant role in determining the output of a drilled well.

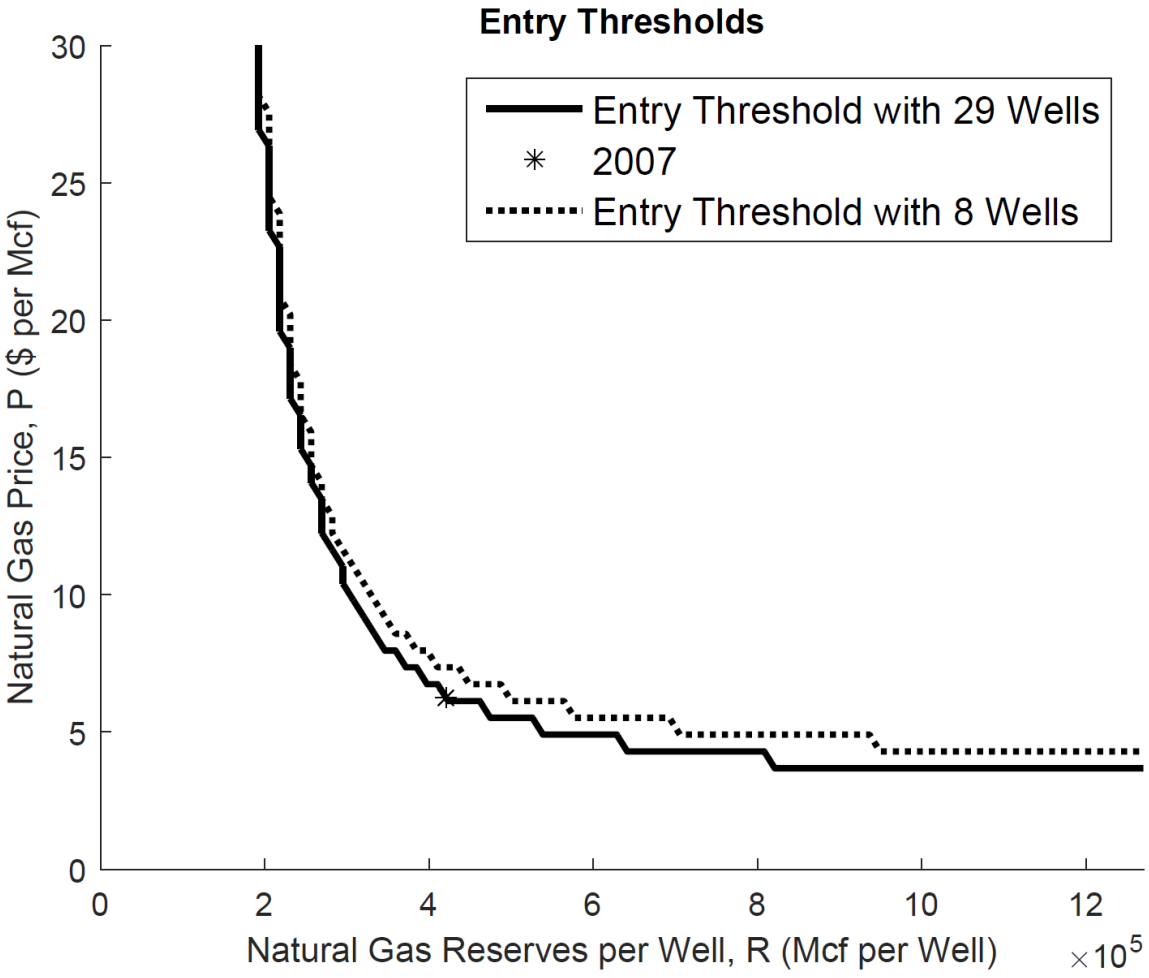

Figure 1.3: Results with baseline expectations: 2007 - The solid line is the entry threshold of a firm drilling 29 wells. The dotted line is the entry threshold of a small firm drilling 8 wells.

Table 1.1: Natural Gas Entry Parameter Values: Pre-Boom Model, Up to 2000

Description	Parameter	Value
Wellhead Price Rate of Reversion	r_P	2.35%
Wellhead Price Long-Run Mean	$\bar{P}$	\$2.12/Mcf
Wellhead Price Volatility	σ_P	6.92%
Proved Reserves Rate of Reversion	r_R	0.000129%
Proved Reserves Long-Run Mean	$\bar{R}$	343, 235 Mcf/well
Correlation Coefficient	ρ	−54.82%
Discount Rate	δ	9.00%
Fixed Exit Cost	$F = 0.30 * 180.71(Q)^{0.75}$	Varies by Q
Variable Exit Cost	$sK = 0.30 * 180.71(Q)^{0.75}$	Varies by Q
Sunk Entry Cost	$K = 180.71(Q)^{0.75}$	Varies by Q
Average Production Costs	$AC(R) = 2\text{x}10^{15}R^{-2.634}$	Varies by R

well and wellhead prices in 2007 are shown by *. Using 2007 as an example, production at the level of the dotted line would require 8 wells given that the average well produced 37,964 Mcf in 2007. [132] demonstrated that new entrants were limited to drilling between one to eight wells. A small firm (with less than 8 wells drilled) would not have found it optimal to enter the natural gas market in 2007 according to our baseline model. A new entrant in 2007 would have needed to produce nearly 1.1 million Mcf in 2007 (29 wells) to justify entering the natural gas market.

We used this same methodology to solve for the number of wells a firm would require to enter the natural gas market in each year from 1989 to 2012 (Table 1.2). Results indicated that firm entry was not likely until 2005 for any size firm. From 2005 to 2008, the minimum number of wells necessary to enter the market trended downward, indicating a potential shift in market structure during this period. The benchmark (pre-boom) model shows that a small firm would have found it optimal to enter the natural gas market in 2008 without any changes to firm expectations or lease appreciation required for entry. Both wellhead prices and average reserves per well were high in 2008. After 2008, the minimum number of wells needed to enter the market began to trend up, although with considerable variation. These results are consistent with available data. While the fracking boom began as early as 2000, few new firms were entering the market early in the shale gas boom. The top four firms in the natural gas market drilled more than 80% of the wells drilled in major shale plays until

Table 1.2: Pre-Boom Benchmark Model - Minimum Firm Size Entering the Natural Gas Market, 2005-2012

Year	Total Production, Q	Average Production per Well	Number of Wells
2005	2,273,180	34,972	65
2006	8,092,382	35,807	226
2007	1,100,956	37,964	29
2008	235,452	39,242	6
2009	41,684,720	38,885	1072
2010	1,426,029	43,213	33
2011	3,129,836	46,027	68
2012	85,207,512	53,288	1599

the early 2000s [132]. However, between 2005 and 2008, the number of active firms increased by 50%. The number of active firms peaked in 2008 at nearly 250 but dropped by almost 38% by 2012.

While this benchmark model captures general trends in the data, it cannot explain the multiple years of small firm entry observed in the data. We investigated three possible hypotheses that may explain the more pronounced entrance of small firms during the boom.

1.4.2 Hypothesis 1: Technological Advances

Technological advances in hydraulic fracturing and horizontal drilling increased $\bar{R}$ in the model by making previously unrecoverable reserves economically viable. Such increases in the long-run mean a reserve level reduces the average cost of production. If small firms believed that the long-run mean level of reserves was 25% higher than historic data indicated due to increased penetration of these technologies, they would begin entering the market. The baseline model is calibrated with the long-run mean at just over 343,000 Mcf per well, which makes the average cost of production, $AC(R)$, \$5.26 per Mcf. A 25% increase in the long-run mean reserve level (429,000 Mcf per well) drives average production costs down to \$2.92 per Mcf. A firm drilling just seven wells would have entered the market between 2005 and 2008 if $\bar{R}$ was about 429,000 Mcf per well (Table 1.3). To compare, average reserves per well were 790,000 Mcf in 2014, meaning that 2014 saw average reserves per well almost double the long-run mean level before the boom. The benchmark model is calibrated with

average production costs 56.6% lower than the average reserves per well in 2014. It is not illogical to conclude that those firms in the 2000s, regardless of size, anticipated long-run average reserve levels would increase.

Another technological advance may have changed firm expectations. During this period, firms began using 3-D seismic imaging to provide a better picture of the structure and properties of rocks below the surface of the ground. This technology served to reduce the variability in well production. To test whether 3-D seismic imaging is a possible explanation for small firm entry, we decreased the value of σ_R, proved reserves per well uncertainty, and solved for the minimum size of firm entrants for 2005-2008, when the largest number of active firms joined the market. The baseline model was calibrated with reserve uncertainty at 3%. The minimum number of wells drilled for a firm that found it optimal to enter in each year as reserve uncertainty is thus reduced by 5, 25, 50, and 75% (Table 1.3). Lowering reserve uncertainty allows smaller firms to enter; however, even when reserve uncertainty was less than one percent, a small firm would not have entered the market during 2005-2008. Reserve uncertainty was low to begin with, so lowering further would not have resulted in a significant change in a firm's incentives to enter the market.

1.4.3 Hypothesis 2: Land Leases

If new entrants in the 2000s were lease speculating, they had to enter the market expecting the value of their land lease would appreciate. This would have allowed them to exit the market with a profit. At the onset of the fracking boom, shale gas plays that were previously not explored for natural gas drilling became the primary focus for firms. The land comprising these plays was wide open for natural gas firms to tap into as long as they acquired the mineral rights from private landowners. Over the course of the boom, the land available for lease from landowners declined. Firms were left to negotiate with other firms, instead of landowners, for land leases. Firms were more able to assert market power to acquire natural gas leases, allowing initial land leases to appreciate in value [129].

The benchmark model assumes costless market exit. A potential entrant expects to recoup 30% ($s = 0.30$) of the lease value, which is just offset by the cost of site remediation F. To test whether lease speculation was responsible for small firm entry, we specified that

Table 1.3: Hypotheses Results, 2005-2008

	Hypothesis 1 $\bar{R}$	Hypothesis 1 σ_R	Hypothesis 2 s	Hypothesis 3 $\bar{P} = \$5.71$
2005				
5%	28	63		
25%	6	59		
50%	4	55		
75%	3	53		
100%			30	
125%			16	
169%				7
2006				
5%	87	220		
25%	7	201		
50%	4	185		
75%	3	174		
100%			104	
125%			51	
169%				8
2007				
5%	15	29		
25%	5	28		
50%	3	27		
75%	3	26		
100%			14	
125%			8	
169%				6
2008				
5%	6	6		
25%	4	6		
50%	3	6		
75%	2	6		
100%			4	
125%			2	
169%				4

leases hold their value without any depreciation (or appreciation), taking $s = 0.30$ to $s = 1$. If a firm expected the value of their lease to remain constant, exiting the market provided a positive payment to the firm. This significantly decreases the minimum firm size required for market entry. The benchmark model indicates that the smallest firm that would have entered in 2007 was drilling 29 wells. The lease price associated with an operation that size is \$6.1 million. In the benchmark model, over \$1.8 million was recouped upon sale of that lease, but site remediation costs F are artificially set to cancel out that payment. However, a lease speculating firm expecting to recoup the full value of the lease $s = 1$ would receive a payment of \$4.3 million for exiting the market.[14] This makes the minimum number of wells for a firm entering the natural gas market in 2007 drop to just 14 wells; there was a similar trend in minimum firm size reduction for 2005-2008 (Table 1.3).

If the lease appreciates in value, the minimum firm size required for entry drops further. An expected 25% appreciation in the lease value at the time of entry would have allowed small firms to enter not only in 2008, but also in 2005 and 2007. Interestingly, we found that once a firm expects to receive a payment just equal to what they paid to enter the market, a firm of any size would find it optimal to enter at any positive wellhead price and reserve per well level. Compared to the benchmark model, a firm would only have to expect the lease value to appreciate by 30% at the time of entry for a small firm to enter in any year.

1.4.4 Hypothesis 3: Prices

Natural gas wellhead prices were relatively stable and hovered between \$1 and \$2 prior to 2000. The results of our pre-boom model are based on a long-run mean wellhead price of \$2.12 per Mcf. When applying wellhead prices from 2000 to 2012, the long-run mean price level increases to \$5.71. New entrants may have viewed the drastic increase in prices after 2000 as a permanent demand shift. These firms would have weighed price data after 2000 relatively more heavily than pre-2000 prices when forming expectations about future wellhead prices (Figure 1.4). Comparing the shift in wellhead price expectations for a firm using historic pre-boom data (1976-1999) and a firm using only price data after 2000, a

[14] The firm would not recoup the entire \$6.1 million because there is still a fixed exit fee F. The benchmark model assumes $sK = 0.30K = F = \$1.8$ million for 2007.

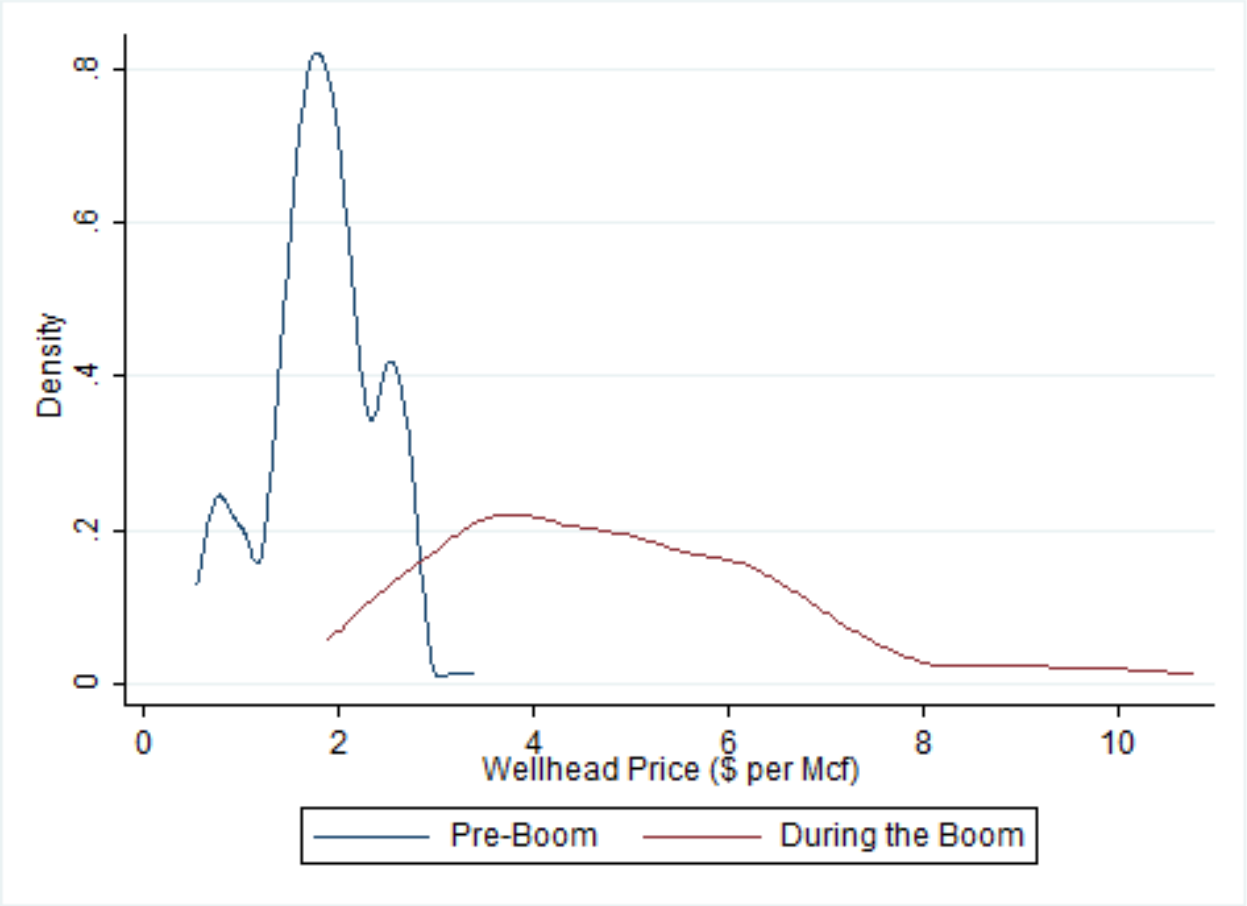

Figure 1.4: Natural Gas Wellhead Price Densities: The pre-boom density and during-the-boom density have different means and standard deviations. The pre-boom density uses EIA data on monthly wellhead prices from 1976-1999, while the during-the-boom density uses data from 2000-2012.

wellhead price of \$6/Mcf would have been viewed as highly unlikely prior to 2000. On the other hand a potential entrant using only more recent price data would have viewed \$6/Mcf as a far more likely outcome.

If firms believed that the long-run mean was \$5.71, small firms would have entered the market from 2005-2008. A firm that formed expectations of future price based only on observed prices after 2000 would have entered the market drilling 7 wells in 2005, 8 wells in 2006, 6 wells in 2007, and 4 wells in 2008 (Table 1.3). This suggests that the wellhead prices observed after 2000 would have been sufficient to entice a small firm to enter the market.

1.5 Conclusion

From the early 2000s, peaking in 2008, to present, the number of active firms in the natural gas market has grown tenfold [132]. New entrants were small in size, which was a drastic change compared to the prior 20 years. If these firms were utilizing historic data and real options theory, our results indicate that they would have found it optimal to enter only at the peak of the boom in 2008. Our results highlight three promising drivers that explain the increased entry of small natural gas firms throughout the 2000s. First, small firms expected an increase in the long-run mean level of natural gas proved reserves due to technological advances. Small firms would find it optimal to enter between 2005 and 2007 if anticipating an increase of about 25% in the long-run mean reserve level from the pre-boom price. In comparison, average reserves per well increased 37% after 2000. Another promising explanation a change in expectations about natural gas wellhead prices. The long-run wellhead price level for years after 2000 was more than double that of historic prices. Small firms that anticipated the regime change in natural gas demand would have found it optimal to enter the market during the boom. Last, small firms could have entered the natural gas market during the boom if they were lease speculating. If a firm believed it could resell a land lease at a price 30% higher than the price paid to acquire the mineral rights, it would have optimally entered drilling just one well for any year during the boom. We found little evidence that the widespread adoption of 3-D seismic imaging could explain the entry of small firms observed in the data between 2000 and 2008. We conclude that a combination of shifts in natural gas proved reserves and prices as well as lease speculating behavior induced small firm entry during the natural gas boom.

Changes in market structure affect not only market power, but a host of other important factors such as firm profits, capitalization, productivity growth, and responses to shocks. If wellhead prices for natural gas and lease speculation were the dominant forces for small firm entry during the boom, the changes would be only transitory. Small firms will move out when larger drilling firms buy out their mineral rights as prices stabilize. If the change in market structure was mainly driven by the shift in natural gas reserves caused by technological advances, then the make-up of the market would have been more permanent. Small firms

will remain in the market. As new data becomes available, it will become more clear whether the change in market structure is permanent or not.

[77] determined that the Henry Hub spot price in the U.S. has a high prevalence of jumps in the data. Incorporating this into our model may lead to the conclusion that natural gas prices have a larger role in the optimal entry decisions of a firm. It should also be noted that the model presented in this paper has several other applications if a stochastic benefit and cost process is involved in decision-making. These topics offer opportunities for future research.

Chapter 2

Drivers of Coal Generator Retirements and their Impact on the Shifting Electricity Generation Portfolio in the U.S.

2.1 Introduction

Coal's share of electricity generation in the United States was over 48% in 2008 and fell to just over 33% in 2015 [121]. Because coal is typically used for baseload generation, this shift in the generation portfolio has been due primarily to utilities choosing to retire coal-fired generators rather than as a result of a decrease in dispatch at these units. These coal-fired generator retirements have consequences for the economy [56, 15] and environment [89, 74, 29, 63, 57, 67, 30]. Regions of the country that rely on coal production for jobs could face increases in unemployment. However, a benefit to retiring coal generators is a reduction in pollution emissions produced by the electricity sector. Formal analyses identifying the forces that drive coal-fired generators into retirement are scarce [39]. Understanding the relative importance of components included in the retirement decision can help isolate what portion of these economic and environmental impacts are attributable to market forces as

opposed to environmental regulation. This research utilizes real options theory and data on coal generator turnover to identify the drivers of coal-fired generator retirement for all coal-fired generators in the U.S.

Coal-fired generators are long-lived assets. Identifying drivers of coal-fired generator retirements is akin to identifying whether the coal-fired generator will exit the market and isolating the factors that shorten the generator's time in the market. According to the National Association of Regulatory Utility Commissioners, engineering estimates of the lifetime of a coal-fired generator account for the time it takes for the net present value of the generator (inclusive of capital costs associated with initial construction) to fall to zero. Once the net present value of operating the generator becomes negative, the generator is assumed to have reached the end of its economic life. These engineering estimates of the lifetime of a coal-fired generator can help identify a set of explanatory variables that could be used in an empirical study of coal retirements.

However, relying solely on econometrics to identify potential drivers of coal-fired generator retirements remains a significant challenge because many of the retirement costs are known only to the utility. News articles and technical reports contain limited and wide-ranging estimates of costs associated with decommissioning but do not account for the cost of offsetting the generation lost to alternative technologies such as natural gas and renewables. These costs may be somewhat offset by the scrap value of machines or the real estate value of land and facilities, but these one-time benefits of retirement are also poorly understood. Retirement cost becomes a key unobservable variable because it acts as a barrier to exit [18, 106].[1] Previous work on market exit attempts to proxy for the sunk costs associated with market exit using a variety of methods. [101] accounted for sunk costs of exiting a market using the ratio of rental payments to assets and primary product specialization ratio, which measures the percentage output actually belonging to an industry. [105] used the proportion of book value of assets that had not yet been depreciated to reflect sunk costs.

Unfortunately, these proxies are unable to account for the full range of costs that influence coal-fired retirement decisions. For example, rental payments and firms' assets will be unable

[1]The literature also suggests lower rates of exit in capital intensive industries [37, 101, 105] and in industries with larger average firm size [37].

to account for decommissioning costs or the costs of offsetting the generation lost from retirement. Previously used proxies will also be unable to account for the option value associated with coal-fired generator retirements [33, 84]. Since many of these retirement costs are sunk, there is an economic value (an option value) to delaying retirement in order to observe the evolution of energy markets and regulation. The irreversibility of the retirement decision and the unpredictability of energy markets can drastically increase the economic lifetime of a coal-fired generator. For example, while the average engineering estimate of the lifetime of a coal-fired generator is 40 years, over 89% of retired coal-fired generators that retired as of 2015 were over 35 years old [117].

Instead of relying on proxies for the costs associated with coal-fired generation retirement, we developed and implemented a real options model of coal-fired generator retirement to back out the net retirement costs implied by the observed timing of 199 retirement decisions between 2009 and 2015. Real options theory [35] treats retirement as an investment option. Utilities and plant operators choose when to incur known retirements costs to permanently shut down a coal-fired generator that is generating an unpredictable flow of profits and losses due to uncertainty in energy markets and regulation. In addition to capturing the influence of irreversibility and uncertainty, a real options approach also allows for heterogeneity across coal-fired generators that can arise due to different electricity market structures. For instance, in our model, coal-fired generator retirement in regulated electricity markets will differ from retirements in deregulated markets due to differences in electricity and coal price expectations [20].[2] Previous applications of real options theory to the electric power sector have focused on market entry decisions such as investments in renewable energy [68, 31], investments in nuclear power plants [102], and electricity plant investments under undefined climate policy [133].[3] To the best of our knowledge, this is the first time real options theory has been used to calculate the costs associated with exiting a market across a large sample of market participants.

[2]Regulated electricity markets feature vertically-integrated utilities that own or control the entire flow of electricity from generation to meter. Deregulated electricity markets divest all ownership in generation and transmission and are only responsible for distribution, operations, and maintenance.

[3][19] provides a comprehensive review of real options theory applications to investments in electricity generation projects.

Since our real options valuation only estimates the retirement costs for generators that have already retired, we employed propensity score matching on observable characteristics from the U.S. Energy Information Administration (EIA), Environmental Protection Agency (EPA), Federal Energy Regulatory Commission (FERC), U.S. Census Bureau, and PJM Interconnection to associate sunk retirement costs with the remaining operational coal generators in the U.S. We then estimated the impact of sunk costs on the probability of retirement using a parametric approach. Our findings indicate the importance of sunk costs in predicting the probability of coal generator retirements for the matched sub-sample. We found that a one standard deviation increase ($45 million) in retirement cost from the mean ($71.6 million) is associated with an increase in retirement probability of 0.2%. Other factors, like nameplate capacity, whether a plant has an ash impoundment, and competing natural gas prices, significantly influence the retirement decision. As a test of our model's predictive power, we compared the predicted probability of retirement for the operating coal fleet to the coal generators reported as retired in the EIA's Form 860 preliminary data for 2016. Almost half of the twenty operating generators with the highest predicted probability of retirement in our sample retired in 2016. When we attempted the same prediction without accounting for sunk costs of retirement, our success rate dropped considerably.

We present our model and describe the data and parameter estimation for the real options model in the next section. Section 3 details the imputed retirement costs for retired coal generators, while Section 4 describes and reports the results for our empirical analysis on the impact of retirement costs on coal generator retirement. Section 5 offers conclusions.

2.2 Coal-Fired Generator Retirement Model

The following illustrates a case where an electricity generation firm (utility) generates electricity with an existing coal-fired generator.[4] The firm receives a flow payoff:

[4]A generator contains all the equipment needed to produce electricity and typically operates independently. Electric power plants can include multiple fuel generators which can use different fuels. For that reason, we conducted our analysis at the generator level. We use the term generator and unit interchangeably.

$$\pi(P_E, P_C) = \Big(P_E(t)q_E(t) - P_C(t)q_C(t) - VC(q_E(t)) - FC \Big) \tag{2.1}$$

where $P_E(t)$ is the wholesale electricity price, $P_C(t)$ is the price of coal, $VC(q_E)$ is the variable operating and maintenance costs, and FC is the fixed levelized capital cost of the generator.[5] q_E is the quantity of electricity supplied by the generator with $\frac{\partial q_E}{\partial P_E} > 0$. q_C is the quantity of coal used to generate q_E with $\frac{\partial q_C}{\partial q_E} > 0$. This relationship between q_C and q_E captures the generation technology of a specific generating unit with newer and more fuel efficient units requiring less fuel to generate an additional unit of electricity.

2.2.1 Price Uncertainty

While current electricity prices and coal prices are known with certainty, future prices are unknown. For example, prices paid for the firm's on-site stock of coal are known, but future prices of coal paid by the firm are unknown. Future coal price uncertainty can be somewhat mitigated by purchasing coal on long-term contracts. Thus, firms that utilize long-term coal contracts will be less exposed to coal price uncertainty than firms that purchase coal on the spot market. [93] determined that energy (coal, crude oil, and natural gas) prices are mean-reverting by testing a century's worth of data. Following [93], future coal prices were treated as random variables and assumed to evolve according to geometric mean reversion (GMR), $\mathrm{d}P_C = r_{P_C}(\bar{P_C} - P_C)P_C\mathrm{d}t + \sigma_{P_C}P_C\mathrm{d}z_{P_C}$. Here, r_{P_C} is the rate of reversion to the mean coal price level, $\bar{P_C}$ is the long-run mean coal price level, and σ_{P_C} is the standard deviation rate. $\mathrm{d}z_{P_C} = \epsilon(t)\sqrt{\mathrm{d}t}$ is the increment of the standard Weiner process, where $\epsilon(t)$ is a standard normal variate. By not reaching 0 in any finite time [65], GMR prevents any negative coal prices. The rate of reversion to the mean, the long-run mean coal price level, and the standard deviation rate are all allowed to vary by generator to capture differences in the types of coal used and the cost of transporting coal to different units. The Weiner process does not vary across generators. This specification assumes coal market uncertainty

[5] Following [10], we define sunk costs as costs that cannot be eliminated even by total cessation of production. In contrast, fixed costs are costs that are not reduced by decreases in output so long as production is not discontinued altogether. Thus, not all sunk costs are fixed and not all fixed costs are sunk.

is universal but responses to coal market shocks are idiosyncratic to the generator. In short, the stochastic differential equation above describes the time-variant coal price distribution facing a particular generator.

Unexpected shifts in supply and demand also influence the prices generators receive for the electricity they produce. For example, electricity demand is sensitive to weather conditions since weather variations lead to large variations in heating and cooling demand. Electricity supply is subject to uncertainty surrounding entry of new and exit of old generating capacities. Uncertainty in both the supply and demand side of electricity markets introduces volatility into electricity prices [61]. To capture this uncertainty, electricity prices evolve randomly around a long-run mean following a geometric mean-reverting process $\mathrm{d}P_E = r_{P_E}(\bar{P_E} - P_E)P_E\mathrm{d}t + \sigma_{P_E}P_E\mathrm{d}z_{P_E}$. The mean reverting process captures the flat load demand in many parts of the country in recent years. Similar to the coal price process, the parameters that govern the stochastic electricity price process are allowed to vary by generator to capture regional differences in regulated and deregulated electricity markets.

2.2.2 The Optimal Timing of Retirement

Based on the expectations of future coal and electricity prices, a coal-fired generator will have a nonzero probability of operating at a loss (i.e. generating negative profits for the firm). At some point in the future, a firm may choose to retire a coal-fired generator when losses become too frequent. Retirement instantly eliminates the flow payoff $\pi(P_E, P_C)$ at some sunk retirement cost, K. Retirement costs can vary widely depending on the level of decommissioning. For instance, retirement costs may be minimal if the generator can be retired while the site is maintained in its current condition with little cleanup needed to meet environmental compliance and ensure safety.[6] In contrast, full decommissioning requires substantial sunk costs associated with dismantling all equipment, demolishing structures, and site clean up including wet and dry disposal areas and coal yards. When a generator is

[6]Environmental compliance includes adhering to the EPA's Interstate Air Pollution Transport Rule, National Emissions Standards for Hazardous Air Pollutants regulations, the industrial waste rule for fossil fuel combustion waste, the Cooling Water Intake Structures rules of the National Pollutant Discharge Elimination System, the Steam Electric Power Generating Effluent Guidelines and Standards, and the most recent Mercury and Air Toxics Standards.

retired, there is no legal requirement to demolish the infrastructure. As long as environmental and safety regulations are followed, the site can remain intact and no decommissioning is required. While the generator is operating, it produces electricity $q_E(t) > 0$ using coal $q_C(t) > 0$ with costs $VC(q_E(t)) > 0$ and $FC > 0$, which generates a flow of profits. If the generator is retired, it produces no electricity and uses no fuel but incurs a sunk retirement cost: $q_E(t) = 0$, $q_C(t) = 0$, $VC(q_E(t)) = 0$, $FC = 0$, and $K > 0$.

The objective of a risk-neutral utility is to determine whether and when to retire an electricity generator, t_R, to maximize the generator's expected discounted profits net of any sunk retirement costs. Using traditional discounted cash-flow analysis, the firm would retire the generator when the expected net present value of generation profits is less than the cost to retire the generator. However, since the costs associated with retirement are sunk, there is an incentive (an option value) to delay retirement longer than suggested by discounted cash-flow analysis. This option value captures the economic value to a firm in being able to respond to new information about coal and electricity markets. The size of this option value is key to determining the timing of coal-plant retirements and will vary by generator depending on coal and electricity market conditions, the efficiency of the coal-fired generation technology currently being utilized, and the sunk costs required to retire.

At each instant in time, the firm must determine whether to continue operating the coal generator or retire it given that all future retirement decisions are made optimally. Given the discount rate δ, the optimal retirement time satisfies the following:

$$V(P_{E_0}, P_{C_0}) = \max_{t_R} \mathrm{E}_0 \left[\int_0^{t_R} \pi\Big(P_E(t), P_C(t)\Big) e^{-\delta t} dt + \Big\{ V\Big(P_E(t_R), P_C(t_R)\Big) - K \Big\} e^{-\delta t_R} \right] \tag{2.2}$$

subject to dP_E, dP_C, $P_E(0) = P_{E_0}$, and $P_C(0) = P_{C_0}$. The evaluation at each instant in time maximizes the expected profits from the coal-fired generator from that point forward by making a choice to continue to generate electricity using coal (whose payoff is defined as V) or to retire and incur K.

Following [35], the retirement decision can be specified as an optimal stopping problem. Treating retirement as an optimal stopping problem will ensure that the retirement decision

maximizes the value of the coal-fired generation asset. While the generator is operating, it not only provides a flow of profits $\pi(P_E, P_C)$, but it also means the firm holds an option - the option to retire the generator when market conditions deteriorate $V(P_E, P_C)$. This option value represents the value of delaying retirement to gain more information about the profitability of coal-fired electricity generation. When the firm retires the generator, sunk retirement costs K are incurred and the option value is terminated - making it an additional opportunity cost of retirement. It is this opportunity cost that causes a more cautious response by the firm in the face of uncertainty.

The generator's unknown value function can be found by employing stochastic dynamic programming with the following Hamilton-Jacobi-Bellman (HJB) equation

$$\delta V(P_E, P_C) \geq \pi(P_E, P_C) + r_{P_E}(\bar{P_E} - P_E)P_E\frac{\partial V(P_E, P_C)}{\partial P_E} + r_{P_C}(\bar{P_C} - P_C)P_C\frac{\partial V(P_E, P_C)}{\partial P_C} + \frac{1}{2}\sigma^2_{P_E}P_E^2\frac{\partial^2 V(P_E, P_C)}{\partial P_E^2} + \frac{1}{2}\sigma^2_{P_C}P_C^2\frac{\partial^2 V(P_E, P_C)}{\partial P_C^2} + \sigma_{P_E}\sigma_{P_C}\rho P_E P_C\frac{\partial^2 V(P_E, P_C)}{\partial P_E \partial P_C} \quad (2.3)$$

ρ is the correlation coefficient between the two stochastic processes $P_E(t)$ and $P_C(t)$: $\rho = corr(\mathrm{d}z_{P_E}, \mathrm{d}z_{P_C})$. In financial terms, the firm faces an obligation to a flow of profits and option value before retirement. The obligation is treated as an asset whose value $V(P_E, P_C)$ must be optimally managed (i.e. maximized). The left-hand side of (2.3) is the return the manager would require to delay retiring the generator over the time interval dt. The right-hand side of (2.3) is the expected return from delaying retirement over the interval dt based on expectations of future coal and electricity prices. This equation acts as an equilibrium condition ensuring a willingness to delay prior to retirement.

The HJB equation in (2.3) is a non-homogenous, second order partial differential equation necessitating numeric methods to approximate the coal retirement value function [81]. Like any differential equation, the HJB equation must be solved subject to a boundary condition. We approximated the solution to the HJB equation by using the well-known value matching condition [35]

$$V(P_E, P_C) = K \quad (2.4)$$

which acts as a boundary condition between the region of the state space where it is optimal to continue operating the coal-fired generator and the region of the state space where it is optimal to retire. We approximated $V(P_E, P_C)$ over a subset of the state space using piecewise linear basis functions [75]. The approximation procedure solves for the 2 x n x m basis function coefficients, which satisfy (2.3) and (2.4) at a set of $n = 50$ and $m = 150$ nodal points spread evenly over the two-dimensional state space.[7]

The combination of (2.3) and (2.4) ensures that the firm holds the option to retire the generator until the value of the option to retire is equal to the cost of retirement. The solution to the HJB equation (2.3) and value matching condition (2.4) can be characterized by a retirement threshold $P_{E_R}(P_C)$ that separates the state space where retirement should occur. Specifically, the retirement curve is the set of points where conditions (2.3) and (2.4) are met. Based on expectations of future profits, the firm optimally retires a generator when a decrease in electricity prices or an increase in fuel prices crosses the threshold curve $P_{E_R}(P_C)$.

A closed-form solution for the retirement threshold only exists in the simplest models where the HJB equation is an ordinary differential equation. With uncertainty in both the coal and electricity prices, the HJB equation becomes a partial differential equation.

2.2.3 Data and Parameter Estimation

The optimal stopping problem described in (2.3) and (2.4) is solved for each coal-fired generator that retired in the U.S. between 2009 and 2015. To identify coal-fired generators that retired as of 2015, we utilized the EIA's 860 data on existing generators. This data identifies generators that retired each year on a cumulative basis. Most retirements in the U.S. as of 2015 occurred in the Northeast, Southeast, and parts of the Midwest (Figure 2.1).

[7] Upwind finite difference approximations were used to construct a linear spline, which approximates the unknown value function. We used Matlab along with the CompEcon Toolbox and the smoothing-Newton root finding method to solve the resulting complementarity problem. The approximated state space ranges from 0 to 15 in the P_C dimension with $n = 50$ nodal points and from 0 to 150 in the P_E dimension with $m = 150$ nodal points. Extending the state space in either the $P_C(t)$ or $P_E(t)$ dimension or increasing the number of nodal points beyond 50 and 150 did not alter our general results.

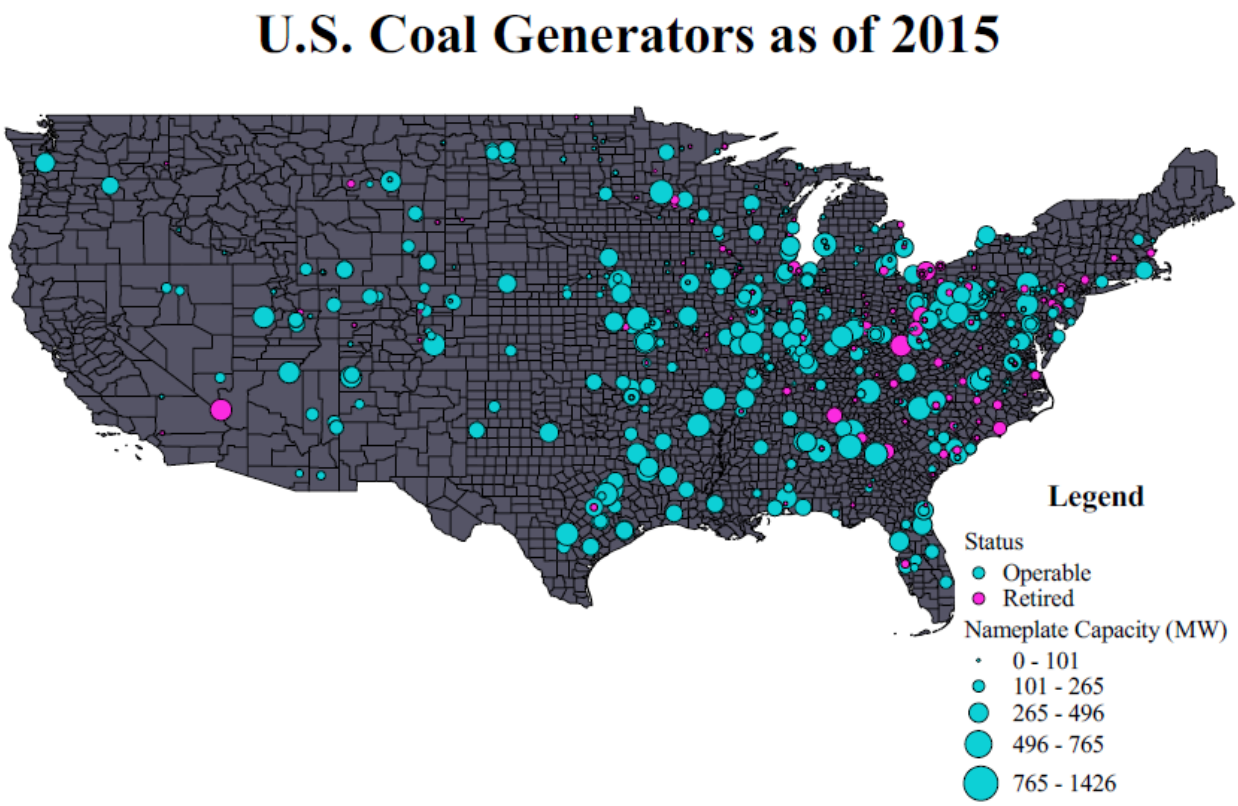

Figure 2.1: Retired vs. Operable Coal Generators

We targeted generators that used some form of coal as their stated main energy source for generation.[8]

We started with a total of 485 retired coal generators as of 2015, restricting retired coal generators to those that have a primary purpose North American Industry Classification System (NAICS) code of 22: electric power generation, transmission, and distribution. For example, the U S Alliance Coosa Pines coal plant retired in 2008, but its sole purpose was to provide electricity for a paper plant. This restriction eliminated generators that produced electricity but were never connected to the electricity grid and dropped the total number of retired coal generators by 77 generators. Uncertainty in electricity and coal prices would play a much smaller role in the decision to retire coal generators not connected to the grid.

A total of 386 coal generators with a primary purpose of electric power generation were reported as retired in the EIA 860 data as of 2015. Of those, over 86% occurred after 2005

[8]Coal includes anthracite, bituminous, lignite, sub-bituminous, waste, refined, and coal-derived synthesis gas.

and almost 26% happened in 2015. We restricted our analysis to coal-fired generators that retired after 2005 because we did not have electricity price data for years prior. In the end, 336 retired coal generators were left in our analysis.

Estimates of the parameters included in (2.3) and (2.4) are required to approximate $V(P_E, P_C)$ over a subset of the state space. While estimating the parameters found in (2.3) and (2.4), we dropped generators that did not have coal or electricity price data that overlapped or used steam load for electricity generation.[9] We also dropped one generator due to lack of data that prevented us from confirming that coal prices delivered to that generator followed geometric mean reversion. We were left with a total of 199 coal generator retirements between 2009 and 2015. Of that total, 56 were in deregulated electricity markets, and the other 143 were in a regulated market. We calibrated the model for each individual coal generator. We followed [55] and assumed firms used a 9% discount rate, $\delta = 0.09$.

Electricity and Coal Prices

A critical step in solving the optimal stopping model is defining firm expectations over electricity prices (P_E) and coal prices (P_C). Coal price and quantity data for plants in our analysis came from the EIA's 923 database, which includes information on monthly fuel receipts such as the quantity and price of fuel delivered to a plant. Since an electricity plant can have more than one fuel delivery per month, we used a weighted average of the fuel-specific quantity and price delivered each month to compile monthly fuel prices ($/MMBtu). Delivered fuel quantities do not accurately describe the fuel used in electricity generation on a monthly basis. Often plants have coal stockpiles, so the quantity of coal delivered to the plant is not all used in the month in which it was delivered.

Real options results critically depend on choosing the correct stochastic process [112]. That process may be one where the price follows a random walk with drift, like geometric Brownian motion, or one where the price reverts back to a trend line, like GMR. Prices that

[9] Of the 336 retired coal generators, 16 reported steam load instead of gross load. Steam load is the rate of steam pressure generated by a unit or source produced by combusting a given heat input of fuel, whereas gross load is the rate of electrical generation of a unit or source produced by combusting a given heat input of fuel. Depending on the generator's technology, it will produce either electrical output or steam output. These two groups are distinctly different, so we focused our attention on generators with electrical output only (Table 2.1).

Table 2.1: Average Coal-Fired Generator Parameters by Market Type

Description	Parameter	Regulated	Deregulated
Coal Price Rate of Reversion	r_{P_C}	10.21%	12.85%
		(6.36)	(17.70)
Coal Price Long-Run Mean	$\bar{P_C}$	\$3.33 per MMBtu	\$2.84 per MMBtu
		(1.05)	(0.52)
Coal Price Volatility	σ_{P_C}	9.88%	11.88%
		(5.42)	(8.13)
Electricity Price Rate of Reversion	r_{P_E}	1.81%	1.84%
		(0.68)	(0.38)
Electricity Price Long-Run Mean	$\bar{P_E}$	\$11.53 per MMBtu	\$13.06 per MMBtu
		(2.43)	(2.59)
Electricity Price Volatility	σ_{P_E}	18.37%	23.28%
		(5.24)	(6.18)
Correlation Coefficient	ρ	-25.5%	-34.36%
		(36.97)	(33.71)
Quantity of Electricity	q_E	$q_E = 17,066P_E$	$q_E = 21,326P_E$
		$(16,566)$	$(21,124)$
Quantity of Coal	q_C	$q_C = 3.07q_E$	$q_C = 2.98q_E$
		(0.39)	(0.29)
Discount Rate	δ	9.00%	9.00%
Variable Costs	$VC(q_E)$	$VC = 2.35q_E$	$VC = 2.35q_E$
Fixed Costs	FC	$FC = 17.58\bar{q_C}$	$FC = 17.58\bar{q_C}$

statistically follow GMR would not be consistent with geometric Brownian motion. Unit root tests provide a platform to determine whether a time series follows a random walk. An augmented Dickey Fuller test was used to check the GMR assumptions for coal prices. We rejected the null hypothesis that the price process follows geometric Brownian motion for each generator in the study (Appendix 2).

With empirical support for our geometric mean reversion assumption, we follow [86] to estimate r_{P_C}, $\bar{P_C}$, and σ_{P_C} using data from 2002-2015 (Appendix 2). The coal-fired generators that retired in a regulated market experienced, on average, less volatile but higher coal prices than those in deregulated markets. The average long-run mean coal price is \$3.33 per MMBtu for regulated coal generators and \$2.84 for the retired coal generators in deregulated markets. Average coal price volatility is 9.88% for regulated markets and 11.88% for deregulated markets. The average rate of reversion to the mean coal price across regulated coal generators is 10.21% and 12.85% for deregulated coal generators. Coal generators that retired in a deregulated market have coal prices that revert back to the long-run mean over 25% faster than their counterparts in a regulated market.

Firm expectations of electricity prices are calculated using historic data on wholesale electricity prices (\$ per MMBtu). For retired coal generators whose balancing authority is PJM, we used the monthly PJM zonal wholesale electricity price data to estimate parameters for $\mathrm{d}P_C$. For all other retired coal generators, we used FERC Form 714 hourly system lambda electricity prices by balancing authority area aggregated to the monthly level. Again, we checked the GMR assumption for electricity prices with unit root tests and rejected that they follow Brownian motion for each balancing authority and PJM zone in our analysis (see Appendix 2). Parameters for r_{P_E}, $\bar{P_E}$, and σ_{P_E} were estimated following Pachamanova and Fabozzi (2011) for the calculations (Appendix 2). Since the electricity price data were either at the balancing authority level or the PJM zone level, the parameters for $\mathrm{d}P_E$ are the same for generators within the same balancing authority or PJM zone. Unlike coal prices, the generators that retired in a regulated market experienced, on average, lower but less volatile electricity prices than those in deregulated markets. Retired coal-fired generators received average long-run electricity prices of \$11.53 per MMBtu in regulated markets and \$13.06 per MMBtu in deregulated markets. The average volatility of electricity prices experienced

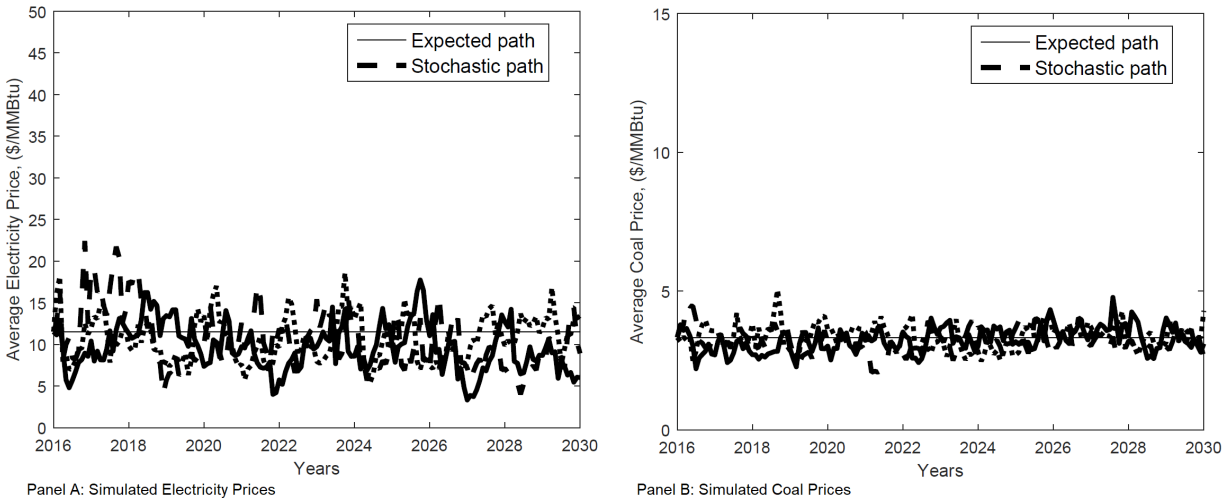

Figure 2.2: Expected and Stochastic Paths for Coal and Electricity Prices. The expected and stochastic paths are simulated using the average dP_E (Panel A) parameters and dP_C (Panel B) for all generators.

by generators in regulated and deregulated markets is 18.37% and is 23.28%. On average, electricity prices revert back to the long-run mean in regulated and deregulated markets at approximately the same rate.

Electricity prices are much more volatile than coal prices (Figure 2.2).[10] The thin horizontal line represents the expected path of electricity prices (Panel A) and coal prices (Panel B) for regulated coal generators given the parameters described above. The dashed, dotted, and thick line represent three possible price paths over the next 13 years. The real options approach captures the effect of the volatility shown in both graphs on the retirement option value.

Because coal is used as an input in electricity generation, there is a clear link between coal and electricity prices. To capture this relationship between the two stochastic processes, we used the cross-correlation function of the two time series discussed above to calculate the correlation between coal prices and electricity prices. The average correlation coefficient for retired coal generators in regulated markets is -0.26 between their levels. In deregulated markets, it is -0.34.

[10]Using the same y-axis for both panels makes the difference between coal price volatility and electricity price volatility appear more extreme.

Electricity and Coal Quantities

The quantity of electricity and fuel at the generator level on a monthly basis is available through the EPA's Continuous Emission Monitoring Systems (CEMS). Electricity quantity data are transformed from megawatt hour to millions of British thermal units. Coal quantity data are reported in MMBtus. We collected monthly EPA CEMS data from 2000 to 2015.

Electricity supplied by a generator q_E follows a simple linear supply function: $q_E = \beta_{P_E} P_E + \epsilon$. We utilized the EPA CEMS data on electricity quantity and the monthly electricity price data described above to estimate the slope of each generator's supply curve using ordinary least squares regressions. We suppressed the constant so a generator would not supply electricity if the price of electricity is zero. The average β_{P_E} for retired coal generators in a regulated market is 17,066, and in a deregulated market, the average β_{P_E} is 21,326. This means that for every 1 dollar per MMBtu increase in electricity prices, the quantity of electricity supplied increases by either 17,066 MMBtus or 21,326 MMBtus. To put that number into context, it takes around 0.064 MMBtus to increase the temperature in a 1600 square foot home with 10 feet ceilings by 50 degrees Fahrenheit. If that house were located in Knoxville, Tennessee, during the winter, it would take only 0.064 MMBtus to heat the home to 70 degrees from 20 degrees. This means that the average quantity of electricity supplied by the retired coal generators in our analysis is highly responsive to changes in electricity prices.

Because coal units are used to generate baseload electricity, owners of coal-fired units seek to minimize the cost of supplying electricity. This means the quantity of electricity generated and the efficiency of the generation technology determines the amount of fuel used. To capture generator technology and efficiency, we used a simple ordinary least squares regression with EPA CEMS data on electricity quantity and fuel quantity (in MMBtus): $q_C = \beta_{q_E} q_E + \epsilon$. Suppressing the constant eliminates the ability of a generator to create electricity without using any coal. This relationship is representative of an inverse production function. The higher β_{q_E}, the less efficient is the unit. For example, a coefficient of 3 for β_{q_E} means that a 1 MMBtu increase in electricity requires 3 additional MMBtus of coal,

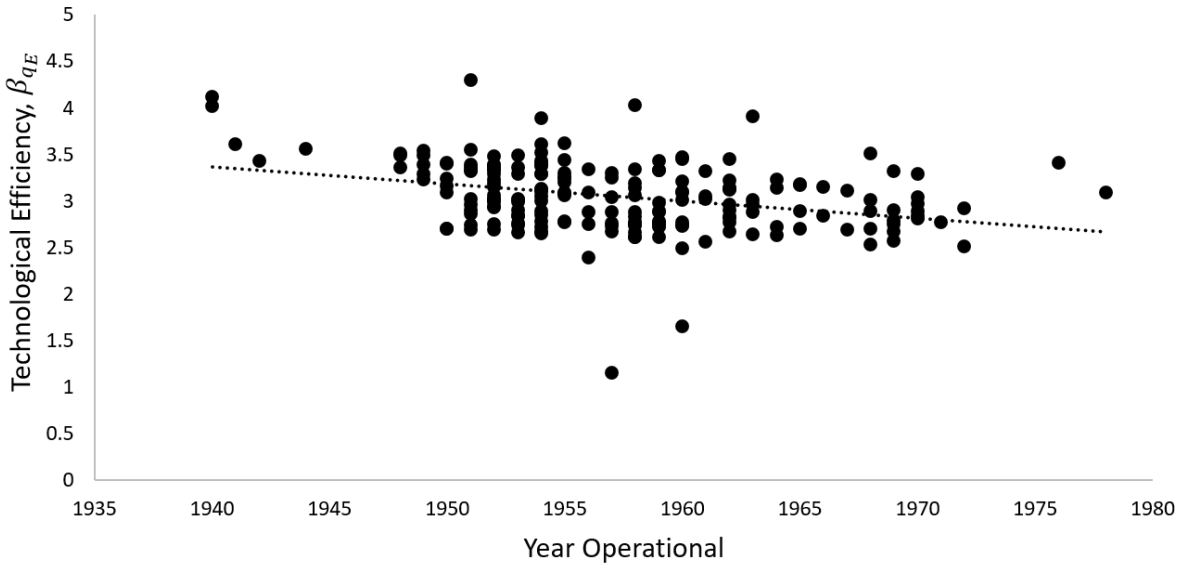

Figure 2.3: Coal Generator Technological Efficiency by Operational Year.

whereas a coefficient of 1 describes a generator that produces the same additional MMBtu of electricity with 2 fewer MMBtus of coal.

An important factor for coal generator retirements is the age of the generator. We were able to account for the effect of age on retirement through β_{q_E} (Figure 2.3). The slope of the inverse production function for each generator mapped against the first operational year of the generator. The negative slope shows that more efficient generators (small β_{q_E}) are associated with younger generators. The average efficiency for retired coal generators in our analysis is 3.07 for those in a regulated market and 2.98 in deregulated markets.

Fixed and Variable Costs

Electricity producers face fixed costs (e.g., capital costs, financing costs) and variable or O&M costs (e.g., fuel, labor, maintenance). According to the EIA, the average operating and maintenance expenses (without fuel) for fossil steam generators between 2005 and 2015 is \$2.35 per MMBtu. The total variable costs associated with each generator in our model is $VC = 2.35q_E$.

We assert that the fixed costs associated with operating a coal-fired generator are equal to the levelized capital costs. Levelized capital costs are calculated by taking the total capital

costs and dividing by the total life of the generator in terms of megawatts of electricity generated. A coal-fired generator that is still operating incurs these levelized capital costs over its productive life (until it retires).

Each year the EIA publishes estimates of levelized costs of electricity for new generation in the Annual Energy Outlook. We took the projected levelized capital costs portion of the estimates for new generation resources in 2019 for conventional coal and transformed it into dollars per MMBtu [116]. The per unit levelized capital costs are $17.58 per MMBtu. Each generator's fixed costs are found by multiplying the per unit levelized capital cost by the average quantity of coal (in MMBtus) used in electricity generation from the EPA CEMS data described above: $FC = 17.58\bar{q_C}$.

2.3 Retirement Cost Analysis

Data on the costs of retiring a coal generator, K, are scarce due to their proprietary nature. Furthermore, retirement costs vary depending on the level of decommission the firm chooses for a generator [123, 54]. If the site is going to be reused for other operations at the power plant, decommission includes removal of equipment and hazardous materials associated with generation. In some cases, full remediation is necessary. The cost and extent of cleanup of hazardous materials depends on the anticipated reuse of the site and the type and location of hazardous materials stored or disposed on the property. If this is the only or last coal generator at the plant, then plants with onsite coal ash ponds or solid waste landfills must follow federal and state permit requirements for closure. The firm can also choose to leave the generator intact and simply maintain environmental permits. Retired generators can be removed and used at other locations the firm owns or sold as scrap.

A variety of news articles and technical reports have published estimates of retirement costs ranging from $4.1 million [54] to $150 million [43]. If a firm's retirement decisions are consistent with real options theory, we can back out an estimate of the retirement costs that justified retiring each generator in our analysis. To do this, we find the electricity price and coal price for the month in which each generator retired (or from the last fuel delivery made to the plant) and compare it to retirement thresholds representative of different retirement

costs.[11] Because coal prices do not vary considerably over our time frame, it is unlikely that using the last coal price for each generator will lead to unreliable estimates of K.

The retirement thresholds, $P_{E_R}(P_C)$, for two randomly selected coal-fired generators are the solution to the optimal stopping problem in (2.3) and (2.4) with specific parameter values (Figure 2.4, Table 2.2). These thresholds represent retirement rules for each coal generator in our analysis whose retirement costs are fully sunk and expectations of electricity prices and coal prices are based on historic data. If the current electricity and coal prices are below a retirement threshold, it is optimal for a firm to retire that generator. Every combination of prices above the threshold is in the continuation region, that is the point at which it is optimal for the generator to remain operational. As expected, retirement becomes optimal when electricity prices fall. How far electricity prices must fall to trigger retirement depends on how much the firm is paying for coal at that generator. Firms will choose to retire a unit at a higher electricity price if they are faced with higher coal prices (a rightward movement along the x-axis); thus, retirement thresholds are upward sloping. A critical retirement cost, $K = K^*$, is found when a retirement threshold goes from being above the electricity-coal price pair from the month-year in which a generator retired (where the firm would find it optimal to retire the unit) to just below the electricity-coal price pair (where the firm would find it optimal to continue operating the generator). To identify K^* for each generator, we varied retirement costs from \$0 to \$160 million in \$500 thousand increments.

The black lines in Figure 2.4 represent retirement thresholds when electricity and coal price volatility is present. The dashed lines are the retirement thresholds without any volatility. The points identify the prevailing electricity prices and coal prices for the month in which each generator retired. Accounting for electricity and coal price uncertainty is important in determining the optimal retirement decision (Figure 2.4). Eliminating the two sources of uncertainty that generators face (the dashed line), a firm would retire the right

[11] One drawback in using the month and year in which a generator retired for coal prices is that plants often receive their last coal delivery months before the generator is technically considered retired. Therefore, it may not be possible observe the coal price for the month in which a generator retires. If a plant operates multiple coal generators and only a subset of these retire, this problem is moot; that plant still receives coal deliveries after the retired generator(s) drops out, meaning the coal price for coal delivered to the plant for use in the remaining operational generators, including the month in which the retired generator(s) leaves the sample. We can use either the coal price from the last delivery made to a plant or the coal price for the month in which a generator retired.

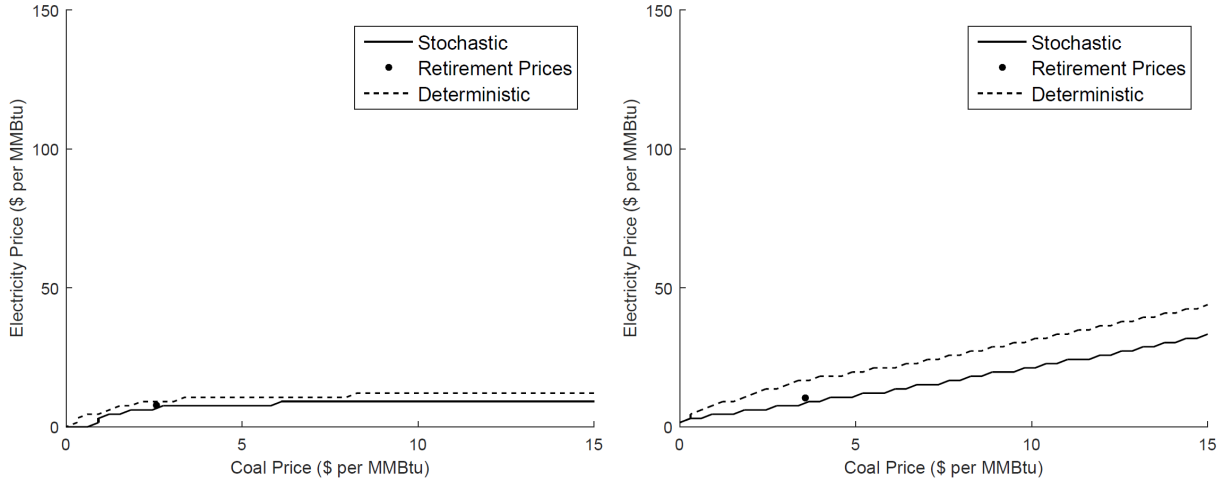

Figure 2.4: Retirement Thresholds for Two Random Generators with and without Electricity and Coal Price Volatility

Table 2.2: Parameter Values for Two Randomly Selected Generators

Description	Parameter	Left Graph	Right Graph
Coal Price Rate of Reversion	r_{P_C}	97.70%	2.80%
Coal Price Long-Run Mean	$\bar{P}_C$	\$2.38 per MMBtu	\$3.93 per MMBtu
Coal Price Volatility	σ_{P_C}	34.50%	10.70%
Electricity Price Rate of Reversion	r_{P_E}	2.54%	1.40%
Electricity Price Long-Run Mean	$\bar{P}_E$	\$9.09 per MMBtu	\$16.35 per MMBtu
Electricity Price Volatility	σ_{P_E}	13.20%	34.20%
Correlation Coefficient	ρ	-37.87%	59.09%
Quantity of Electricity	q_E	$q_E = 16,350P_E$	$q_E = 21,533P_E$
Quantity of Coal	q_C	$q_C = 3.09q_E$	$q_C = 2.82q_E$
Discount Rate	δ	9.00%	9.00%
Variable Costs	$VC(q_E)$	$VC = 2.35q_E$	$VC = 2.35q_E$
Fixed Costs	FC	$FC = 17.58\bar{q_C}$	$FC = 17.58\bar{q_C}$
Critical Sunk Cost	K^*	\$20 million	\$45.5 million

generator at electricity and coal prices much higher than is prescribed by real options theory (the black line). Because there is a positive benefit to delaying retirement (the option value) due to price uncertainty, the black line lies below the dashed line. Real options theory tells us that larger drops in electricity prices are necessary to drive a generator into retirement when those prices are uncertain. The estimates of implied retirement costs are inclusive of the retirement option value. If we did not consider electricity and coal price uncertainty in our analysis, we would over-estimate the implied critical retirement costs. Electricity and coal price volatility have heterogeneous effects across generators (Figure 2.4). The generator's retirement thresholds on the left show minimal changes in the presence of price volatility compared to no volatility.

Figure 2.5 shows the frequency distribution of implied retirement costs across the 199 retired coal-fired generators in our study. The gray (black) bars show the number of deregulated (regulated) coal generators that have retirement costs in each bin. Over 93% of the 199 retired coal generators have critical K values less than \$80 million. Eight coal generators (five in regulated markets) have critical retirement costs equal to \$0. The barriers to retirement were the lowest for these eight coal generators; plant managers had to pay nothing to pull those generators out of the market. On the other end of the distribution, for the most part coal generators that retired in a regulated market have the highest retirement costs. We estimated that the largest coal generator at Widows Creek coal plant in Alabama had the largest retirement costs in our sample, totaling over \$150 million when it retired.

The distribution of critical K values is skewed to the left and takes the general shape of a logistic distribution for generators that retired in either market type. The K^* distribution for regulated coal generators that retired between 2009 and 2015 appears to have a long right tail, and a large mass at \$30-\$40 million. The retirement cost distribution for coal generators that retired in a deregulated electricity market is similar to that of their counterparts, but the highest frequency of retirement costs occurs between \$10-\$20 million. The weighted average of retirement costs for all coal generators in our analysis is \$37.35 million. The retirement costs are a barrier to exit for coal-fired generators and hold potentially informative signals about the likelihood of coal-fired generator retirements (Figure 2.5). In the following section,

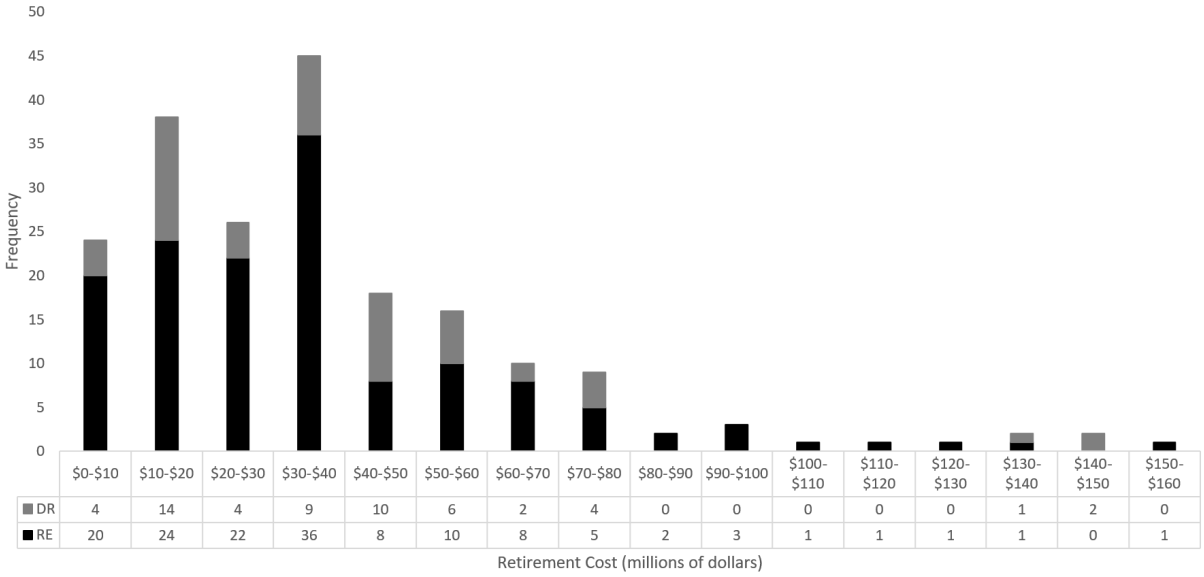

	$0-$10	$10-$20	$20-$30	$30-$40	$40-$50	$50-$60	$60-$70	$70-$80	$80-$90	$90-$100	$100-$110	$110-$120	$120-$130	$130-$140	$140-$150	$150-$160
DR	4	14	4	9	10	6	2	4	0	0	0	0	0	1	2	0
RE	20	24	22	36	8	10	8	5	2	3	1	1	1	1	0	1

Figure 2.5: Retirement Cost Distribution by Market Type

we quantified the improvements in generator retirement prediction afforded by retirement cost estimates.

2.4 Empirical Analysis

We developed a procedure to estimate the retirement costs at active coal-fired generators based on the imputed retirement costs from the real options model described above. Using the data needed to estimate retirement costs, as well as the matching procedure used to generate a sample with common support across retired and active generators, we estimated retirement costs for active plants. While some utilities may have assessed these costs for some of their coal-fired generators, that information is not public.[12]

The result demonstrates that the estimated retirement costs are reasonable and that those retirement costs can be an important input into determining which generators are likely to retire. Analyzing how the retirement costs, derived from fuel and electricity prices, impact actual retirement decisions indicates that our estimated retirement costs contribute to explaining retirement decisions even after controlling for generator observables,

[12] Anecdotally, much of the costs of retiring these generators is unknown even to their owners until they begin the process of evaluating their capacity needs and consider retiring plants.

local sociodemographics and natural gas prices. Simple back-of-the-envelope counterfactual analysis demonstrates the importance of considering retirement costs when predicting the exit of coal fired generators.

2.4.1 Data

Several different data sets were used to estimate the probability of retirement for active coal generators. Generator characteristics from EIA 860 data, local area information from the census, pollution emissions from EPA CEMS data, and fuel prices from EIA build an electric generator panel by the month.[13] The EIA 860 data contains unique identifiers that can be used to link EIA and EPA generator records. Both data sets include the geographic location of the generator, which allowed us to combine generator characteristic data with data from the 2000 Census on the socioeconomic characteristics of the communities where generators were located. We used the 2000 census, rather than contemporaneous data, in order to avoid conflating the impacts of a generator retirement on the sociodemographic status of a county with the impact of the county's sociodemographics on the retirement decision. FERC Form 423 collects monthly data on the electric power price of natural gas and is included this analysis to account for cheaper and more competitive natural gas brought about by the hydraulic fracturing boom in the U.S. (Table 2.3).

[13]Most coal-fired power plants are made up of multiple electric generators. The mean plant in our sample had 3.4 coal-fired generators. Generators can be retired separately or taken down in groups as part of a plant retirement.

Table 2.3: Data Dictionary

Variable	Description	Source
Retire	Indicator for retired generators	EIA 860
Retirement Costs	Imputed cost to retire generator	Imputed from real options model
Efficiency (100%)	Generator efficiency when operating at full capacity	EIA 860
Nameplate Cap (MW)	Generator reported nameplate capacity	EIA 860
Regulated	Indicator for generator in a regulated wholesale electricity market	EIA 860
Ash Impoundment	Indicator for a generator with an ash impoundment	EIA 860
Nonattain	Indicator for a generator located in a county that is regulated under the Clean Air Act for any pollutant	EIA 860
Mercury Control	Indicator for a generator with mercury abatement technology installed in 2009	EIA 860
Median Income	Median county income	2000 Census
Pop. Density	Population Density (pop/sq. mile)	2000 Census
Unemployment Rate	County unemployment rate	2000 Census
Male Higher Ed.	Fraction of males with at least a college degree	2000 Census
Female Higher Ed.	Fraction of females with at least a college degree	2000 Census
Lagged Gas Price	Annual average electric power price for natural gas price lagged 12 months	FERC 423

2.4.2 Estimate Retirement Costs for Active Plants

This section describes the procedure for estimating retirement costs for active generators from the real option model's imputed retirement cost for retired plants. In order to impute the retirement costs described previously, it is important to know the electricity and coal prices for the month in which each generator retires. Because active generators have not yet retired, this information was not available. Thus, we employed propensity score matching to assign retirement costs to operating generators based on a host of observable characteristics. The propensity score is the conditional probability that any generator retires, given the observed generator specific characteristics and sociodemographic characteristics. Specifically, we matched retired coal generators to operating generators based on the variables (except lagged gas price) as well as operational year, utilization rate for 2007, electricity price parameters, coal price parameters, and the total cost of existing abatement technology (Table 2.3). Data on year operational and the total cost of existing abatement technology come from EIA 860.[14] Electricity and coal price parameters were calculated for active generators in the same way as described for the real options analysis above from wholesale electricity prices reported by FERC and PJM and from fuel delivery prices and quantities reported by EIA. The utilization rate for 2007 was calculated by taking the total grossload generated in 2007 (EPA CEMS data) and dividing it by the total capacity (EIA 860) that generator could have produced in a year if it operated at 100% every day of 2007.

The probability that each generator retires were estimated as a function of the variables just described using a logit regression. Propensity scores were calculated as the estimated probability of retiring. Each retired coal generator was matched to an active generator with the closest propensity score. We allowed for matching with replacement in order to reduce the sample size by more than half. For each matched pair, the imputed retirement costs of the retired coal generators was assigned to the matched active generator. The distribution of retirement costs for retired and operating coal generators as of 2015 indicated that coal generators that retired have a large mass where retirement costs are low, and operating coal

[14]The total cost of existing abatement technology is reported in the EIA Form 860 in nominal dollars. We used the producer price index reported by the Federal Reserve Bank of St. Louis to convert nominal dollars into real 1982 dollars.

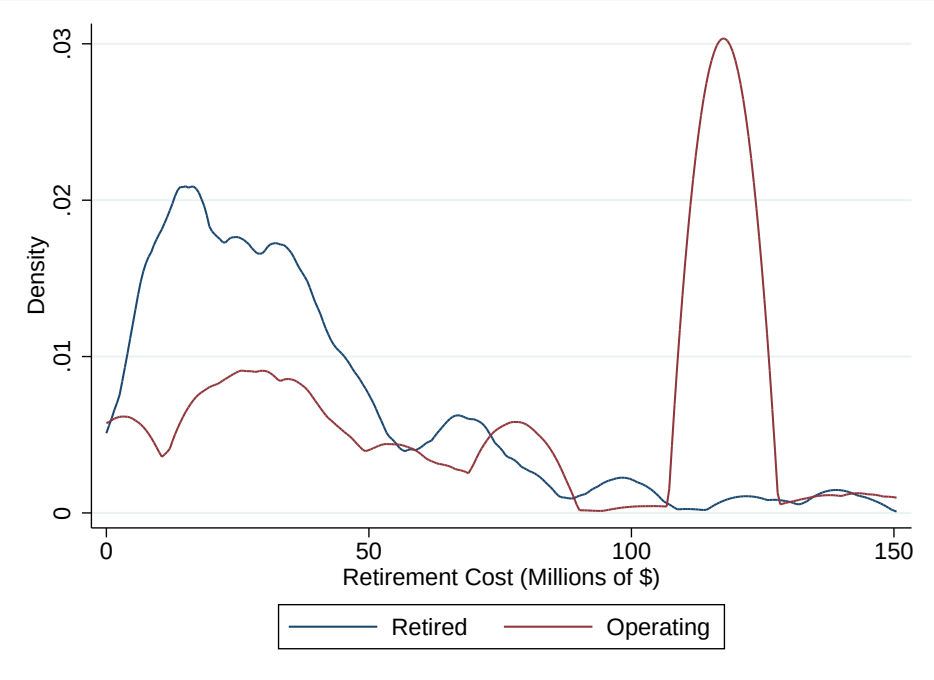

Figure 2.6: Retirement Cost Density

generators have much higher retirement costs (Figure 2.6). This is consistent if active coal generators face higher barriers to exit because the costs of retiring are high.

For reference, the logit model of the probability of each generator retiring takes the form

$$\Pr(Retire_{it} = 1) = \frac{e^{\beta' X}}{1 + e^{\beta' X}} \tag{2.5}$$

where $Retire_{it} = 1$ if generator i retires in time t and is zero otherwise, and X is a matrix of the generator and sociodemographic variables mentioned previously. We used the matched imputed retirement costs in a separate analysis to determine the impact of sunk costs on coal generator retirement.

2.4.3 Analysis of Estimated Retirement Costs

We began by estimating the probability of retirement for each generator in the sample as a function of o imputed retirement costs, generator and local area characteristics. The estimation procedure is:

$$Logit[Prob(Retire_{ist} = 1)] = \alpha RetireCost_i + \beta GenChar_{it} + \gamma Census_i + \theta Gas_t + \delta_s + \zeta_m + \eta_y + e_{ist} \tag{2.6}$$

The dependent variable is an indicator for whether a specific generator (i) is retired in time period t. $RetireCost_i$ is the imputed retirement cost described above, and the parameter α measures how changes in retirement cost affect the probability of retirement all else equal. $GenChar$ is a matrix of generator characteristics including the efficiency of the generator and the nameplate capacity. The $GenChar$ matrix also includes separate indicators for whether the generator is in a regulated wholesale electricity market, has an ash impoundment, has mercury controls, and is in a county that is in non-attainment for any criteria pollution for any year during the sample period. Local sociodemographic conditions may play a role in an operator's decision to retire a generator. For that reason we also included a matrix of variables from the census measured at the county level. The variables include median income, population density, unemployment rate, and the percentage of males and females with bachelor's degrees or higher. We also included lagged natural gas prices to measure competition from other fuel types. Last, we included state (δ_s), month of year (ζ_m), and year of sample (η_y) fixed effects.[15]

The results of this estimation are presented in Table 2.4. Column 1 reports a univariate regression of our imputed retirement costs on the probability of retirement. We found that increases in retirement cost are associated with a reduced probability of retirement, which is evidence that the imputed retirement costs are good proxies for actual retirement costs. In Column 2 state, month, and year fixed effects add to the specification to control for unobserved spatial, seasonal, and time confounds. The estimated impact of retirement costs on the probability of retirement does not change. Column 3 adds generator characteristics; larger generators with ash impoundments are less likely to be retired. Including these generator characteristics moderates the relationship between the imputed retirement cost measure and the probability of retiring somewhat. Column 4 adds county

[15]We did not include the same variables used in the propensity score matching procedure due to multicollinearity issues. For example, the total cost of existing abatement technology and the indicator for having an ash impoundment measure a generator's level of abatement and are highly correlated.

Table 2.4: Estimates of the Probability of Retirement

	1	2	3	4	5
Retirement Costs	-0.02***	-0.02***	-0.01**	-0.01**	-0.01**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Efficiency(100%)			16.60	13.05	13.11
			(13.57)	(14.92)	(14.94)
Nameplate Cap (MW)			-0.01***	-0.01***	-0.01***
			(0.00)	(0.00)	(0.00)
Regulated			-0.59	-0.76	-0.76
			(0.61)	(0.61)	(0.61)
Ash Impoundment			-2.41***	-1.89***	-1.89***
			(0.58)	(0.55)	(0.55)
Nonattainment			0.29	-0.05	-0.05
			(0.43)	(0.47)	(0.47)
Mercury Control			0.42	0.78*	0.78*
			(0.47)	(0.47)	(0.47)
Median Income				0.00	0.00
				(0.00)	(0.00)
Pop. Density				-142.13	-139.82
				(877.47)	(878.27)
Unemployment Rate				17.18	17.12
				(46.43)	(46.45)
Male Higher Ed.				0.00	0.00
				(0.00)	(0.00)
Female Higher Ed.				-0.00	-0.00
				(0.00)	(0.00)
Lagged Gas Price					-0.15**
					(0.06)
Constant	-1.12***	-0.93	-11.55	-10.62	-9.76
	(0.19)	(0.85)	(11.73)	(12.70)	(12.68)
Month FE	No	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes	Yes
State FE	No	Yes	Yes	Yes	Yes
Pseudo R^2	0.086	0.342	0.462	0.470	0.471
Observations	43,056	43,056	43,056	43,056	43,056

Note: The dependent variable in each regression is an indicator equal to 1 if the coal fired generator was retired in that month. Retirement costs were imputed from a real options model as described above. Nameplate capacity through mercury control are generator characteristics collected from EIA 860 data. Median income through female higher education are for the generator's county collected from the 2000 census. Lagged gas prices are monthly national average electric power prices for natural gas as reported by EIA lagged by one year. Standard errors, clustered at the plant level are reported in parentheses below the coefficients. * p<.10, ** p<.05, *** p<.01.

level sociodemographics to the estimation. We found no strong relationship between county characteristics and the probability of generator retirement. Finally, Column 5 adds natural gas prices lagged six months to the estimation. The coefficient is negative, statistically significant, and large in magnitude, suggesting that low natural gas prices increase the probability of retirement.

The impact of estimated retirement costs on the probability of retirement are consistently negative and statistically significant, which holds true even after including a number of the controls used in the estimation of the retirement costs at active generators. This can be considered to be reassuring suggestive evidence that the estimated retirement costs are a good proxy for the actual retirement costs faced by the owners of coal-fired electric generators. Including logit coefficients means the marginal effects of changes in retirement costs are not obvious. To facilitate interpretation of these coefficients, the estimated marginal effect of retirement costs on the probability a coal-fired generator retires is reported across the support of estimated retirement costs (Figure 2.7). The marginal effects of increases in retirement costs are negative across the distribution, but less precisely estimated at low levels of retirement costs. There is a small reduction in the magnitude of the marginal effect of increases in retirement costs on the probability of retirement, but any single ten million increase is not significantly different from the previous. The mean retirement cost across the sample is \$71.6 million and the standard deviation is \$45 million. A one standard deviation increase in retirement cost from the mean is associated with an increase in retirement probability of 0.2%.

The logit estimation described above can generate a predicted probability of generator retirement for all the remaining active generators in the sample. We captured predicted values from the logit estimation presented in Column 3, which includes generator characteristics and fixed effects but not county sociodemographics or fuel prices. The density for retired generators includes all generators that retire during our sample period (Figure 2.8). The operating generator density was estimated on a sample of generators identified as similar to the retired generators across the observable generator and county characteristics reported in the previous section. The operating generators' retirement probabilities are clustered at low

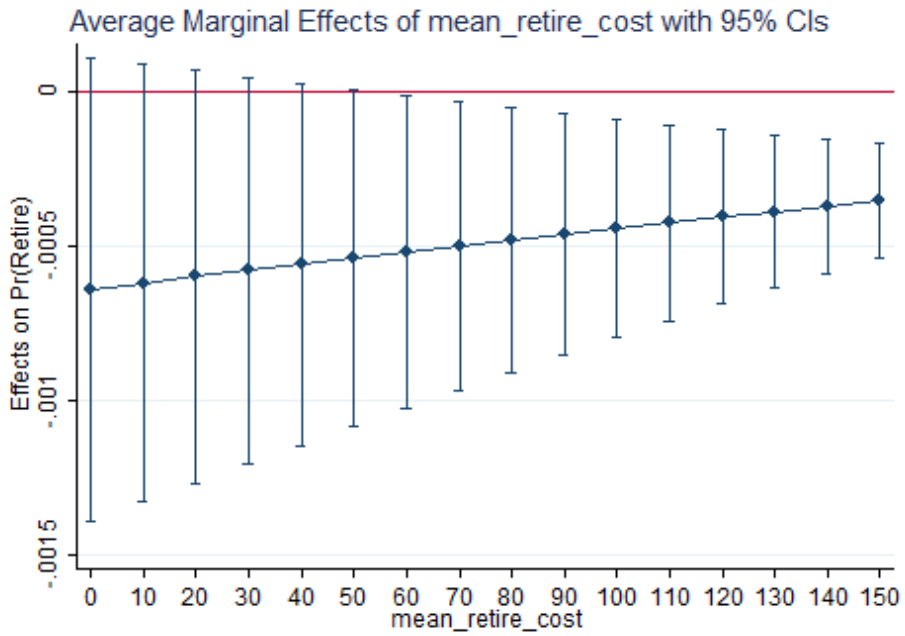

Figure 2.7: Marginal Effect of Retirement Costs

Note: This figure reports the marginal effect of an increase in estimated retirement costs on the probability of retirement as estimated from the logit reported in Column 5 of Table 2.4. The vertical lines represent the 95% confidence interval for that marginal effect estimate. The horizontal axis is measured in millions of dollars. The vertical axis is the marginal effect of retirement costs on the probability of retirement. This is estimated at every $10 million interval from $0 to $150 million.

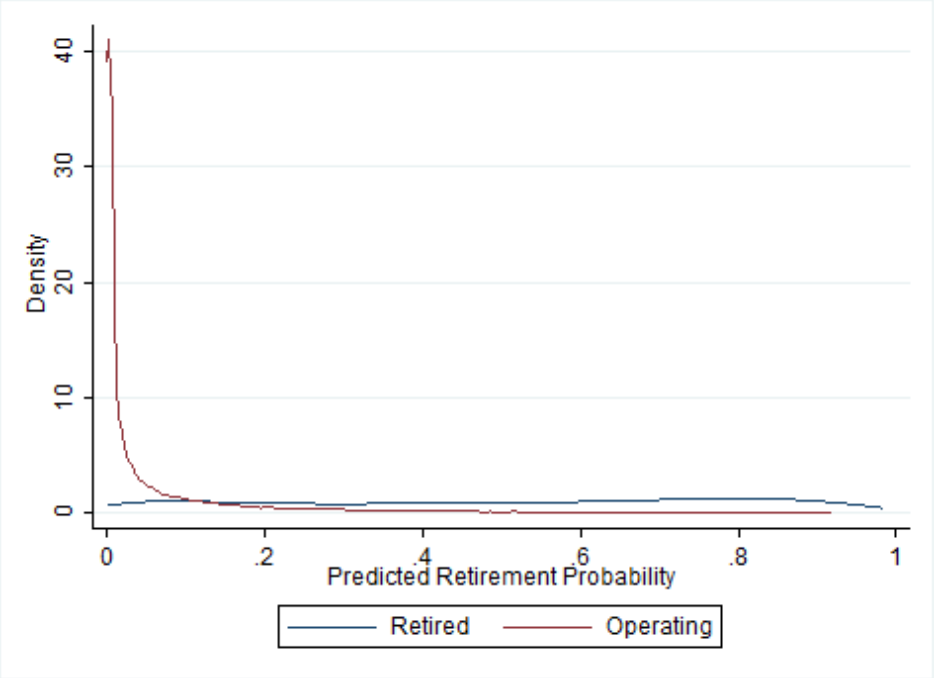

Figure 2.8: Predicted Retirement Density

Note: This figure reports two separate density regressions of the predicted value of retirement: one for generators that retired during our sample period, and a second for generators that survived throughout the sample. All generators that retired during the sample are included in the density of the probability of retirement for retired generators. The operating generators are limited to those that match across observables to generators that retired during the sample period described above.

probabilities. The distribution of retired plants is much more even with a tail of probabilities that are beyond the highest estimated probability for active generators.

As a test of the model's predictive power, we estimated a logit regression using all the variables included in the propensity score matching analysis, as well as retirement costs, and captured the residuals. The residuals represent the predicted probability of retirement for each generator in our analysis. As expected, the generators with the highest predicted probability of retirement are those that in fact retired. More important are the active coal generators with the highest predicted probability of retirement. We compared the top 20 active coal generators with the highest predicted probability of retirement to EIA Form 860's preliminary data on electricity generator retirements for 2016. Of the top 20 retirements predicted based on this analysis, almost half of them did in fact retire in 2016.

To isolate the predictive power of retirement costs, we estimated the same logit regression without retirement costs and collected the residuals. Using those predicted probabilities of retirement, we found the top 20 active coal generators with the highest predicted probability of retirement. They are not the same as those when we included retirement costs, which leads to the conclusion that predicted retirements for active coal generators predicts even fewer actual retirements that occurred in 2016.

2.5 Conclusion

Coal-fired generators have been retiring at an increasing rate since 2009. These retirements have shifted the electricity generation portfolio across the country. In 2015, natural gas surpassed coal as the primary fuel source for electricity generation. This shift has multiple consequences for the economy (increases in unemployment in part of the country reliant on coal production) and the environment (reduction in carbon dioxide emissions produced by the electricity sector). Given the impacts of these shifts in electricity generation, economic analyses are needed to understand the drivers of these retirement decisions. This study explored how the sunk costs associated with retiring a coal generator impact the probability of retirement.

The costs associated with retiring electricity generators are scarcely reported and not publicly available. Instead of proxying for retiring costs with variables used in the past, like the amount of non-depreciated capital, we estimated the retirement costs for almost 200 coal generators that retired between 2009 and 2015 by using a real options model. We matched imputed retirement costs for retired coal generators to active generators using propensity score matching on observable generator and sociodemographic characteristics and then estimated a logistic regression to determine the impact of sunk retirement costs on the probability of retirement. Our results indicate the significant impact retirement costs have in determining market exit for coal generators in the U.S. Furthermore, accounting for retirement costs improves our ability to accurately predict coal generator retirements that occurred in 2016.

To the best of our knowledge, this is the first study to back out estimates of sunk costs of market exit. The retirement decision of the average generator in our study is consistent with retirement costs greater than $37 million. Retirement costs vary across regulated and deregulated markets, across plants within the same market, and within plants. Matching retirement costs of retired generators to active generators indicates that the existing coal fleet still operating in the U.S. have retirement costs greater than $100 million.

Future applications of our real options model and subsequent empirical analysis include incorporating changes in environmental regulation. For example, our framework provides us the opportunity to add environmental regulation like a carbon tax and determine the set of existing coal generators that have the highest risk of being pushed out of the electricity market. Another application would be to create a counterfactual world in which the hydraulic fracturing boom never happened in order to determine whether the coal generators that have already retired would still be operating today. We leave these topics to future research.

Chapter 3

Abate or Exit? The Impact of Mercury Regulation on Coal Generator Retirements

3.1 Introduction

Air emissions have been regulated by the Environmental Protection Agency (EPA) since the passage of the Clean Air Act (CAA) in 1970. Hazardous air pollutants or air toxics are known to be or suspected of causing serious health effects and include mercury, arsenic, chromium, nickel, hydrochloric acid, and hydrofluoric acid. The electric power sector, and coal power generation in particular, has been the largest source of mercury air emissions in the United States since the 1990s [78], but it was not until December 2000 that the EPA deemed it "appropriate and necessary" to regulate fossil fuel-fired electric generating units (i.e. electricity generators) under the CAA.[1] After the D.C. Court of Appeals vacated the first attempt at mercury regulation, the EPA announced the Mercury and Air Toxics Standards (MATS) in December 2011.[2] The final rule establishes numerical emission limits for mercury, particulate matter, and hydrochloric acid for all existing and new coal-fired

[1] An electricity generator contains all the equipment needed to produce electricity and typically operates independently. Electric power plants can include multiple electricity generators that can use different fuels.

[2] A more complete history of mercury regulation in the U.S. is provided in the next section.

electricity generators, which can be satisfied by implementing a range of technologies such as wet and dry scrubbers, dry sorbent injection systems, activated carbon injection systems, and fabric filters [122].

The federal government is not the only source of mercury regulation in the U.S. Twenty-one states had adopted their own set of standards for power plant mercury emissions by the time MATS was announced [79]. State mercury regulations were typically more stringent than the federal regulation. Owners of coal-fired electricity generators had to choose between investing in expensive abatement technology or retiring the generator in order to comply with mercury regulation that was enacted at the state and federal levels from 2000 to present.[3] This papers investigates the impact of environmental regulation on coal generator retirements in the U.S. Much of the existing literature has studied the effect of environmental regulation on productivity [53, 13, 47, 104, 69, 12, 50, 103, 4] and location choices [11, 71, 130] in the manufacturing and industrial sectors.[4] The literature that examines the effect of environmental regulation on exit decisions is considerably less extensive.[5]

[48] utilized the Clean Air Act Amendments of 1970 and 1977 and the Census of Manufactures data to assign manufacturers into regulated and unregulated groups based on their county's regulatory status, their emissions of the regulated pollutant, and year.[6] Identification relied on assigning 1.75 million plants to either a regulated group or unregulated group based on the county in which the plant was located. Each county receives a pollutant-specific regulation designation every year that is determined by the concentration of four controlled pollutants in their air that exceeds a set of national ambient air quality standards. Many industries included in that analysis emitted multiple pollutants and many counties were deemed to be high regulation for multiple pollutants, where a county's pollutant-specific regulatory status changed over time. As a result, the study exploited cross-sectional and longitudinal variation in which plants were regulated and found

[3] For the remainder of the paper, "mercury regulation" refers to both state and federal regulation of mercury.

[4] [32] provided a review of the empirical literature on the impacts of environmental regulation on firm competitiveness. While the focus is on firm competitiveness, the article covers several additional outcomes.

[5] This study categorizes coal generator retirements as a form of firm exit, where each generator is a unique entity to be retired individually from a power plant.

[6] [48] is a previous version of [49] that included a section dedicated to differential impacts on firm births, deaths, and stayers.

that after the Amendments became law, high regulation counties lost jobs, capital stock, and output compared to low regulation counties.[7] Decomposing those effects on births, deaths, and stayers, [48] found that the CAA Amendments' regulations by emission type significantly reduced employment, investment, and shipments for each firm type. [109] used plant-level data to econometrically evaluate the effect of environmental regulation on the adoption and exit behavior of chlorine manufacturers and found that the impact of regulation on technological change was due to the adoption of cleaner technologies at existing plants and new plants, as well as the closing of facilities using older, dirtier technology. The results indicate that regulatory factors have not had a significant effect on adoption for existing plants. However, indirect regulation of the end-users of chlorine appears to have quickened facility closures.

[72] analyzed how the CAA's New Source Review (NSR) requirement impacted modification rates and closure rates in existing plants. Existing plants seeking to modify operations are forced to bring the entire plant into pollution control compliance for new sources with the NSR requirement. The study concluded that substantial modifications of existing pollution-intensive plants in nonattainment counties had been delayed and saw minimal closure rates of existing dirty plants. [134] explored whether underground storage tank (UST) regulations unevenly impacted petroleum retail outlets as well as the reasons that some outlets were more likely to exit the market than others under UST regulations. This study considered whether economies of scale or liquidity constraints could explain the asymmetric impact of UST regulations on petroleum retail outlets. [80] provided a survey of the literature on the effect of environmental regulation on market structure.

The current study differs from the previous literature in several ways. First, it focuses on the impact of environmental regulation on a finer scale entity. Electric power plants typically contain multiple electricity generators that can use various types of fuels to produce electricity. It is more common for individual generators than an entire power plant to be retired. The previous studies noted above papers explored effects at the manufacturing or industrial plant level, whereas this study looks at retirement decisions for individual coal

[7] High regulation counties are those in nonattainment based on the CAA. Those counties have air quality worse than the National Ambient Air Quality Standards defined in the CAA.

generators within a power plant, which can thus exploit within plant-level heterogeneity that the previous literature overlooked. Second, this study examines the direct impact of environmental regulation on coal generator exit. Greenstone (1998) and Snyder et al. (2003) focused on the indirect impact of environmental regulation on plant exit, while List et al. (2004) and Yin et al. (2007) emphasized environmental modifications and pathways for exit, respectively. Additionally, the focus here is on a different type of environmental mercury emission regulation that targets the electric power sector, not the manufacturing or industrial sector.

The analytic sample consists of 1,201 retired and operating coal-fired electricity generators in the U.S. The EPA anticipated approximately 1,100 existing coal generators would be affected by MATS in December 2011 [122], but even as early as 2007, 18 states had established mercury emission limits with an additional four states in the regulation development stage [79].[8] Many coal generators located in these states adopted mercury control technologies before the federal regulation was announced by the EPA in 2011. The earliest mercury regulation at the state or federal level was adopted in June 2003 by Connecticut. Therefore, any coal generator that had mercury abatement technology before 2004 acts as a control group for comparison against the coal generators that were forced to adopt abatement technology or retire due to mercury regulation.[9] A total of 436 coal generators in the sample retired after the start of various state and federal mercury regulations. There are two other potential reasons why these coal generators were pulled offline indefinitely. First, natural gas prices have plummeted since the mid 2000s due to technological advancements in hydraulic fracturing and horizontal drilling making electricity generation fueled by natural gas cost competitive relative to coal. Second, about 73% of all coal-fired capacity was 30 years or older at the end of 2010 [114]. Coal generators that

[8]The states with regulations already enacted in 2007 are Arizona, Colorado, Connecticut, Delaware, Florida, Illinois, Maryland, Massachusetts, Minnesota, Montana, Nevada, New Hampshire, New Jersey, New York, North Carolina, Oregon, Pennsylvania, and Virginia. Georgia, Michigan, Washington, and Wisconsin were in the proposal stage of mercury regulation at the time.

[9]Because Connecticut enacted mercury regulation in June 2003, coal generators assigned to the control group must have adopted abatement technology beforehand. No coal generators adopted abatement technology in 2003. Therefore, all months in 2003 are designated as in the pre-regulation time period to extend the pre-regulation observations.

retired after the start of mercury regulations did so because they had come to the end of their useful life.

In this case, environmental regulation targets coal- and oil-fired electricity generators as they are the primary source of mercury emissions in the U.S. A consequence of such regulation is the retirement of coal generators, where owners of such units have the option to adopt abatement technology or retire the generator. Results indicate that mercury regulation increased the probability a coal generator was retired if it did not have abatement technology before the start of regulation. However, when compared with a group of control coal generators that already had the abatement technology before mercury regulation began, the probability a coal generator was retired due to such regulation is muted. This indicates that in the absence of mercury regulation, the probability a coal generator was retired is not significantly different than the probability a coal generator was retired in the presence of mercury regulation. Instead, cheap natural gas and the vintage of the generator contribute to the probability that a coal generator is retired. Findings are informative for policymakers seeking to understand the response of the electric power sector to environmental regulation. If the goal of mercury regulation was to drive coal generators into retirement, then the current set of policies are not successful. However, if the goal was to force coal generators to adopt abatement technology, the policy could be seen as effective.

The remainder of this paper is organized as follows. The next section provides an overview of mercury regulation in the U.S. Section 3.3 describes the details of the data and empirical method used, and Section 3.4 discusses the impact of MATS on coal generator retirement. Extensions are provided in Section 3.5, which include multiple event study analyses. Lastly, the paper concludes with a discussion and conclusion in Section 3.6.

3.2 Background on Mercury Regulation in the U.S.

As power plants operate, mercury remissions are released into the air. Once mercury from the air settles in water sources, microorganisms transform it into methylmercury, which is a highly toxic form that can build up in fish. The primary source of exposure to humans is through eating contaminated fish, which is a main concern for women of childbearing age,

unborn babies, and young children. Several studies have linked high levels of methylmercury to damage to the developing nervous system [122]. A complete history of federal mercury regulation in the U.S. can be found on the EPA's website: https://www.epa.gov/mats/history-mats-regulation. In addition, a Congressional Research Service (CRS) study provides an overview of state mercury regulation as of February 2007. The summary here is based wholly on [126] and [79], unless otherwise cited.

3.2.1 Federal Regulation History

Under the CAA, several steps must be taken by the EPA in order to regulate air toxics emissions, like mercury, from power plants. Section 112 of the CAA was revised in 1990 to require issuance of technology-based standards for major sources of hazardous air pollutant, where "major sources" include stationary sources like power plants [127]. In October 1994, the EPA entered a settlement agreement with Congress to complete a "Utility Air Toxics Study." The study was to reveal whether it was "appropriate and necessary" to regulate power plants under the CAA section 112. The EPA specifically reviewed mercury emissions from power plants and other industrial sources as well as the health and environmental impacts of such emissions, noting available control technologies. The findings were provided to Congress in December 1997. Two months later, the EPA released its "Utility Air Toxics Study" report to Congress, which considered not just mercury but all potential emissions of toxic air pollutants. The EPA did find it "appropriate and necessary" to regulate coal- and oil-fired electric utilities under the CAA's section 112 and announced this finding in December 2000.

The EPA's finding under the "Utility Air Toxics Study" prompted a requirement for the EPA to propose regulations to control such air toxics emissions. After going through the typical proposal-to-rule steps, the EPA issued the Clean Air Mercury Rule (CAMR) in March 2005. It is important to note that before the CAMR was issued, the EPA removed coal- and oil-fired electricity generators from the Clean Air Act's list of sources of hazardous air pollutants, which are regulated under section 112 of the CAA. Instead, the CAMR was issued under section 111 of the CAA. The key difference between section 111 and section 112 of the CAA is the following: section 111 establishes a mechanism for controlling air

pollution from stationary sources, whereas section 112 establishes emission standards for sources of hazardous air pollutants. The CAMR limited mercury emissions for new and existing utilities by creating a voluntary market-based cap-and-trade program. The D.C. Circuit vacated the CAMR in February 2008 for two reasons: (a) the EPA's removal of coal- and oil-fired electricity generators from the list of sources of hazardous air pollutants inherently goes against section 112, which requires the EPA to make specific findings before removing a source listed under section 112; the EPA conceded that it never made such findings, and (b) because coal-fired electricity generators are listed sources under section 112, regulation of existing coal-fired electricity generators' mercury emissions under section 111 is prohibited [113].

In December 2009, the EPA approved an Information Collection Request (ICR). The ICR required all U.S. power plants with coal- or oil-fired electricity generators to submit emissions information to the EPA for use in creating air toxics emissions standards. With that information, the EPA proposed a rule in March 2011, which would eventually become MATS, setting standards to limit mercury, acid gases, and other toxic pollutants from power plants.

3.2.2 Current Federal Regulation

The rules contained in MATS are technology-based emissions limitation standards. Coal- and oil-fired electricity generators with a capacity of 25 megawatts or greater are subject to MATS [125].[10] Existing sources had up to four years to comply, while three years are provided under the CAA with a potential additional year granted by state permitting authorities. The first deadline for MATS was in April 2015, and the second, final deadline a year later. If a coal generator did not adhere to MATS at that time, it was forced to retire. Figure 3.1 displays coal generator retirements mapped onto the historical timeline of MATS. The black line tracks the number of coal generators that retired in each month from 1993-2016, and the vertical lines represent landmarks in the history of mercury regulation in the U.S.[11]

[10] Burning natural gas for electricity generation emits negligible amounts of mercury [124].

[11] Figure 3.1 contains coal generators that were reported as "retired" in EIA Form 860 for 2016.

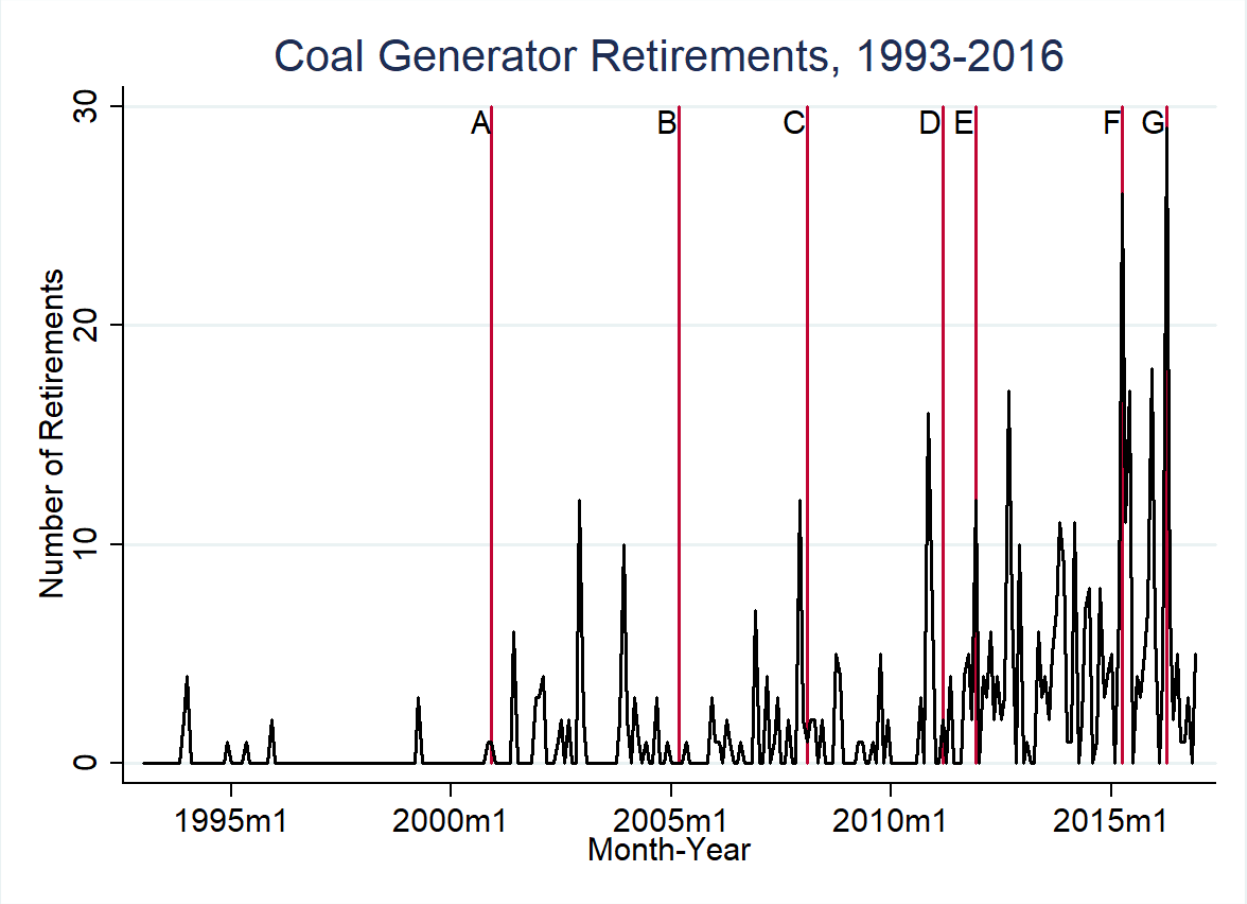

Figure 3.1: Number of coal generator retirements between 1993 and 2016. Data are from the Energy Information Administration Form 860 for 2016. A: EPA deemed it "appropriate and necessary" to regulate coal generators, December 2000. B: Final CAMR issued, March 2005. C: D.C. Circuit vacated CAMR, February 2008. D: EPA proposes MATS, March 2011. E: Final MATS issued, December 2011. F: First MATS deadline, April 2015. G: Final MATS deadline, April 2016.

3.2.3 State Regulation Overview

By February 2007, 18 states had established mercury emission limits that would have taken effect sooner than the EPA's CAMR. Four other states were developing regulations that would do the same: Georgia and Michigan adopted mercury regulations in 2007, Wisconsin in 2008, and Washington in 2011. The effective date of the 18 states that adopted mercury regulation by the time of the CRS study in 2007 ranged from 2007 to 2015. Most state programs require reductions of 80% to 90% in mercury emissions once fully implemented, which are more stringent than those of the CAMR. There is no documentation of a recent examination of state mercury regulations, but the purpose of covering [79] is to highlight the fact that states implemented their own mercury regulations during the study period.

3.3 Data and Methods

Electricity generator data are available from 1990 to present, which covers the time period for which states and the federal government implemented mercury regulations. Therefore, this study used a panel data model that mimics a difference-in-differences identification to determine how owners of coal generators responded to the start of mercury emissions regulation. All coal generators that had already adopted mercury abatement technology before the start of the regulations served as the control group.[12] For the purpose of this study, mercury regulations started in 2003.[13] Observations starting in 2004 were considered to be within the treatment time frame, and observations before 2004 in the pre-treatment time frame. There were no coal generators that installed mercury abatement technology in 2003, so there is no difference in specifying the treatment time frame in 2003 versus 2004.

3.3.1 Data

The EIA collects generator-level specific information about existing and retired generators as well as their associated environmental equipment at electric power plants with 1 megawatt

[12]Difference-in-differences models measure the difference in outcome over time for the treatment group compared to the difference in outcome over time for the control group.

[13]The start date is based on the year the regulation was adopted, not the year the regulation became effective.

or greater of combined nameplate capacity in their Form EIA-860 data [120]. Form EIA-860 is comprehensive, meaning it contains static information on coal generators that retired starting in 1990 up to the current release for 2016. The coal generator retirements with this data, including the month and year of retirement, are identified for this study. Because Form EIA-860 encompasses associated environmental equipment at the boiler level, it is possible to ascertain the year in which a coal generator adopted mercury emissions abatement technology.[14] In addition, the form reports the total nominal cost of all existing abatement technology and provides the year in which it was installed. The year a generator became operational is documented in the same data, which permits calculations of each generator's age.

The EIA also collects detailed electric power data on monthly fuel receipts and costs at the power plant level on their Form EIA-923. However, only data for generators operating in a regulated wholesale electricity market are reported in the publicly available data.[15] Form EIA-923 supersedes Form EIA-423 and FERC-423 starting in 2008, where FERC stands for Federal Energy Regulatory Commission. Coal prices for regulated electricity plants are calculated by using a weighted average of the coal deliveries each month, weighted by the quantity of coal for each delivery.[16] Coal prices for plants operating in a deregulated wholesale electricity market are indexed to the average monthly coal price by North American Electric

[14]Each coal generator has at least one boiler attached to it. The electricity generating process is as follows. Coal is burned within the boiler, which creates heat that turns water within the boiler's pipe system into steam. The steam reaches extremely high temperatures and pressure, which pushes against a turbine shaft. The turbine shaft turns as a result of this pressure and spins the magnets within wire coils of a generator to produce electricity [36]. Mercury abatement technology is associated with a boiler that is attached to a coal generator. Form EIA-860 reports environmental equipment at the boiler level and provides a link between boilers and generators.

[15]Utilities that operate in regulated wholesale electricity markets are responsible for system operations and management and for providing power to retail customers. They are typically vertically integrated in that they own the generation, transmission, and distribution systems. In deregulated wholesale electricity markets, independent system operators operate the transmission system independently of wholesale market participants [128].

[16]Regulated power plants do not receive fuel deliveries each month for coal; therefore, linear interpolation is used to fill in missing data.

Reliability Corporation (NERC) region.[17] The same approach was used to calculate monthly natural gas prices and monthly oil prices.[18]

To control for coal generator retirements driven by electricity prices, state-level average retail electricity prices for all sectors (residential, commercial, industrial, transportation, and other) at the monthly level were included the data came from the EIA's Electricity Data Browser. Other controls included changing yearly state-level natural gas- and oil-fired electricity generator capacity, which came from the EIA's Electric Power Annual.[19] This results in a total sample of 1,201 coal generators at 473 electric power plants with monthly observations between 2001 and 2016 (Table 3.1).

[17]Approximately 34.14% of the coal generators in the current sample operate in a deregulated wholesale electricity market, thus NERC coal prices were used for those coal generators. Linear interpolation was used for NERC regions with missing data. The Northeast Power Coordinating Council (NPCC) does not have coal prices for November and December of 2016.

[18]NPCC does not have natural gas prices for 2001-2007. This matters for 38 (3.16% of the total) coal generators, 32 of which are in the treatment group.

[19]Capacity refers to the maximum amount of electricity a generator can produce in megawatts.

Table 3.1: Data Dictionary

Variable	Description	Units
Retire	Indicator for retired generators	1 or 0
Mercury	Indicator for treatment generators	1 or 0
Age	Difference between current year and generator birth	years
Age^2	Age squared	years
Total Abatement Cost	Monthly installed cost of abatement technology	$1,000
Coal Price	Monthly delivered price of coal at a power plant	$ per MMBtu
Natural Gas Price	Monthly delivered price of natural gas at a power plant	$ per MMBtu
Oil Price	Monthly delivered price of oil at a power plant	$ per MMBtu
Electricity Price	Monthly state average retail price of electricity	¢ per kWh
Natural Gas Capacity	Yearly state existing nameplate capacity for natural gas	MW
Oil Capacity	Yearly state existing nameplate capacity for oil	MW

Note: MMBtu stands for million British thermal units and is defined as the amount of heat required to raise the temperature of one pound of water by one degree Fahrenheit. kWh stands for kilowatt hour is a composite unit of energy equivalent to one kilowatt of power sustained for one hour. MW stands for megawatt and is one million watts, which is a unit of power.

3.3.2 Empirical Strategy

The impact of mercury regulation on coal generator retirement decisions was determined by estimating the following fixed effect model.

$$Retire_{ipst} = \beta_0 + \beta_1 Mercury_{ipst} + P_{pt} + C_{st} + X_{it} + \delta_t + \delta_i + \epsilon_{ipst} \tag{3.1}$$

The dependent variable in Equation 3.1 is a binary indicator equal to one for a coal generator i in plant p in state s in month t that is retired and zero otherwise. $Mercury_{ipst}$ is a binary indicator equal to one for the treatment coal generators that retired after mercury regulations began, all observations after 2003, and zero otherwise. P_{pt} contains controls for variation in fuel prices across plants over time, including monthly coal prices, monthly natural gas prices, and monthly oil prices. C_{st} includes controls for state-level variation in monthly average retail prices of electricity, natural gas capacity, and oil capacity. X_{it} contains generator controls, including the age of the generator, the age of the generator squared, and the total abatement cost. δ_t and δ_i are time fixed effects and generator fixed effects, respectively. All price and cost controls were converted to 2010 dollars using the Consumer Price Index. Robust standard errors are clustered by electricity plant.

The coefficient of interest, β_1, describes the average treatment effect of mercury regulation on generator retirement decisions. The main identification threat is unobserved heterogeneous trends. More specifically, a potential concern is there might be another reason for coal generator retirements that are coincident with the introduction of mercury regulation programs. One potential driver of these coal generator retirements is more competitive natural gas prices over the study period. Technological advances in hydraulic fracturing and horizontal drilling made previously unrecoverable natural gas reserves recoverable in shale formations. This led to the natural gas boom of the 2000s and, consequently, low natural gas prices. Natural gas-fired electricity generation became cost competitive with coal-fired electricity generation. This study addresses this concern by controlling for natural gas prices received at the plant level over time as well as natural gas capacity at the state level over time. Price variables account for the competitiveness of other fuel types, and capacity variables control for changes in the quantity of electricity generation by fuel type.

Using a difference-in-differences identification strategy assumes the control and treatment generators have parallel pre-treatment trends. Section 3.5 presents an event study analysis to confirm that this assumption is consistent with the data.

Because the outcome of interest is a binary indicator for whether a coal generator retired or not, a standard procedure is to utilize a non-linear estimation procedure and assuming a distribution for $Retire_{ipst}$, like a logistic or normal distribution. However, a growing body of literature has pointed out a problem with utilizing a non-linear estimation approach with interaction terms, such as a difference-in-differences setting [3, 98, 64]. $Mercury_{ipst}$ in Equation 3.1 is an interaction term: the interaction between being a treatment or control coal generator and whether the observation occurred before or after mercury regulation. In short, the interaction effect of two independent variables is the cross-derivative of the expected value of the dependent variable. In a linear model, the coefficient on the interaction term, β_1, is exactly this interaction effect. The same is not true for non-linear models due to the transformation of the dependent variable using a distribution bounded between 0 and 1.

[3] provided a clear explanation with mathematical examples using a probit and logit model. [98] provided the calculation of the treatment effect in a nonlinear difference-in-differences model with a strictly monotonic transformation function. [64] extended the work of [3] to difference-in-differences models, models with higher powers of explanatory variables, other nonlinear models, and panel data models. This study utilized a linear probability model in a difference-in-differences setting to avoid the issues described above. The main reluctance to use a linear probability model is that is can lead to predictions outside the 0 to 1 range. This is only a concern when looking at the coefficient on continuous covariates where the linearity assumption matters. With dummy variables and their interactions, like $Retire_{ipst}$, the focus is on the mean differences between the treatment and control group, and hence, no over- or under-prediction. However, when continuous covariates not bounded between 0 and 1 are introduced, the linear probability model can predict retirement probabilities outside the 0 to 1 range.

3.3.3 Descriptive Statistics

Summary statistics are presented separately for treatment (1,155) and control (46) coal generators (Table 3.2). Treatment and control generators are comparable in terms of generator specific control variables as well as the outcome of interest, retirement. For the last month in the panel data, treatment coal generators were slightly older than control coal generators, where this difference is significant at the 10% level. Both age squared and total abatement cost for December 2016 are not statistically different between the two groups. Treatment and control generators face similar coal prices but encounter significantly different natural gas and oil prices at the plant level. The same is true for state level average retail electricity prices and oil capacity in the state, but natural gas capacity in a state for December 2016 was similar for treatment and control coal generators.

3.4 Results

Equation 3.1 is estimated where the control group consists of coal generators that had mercury abatement technology before the start of mercury regulation at either the federal or state level, and the treatment group consisted of coal generators that had to adopt abatement technology or retire due to various forms of mercury regulation. The data covers 1,201 coal generators (1,155 treatment and 46 control) from 2001 to 2016 on a monthly time step and is a strongly balanced panel. The results are shown in Table 3.3, and robust standard errors, clustered by plant, are in parentheses. Column I lists the estimate of the effect of mercury regulation on the probability of coal generator retirement without generator, plant, or state controls. Columns II through IV add in each consecutive set of controls.

Estimates from the basic difference-in-differences model in Column I show that coal generators without mercury abatement technology before regulation have a 0.09 percentage point increase in the probability of retirement after mercury regulation relative to what that probability would have been if that coal generator had already had mercury abatement technology. Adding in generator controls (age, age squared, and total abatement costs) removes all significance on the effect of mercury regulation on the probability of coal

Table 3.2: Summary statistics for generator outcome, generator controls, plant controls, and state controls in December 2016

	I Treatment Generators	II Control Generators
Retire	0.37	0.28
	(0.48)	(0.46)
Age	48.00	43.80
	(15.60)	(19.19)
Age^2	2547.10	2279.20
	(1416.07)	(1892.98)
Total Abatement Cost	194.96	0
	(4887.16)	(0)
Coal Price	2.22	2.26
	(0.40)	(0.34)
Natural Gas Price	5.61	8.51
	(2.43)	(10.13)
Oil Price	11.40	11.78
	(0.70)	(0.72)
Electricity Price	8.65	9.43
	(1.21)	(1.95)
Natural Gas Capacity	13022.86	13722.21
	(14655.35)	(10450.98)
Oil Capacity	865.15	1673.52
	(1093.11)	(1691.14)
Observations	1,155	46

Note: Columns I and II list the mean of generator outcome variable, generator controls, plant controls, and state controls in December 2016 for treatment and control coal generators, respectively. All financial variables are in 2010 dollars. Standard deviations are in parentheses.

generator retirement and lowers the size of the estimate. This holds with the addition of plant and state controls.

An often cited reason for the retirement of coal generators in the 2000s and 2010s is that most coal generators are old. Accounting for non-linear effects of age on coal generator retirement probability reveals that when a coal generator is first operational, the probability that it is retired decreases by 0.03 percentage points for an additional year of operation. However, as that coal generator continues to age, the effect on the probability of retirement switches signs, meaning that at some point, an additional year of operation increases the probability that a coal generator is retired. Figure 3.2 displays this relationship between age and the predicted probability of retirement from estimating Equation 3.1.[20] This impact is significant and consistent even after adding in plant and state controls.

Column III explores one other potential reason that coal generator retirements increased during the study period. The hydraulic fracturing boom brought with it cheap natural gas prices, which made electricity generation fueled by natural gas competitive with coal-fueled generation. A $1 per MMBtu decrease in the price of natural gas is associated with an increase in the probability a coal generator is retired by 0.00005 percentage points. To further show the significant impact cheap natural gas had on coal generator retirement probabilities, Column IV displays that a one megawatt increase in natural gas capacity within a coal generator's state is consistent with a 0.00001 percentage point increase in the probability a coal generator is retired.

The price of oil is positively related to the probability of coal generator retirement, shown in columns III and IV. Oil is typically used as a "startup fuel" for coal-fired electricity generation, meaning that oil is used to initiate combustion that will eventually burn coal for electricity generation. In this sense, oil can be thought of as a complement in production/generation to coal. An increase in the price of oil by $1 per MMBtu is associated with an increase in the probability a coal generator is retired by 0.002 percentage points. However, additional oil capacity on the grid should not be seen as a complement to coal-fired electricity generation; a startup fuel is very different from the primary fuel used in

[20] One potential problem with using a linear probability model with a binary outcome is that linear predictions are not bounded between 0 and 1. This is shown in Figure 3.2, where the marginal effect of age on the probability of retirement dips below 0.

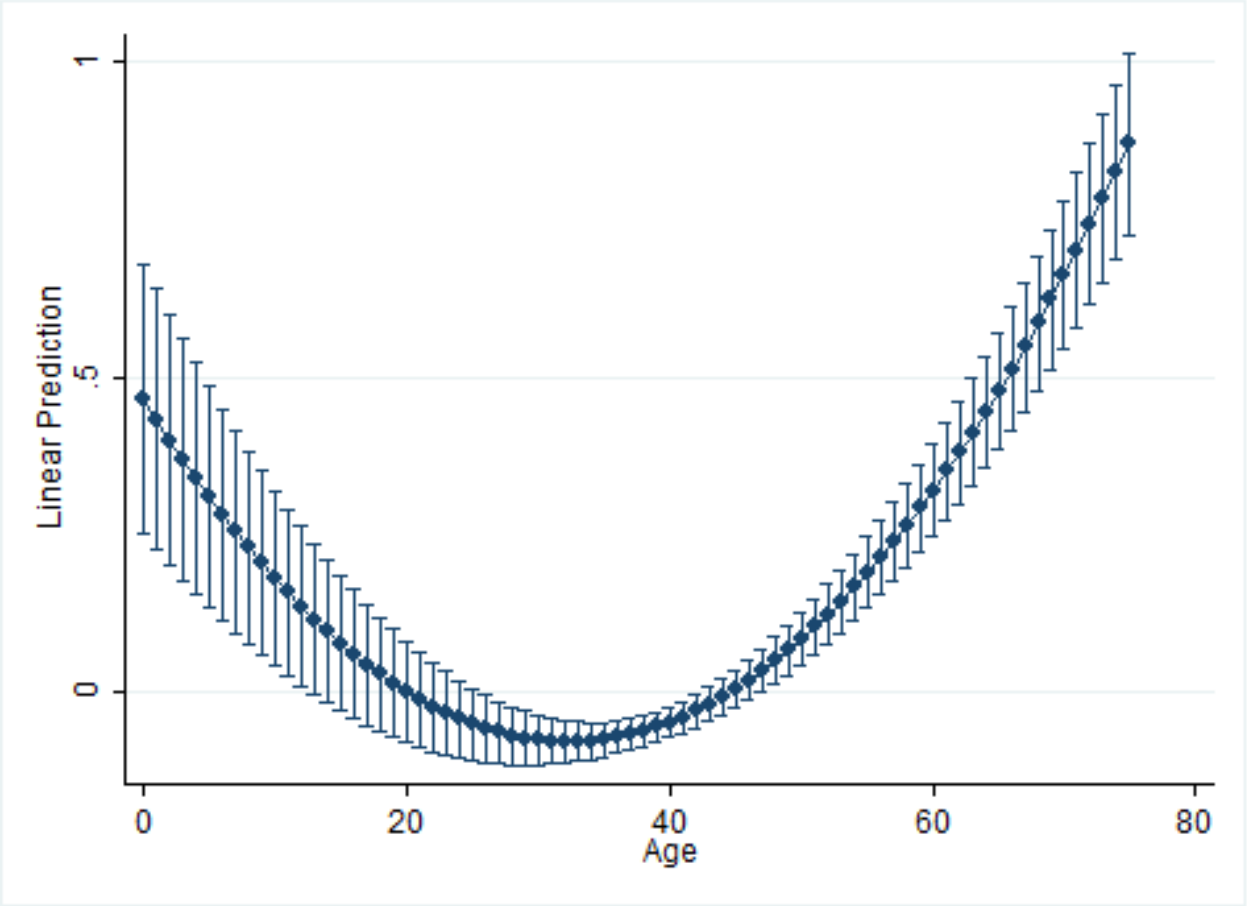

Figure 3.2: The Marginal Effect of Age on Coal Generator Retirement

Note: Figure 3.2 plots the marginal effect of age on the probability of coal generator retirement from the first year operational to 75 years old based on estimating Equation 3.1. Age is in years. The vertical lines show the 95% confidence intervals.

electricity generation. The more oil capacity within a state, the higher the probability that a coal generator is retired. While the price of oil affects the probability of retirement in one direction, the capacity of oil in the state affects it in the opposite direction. The owner of a coal generator has to pay for coal and oil (if that is the startup fuel of choice) to produce electricity from a coal generator, so a higher input oil price is associated with an increase in the probability of retirement. On the other hand, the more generating capacity fueled by oil within a state also leads to an increase in the probability a coal generator is retired. This is due to the fact that oil-fired electricity generation is crowding-out coal-fired electricity generation. The difference between oil prices and oil capacity is that one is a complement in production (startup fuel) and the other is a substitute for generation (primary fuel).

The sign on the effect of coal price is perplexing. Theory would say that the higher the fuel cost, the lower the profits, the higher the probability of coal generator retirement. Because oil is considered a complement of generation for a coal generator, Column III could show oil prices picking up the impact of coal prices. However, estimating Equation 3.1 without oil price still results in a negative coefficient on coal prices. All other controls statistically help explain coal generator retirement decisions except for total abatement cost. This is potentially due to the way in which the control is measured as the total abatement cost installed in a given month. A different way to account for environmental equipment is to measure the total cost of all abatement technology for some lagged aggregate time frame.

3.5 Extensions

3.5.1 Event Study Analysis

The difference-in-differences identification strategy used to estimate the impact of mercury regulation on coal generator retirement assumes that the treatment and control generators have parallel trends before any mercury regulation is implemented. An event study analysis is presented here in order to identify differences in trends in addition to levels between treatment and control coal generators [1]. Following [38], this study estimated the following

Table 3.3: Estimates of the Probability of Retirement

	I	II	III	IV
Mercury	0.08898***	0.04504	0.03997	0.04331
	(0.01)	(0.04)	(0.04)	(0.04)
Age		-0.03484***	-0.03261***	-0.03368***
		(0.00)	(0.00)	(0.00)
Age Squared		0.00053***	0.00051***	0.00052***
		(0.00)	(0.00)	(0.00)
Total Abatement Cost		0.00000	0.00000	0.00000
		(0.00)	(0.00)	(0.00)
Coal Price			-0.04544***	-0.04647***
			(0.01)	(0.01)
Natural Gas Price			-0.00005**	-0.00005**
			(0.00)	(0.00)
Oil Price			0.00227*	0.00241*
			(0.00)	(0.00)
Electricity Price				-0.01381***
				(0.01)
Natural Gas Capacity				0.00001***
				(0.00)
Oil Capacity				0.00002***
				(0.00)
R^2	0.199	0.293	0.292	0.296
Observations	230,592	230,592	227,324	227,324

Note: The dependent variable in each regression is an indicator equal to 1 if the coal-fired electricity generator was retired in that month. Standard errors, clustered at the plant level are reported in parentheses below the coefficients. * p<.10, ** p<.05, *** p<.01.

flexible specification that looks for breaks in any pre-existing differences between treatment and control coal generators, especially at the time of mercury regulation execution.

$$Retire_{ipst} = \lambda_0 + \sum_{d=-36}^{36} 1(t = t_m + d) \cdot EverMercury_{ips} \cdot \lambda_d + P_{pt} + C_{st} + X_{it} + \delta_t + \delta_i + \epsilon_{ipst} \quad (3.2)$$

The month that mercury regulation was implemented is represented by t_m. $EverMercury_{ips}$ is a binary indicator equal to one for all treatment coal generators, those that did not have mercury abatement technology before January 2004, and is interacted with indicators for up to 36 months prior and since the start of mercury regulation.[21] All other variables including plant controls, state controls, generator controls, generator fixed effects, and time fixed effects are defined as in Equation 3.1. The flexible specification shows the differences in trends between treatment and control coal generators. Identification relies on the assumption that pre-existing trends would have persisted in the absence of mercury regulation. The coefficients of interest are the λ_d's, which describe the changes in trends for generator-level outcomes for treatment coal generators relative to control coal generators.

Figure 3.3 maps the λ_d's against 36 months before and after the start of mercury regulation. Prior to mercury regulation, trends for treatment coal generators were flat relative to control coal generators. This indicates that pre-treatment trends were not significantly different between treatment coal generators and control coal generators. After mercury regulation, this pattern continues for the next three years. It is important to note that Figure 3.3 stops before the end of the panel in order to be symmetrical to the time frame before mercury regulation. The flexible event study results do not provide strong evidence that there are significant changes in trends around mercury regulation implementation. Continuing in the spirit of [38], this study estimated the following parametric event study specification to compare to the existing results.

$$Retire_{ipst} = \alpha_0 + \alpha_1 \cdot (t - t_m) \cdot Mercury_{ipst} + P_{pt} + C_{st} + X_{it} + \delta_t + \delta_i + \epsilon_{ipst} \quad (3.3)$$

[21] The first three years of the panel are pre-treatment/pre-mercury regulation. Because the data are on a monthly time-step, this translates to 36 months before the start of mercury regulation.

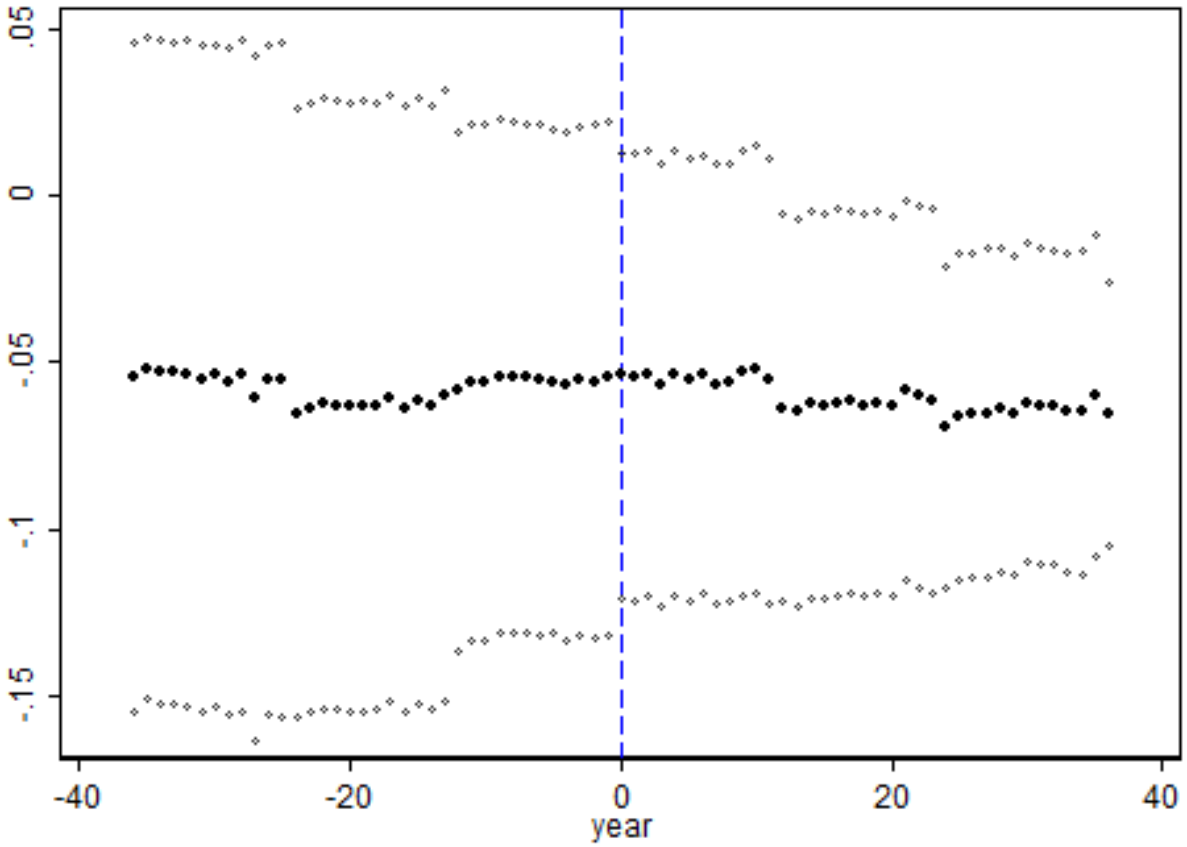

Figure 3.3: Flexible Event Study Estimates

Note: Figure 3.3 graphs λ_d's from estimating Equation 3.2 for coal generator retirement decisions. The vertical line at year zero indicates the year that mercury regulation was introduced. The grey diamonds show the 95% confidence intervals. Robust standard errors are clustered by plant.

Table 3.4: Parametric Event Study Estimates

	I	II
Mercury	0.04331	0.00083**
	(0.04)	(0.00)
Age	-0.03368***	-0.03314***
	(0.00)	(0.00)
Age Squared	0.00052***	0.00052***
	(0.00)	(0.00)
Total Abatement Cost	0.00000	0.00000
	(0.00)	(0.00)
Coal Price	-0.04647***	-0.04639***
	(0.01)	(0.01)
Natural Gas Price	-0.00005**	-0.00005**
	(0.00)	(0.00)
Oil Price	0.00241*	0.00239*
	(0.00)	(0.00)
Electricity Price	-0.01381***	-0.01379***
	(0.01)	(0.01)
Natural Gas Capacity	0.00001***	0.00001***
	(0.00)	(0.00)
Oil Capacity	0.00002***	0.00002**
	(0.00)	(0.00)
R^2	0.296	0.297
Observations	227,324	227,324

Note: The table lists the event study estimates of the effect of mercury regulation on coal generator retirement decisions. Column I lists the baseline estimates from Equation 3.1. Column II lists the estimates from Equation 3.3. Robust standard errors, clustered at the plant level, are reported in parentheses below the coefficients. * p<.10, ** p<.05, *** p<.01.

$Mercury_{ipst}$ is a binary indicator equal to one for treatment coal generators after mercury regulation implementation and zero otherwise. The gap between month t and the month that mercury regulation was implemented is represented by $(t-t_m)$. All other variables are defined as in Equations 3.1 and 3.2. The coefficient of interest is α_1, which describes the change in the slope of coal generator retirement probabilities after the introduction of mercury regulation. The results are presented in Table 3.4. Column I lists baseline results from estimating Equation 3.1, while Column II lists the event study results from estimating Equation 3.3. Results from the more parametric event study analysis are similar to the baseline results for all control variables. However, the impact of mercury regulation on coal generator retirement

decisions is significant in the event study analysis. The key difference between the baseline, difference-in-differences approach, and the event study is the identification strategy. The event study coefficient of interest, α_1, is significant if the slope of coal generator retirement probabilities changed after mercury regulation was implemented, where the difference-in-differences strategy relies on changes in coal generator retirement probabilities between a treatment and control group. The event study does not use a control group to account for aggregate effects that have nothing to do with mercury regulation. The results indicate that mercury regulation significantly increased the probability of coal generator retirement, but after controlling for the impact of mercury regulation on a set of coal generators that were, in a sense, not subject to the regulation, mercury regulation is not found to have causally increased the probability of coal generator retirements for those that were forced to abate or retire (Table 3.4). In addition, a visual inspection of the retirement trends for control and treatment coal generators does not indicate a violation of the parallel trends assumption, which is pertinent for the difference-in-differences identification used in Equation 3.1.

3.6 Conclusion

Air emissions have been regulated for nearly 50 years in the U.S., but air pollution from the electric power sector has only been regulated for not quite 20 years. Hazardous air pollutants, like mercury, are emitted through fossil fuel-fired electricity generation. As early as 2003, states began regulating the amount of mercury released from coal-fired electricity generators. The federal government and other states followed suit and passed mercury regulations throughout the 2000s and 2010s. An extensive quantity of research has investigated the impact of environmental regulation on firm productivity and location choices, but much less is known about the exit decisions resulting from such regulation. This study examines whether mercury regulation affects the coal generator retirement decision and contributes to the literature by analyzing direct retirement decisions of over 1,200 coal generators and exploring other potential drivers of such retirements. A 2001-2016 panel of coal generators reveals that owners of coal generators did not seem to be retiring these units due to various mercury regulations. Instead, cheap natural gas prices and the vintage of the generator significantly

influenced whether a coal generator is retired or not. Additionally, an extension is provided that utilizes an event study to check for model assumptions. However, two open questions still remain: Did state and federal regulation reduce the amount of mercury released from the electric power sector? If so, was this because electricity plant managers chose to invest in abatement technology or retire their coal generators? Exploring the effectiveness of potential pathways for reduced air pollution through environmental regulation is useful for guiding future public policy decisions.

Chapter 4

Conclusions

This dissertation is comprised of three studies in energy economics. Chapter 1 examines the potential pathways by which firms were able to enter the natural gas market during the hydraulic fracturing boom. The natural gas boom of the 2000s was characterized by the highest prices and production in history. Advances in horizontal drilling, 3-D seismic imaging, and hydraulic fracturing made it highly profitable for firms to produce large quantities of shale gas. An often-overlooked source of increased production is the increase in the number of active firms in the market. The U.S. natural gas market was historically defined by large firms, but the boom enabled a large number of small firms - those drilling between one and eight wells - to enter the market. To better grasp the effects of changing market structure, this study developed a real options model of market entry and used data on natural gas turnover to test three potential explanations for small firm entry during the boom: (a) technological advances, (b) land lease speculation, and (c) regime changes in natural gas demand. The analysis revealed mixed support for the first explanation but strong support for the last two.

Chapter 2 investigates drivers of coal generator retirement using data on coal generator turnover, delivered coal prices, and wholesale electricity prices by estimating the impact of sunk retirement costs on the probability of coal generator retirement. By pairing an optimal stopping model of firms' generator retirement decisions with the retirement timing of almost 200 coal-fired generators across the U.S., it is possible to estimate implied retirement costs that are not typically disclosed by firms and are not publicly available. Because the

real options model cannot impute the retirement costs for coal generators that have not retired, I utilized propensity score matching to assign retirement cost amounts to active coal generators. With this data, a parametric approach can estimate the impact of retirement costs on the probability of retirement and finds that a one standard deviation increase in retirement costs results in a 0.2% reduction in the probability of retirement. Findings are robust across several specifications. A comparison of the predicted probability of retirement for active coal generators against the U.S. Energy Information Administration's reported retirements for 2016 found that almost half of the top 20 active coal generators with the highest predicted probability of retirement did indeed retire in 2016.

The final chapter examines how mercury regulation has impacted coal-fired electricity generator retirements. Previous research has emphasized the effect of environmental regulation on firm productivity, location choices, and labor market outcomes. Yet much less is known about how federal and state environmental regulations have affected firm exit decisions or, as in this case, coal generator retirements. I used generator level data from the Energy Information Administration and the Federal Energy Regulatory Commission, including mercury control technology, fuel costs, age, and abatement costs. The results suggest that mercury regulation was not the sole driver of coal generator retirements; instead, competitive natural gas prices and generator age substantially explain coal generator retirement decisions.

Bibliography

[1] Abadie, A. (2005). Semiparametric difference-in-differences estimators. *Review of Economic Studies*, 72(1):1–19. 76

[2] Acs, Z. and Audretsch, D. (1990). *Innovation and Small Firms*, volume 1. The MIT Press, 1 edition. 2

[3] Ai, C. and Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, 80(1):123–29. 71

[4] Albrizio, S., Kozluk, T., and Zipperer, V. (2017). Environmental policies and productivity growth: Evidence across industries and firms. *Journal of Envrionmental Economics and Management*, 81(1):209–226. 59

[5] Anderson, S. T., Kellogg, R., and Salant, S. W. (2017). Hotelling under pressure. Journal of Political Economy. forthcoming. 16

[6] Asplund, M. and Nocke, V. (2006). Firm turnover in imperfectly competitive markets. *Review of Economic Studies*, 73(2):295–327. 1

[7] Bain, J. S. (1954). *Conditions of Entry and the Emergence of Monopoly*, pages 215–241. Palgrave Macmillan UK, London. 2

[8] Baldwin, R. E. (1989). *Sunk-Cost Hysteresis*. National Bureau of Economic Research, Cambridge, MA. Working Paper No. 2911. 2

[9] Baumol, W. J. (1982). Contestable markets: An uprising in the theory of industry structure. *American Economic Review*, 72(1):1–15. 2

[10] Baumol, W. J. and Willig, R. D. (1981). Fixed costs, sunk costs, entry barriers, and sustainability of monopoly. *Quarterly Journal of Economics*, 96(3):405–31. 5, 30

[11] Becker, R. and Henderson, V. (2000). Effects of air quality regulations on polluting industries. *Journal of Political Economy*, 108(2):379–421. 59

[12] Becker, R. A. (2011). Local environmental regulation and plant-level productivity. *Ecological Economics*, 70(12):2516–2522. 59

[13] Berman, E. and Bui, L. T. (2001). Environmental regulation and productivity: Evidence from oil refineries. *Review of Economics and Statistics*, 83(3):498–510. 59

[14] Berry, S. and Reiss, P. (2007). *Empirical Models of Entry and Market Structure*, volume 3, chapter 29, pages 1845– 86. Elsevier. 3

[15] Black, D., McKinnish, T., and Sanders, S. (2005). The economic impact of the coal boom and bust. *The Economic Journal*, 115(503):449–76. 26

[16] Brekke, K. A. and Øksendal, B. (1994). Optimal switching in an economic activity under uncertainty. *SIAM Journal on Control and Optimization*, 32(4):1021–1036. 4, 9

[17] Brennan, M. J. and Schwartz, E. S. (1985). Evaluating natural resource investments. *The Journal of Business*, 58(2):135–57. 4

[18] Caves, R. E. and Porter, M. E. (1976). Barriers to exit. In Bain, J. S., Masson, R. T., and Qualles, P. D., editors, *Essays on Industrial Organization in Honor of Joe S. Bain*, pages 36–69. Ballinger Publishing Company. 27

[19] Cesena, E. M., Mutale, J., and Rivas-Davalos, F. (2013). Real options theory applied to electricity generation projects: A review. *Renewable and Sustainable Energy Reviews*, 19:573–81. 28

[20] Cicala, S. (2015). When does regulation distort costs? lessons from fuel procurement in u.s. electricity generation? *American Economic Review*, 105(1):411–44. 28

[21] Clarke, H. R. and Reed, W. J. (1989). The tree-cutting problem in a stochastic environment. *Journal of Economic Dynamics and Control*, 13(4):569–95. 4

[22] Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica*, 81(3):1003–37. 2

[23] Conrad, J. M. (1997). Global warming: When to bite the bullet. *Land Economics*, 73(2):164–73. 4, 101, 105

[24] Conrad, J. M. and Kotani, K. (2005). When to drill? trigger prices for the arctic national wildlife refuge. *Resource and Energy Economics*, 27(4):273–86. 4

[25] Cortazar, G. and Casassus, J. (1998). Optimal timing of a mine expansion: Implementing a real options model. *The Quarterly Review of Economics and Finance*, 38(3):755–69. 4

[26] Cortazar, G. and Schwartz, E. S. (1997). Implementing a real option model for valuing an undeveloped oil field. *International Transactions in Operational Research*, 4(2):125–37. 4

[27] Cortazar, G., Schwartz, E. S., and Casassus, J. (2002). Optimal exploration investments under price and geological-technical uncertainty: A real options model. *R&D Management*, 31(2):181–89. 4

[28] Costello, C. and Kolstad, C. D. (2016). *Mining and Quasi-Option Value*. National Bureau of Economic Research, Cambridge, MA. Working Paper No. 21325. 4

[29] Cullen, J. (2013). Measuring the environmental benefits of wind-generated electricity. *American Economic Journal: Economic Policy*, 5(4):107–33. 26

[30] Cullen, J. A. and Mansur, E. T. (2017). Will carbon prices reduce emissions in the u.s. electricity industry? evidence from the shale gas experience. American Economic Journal: Economic Policy. forthcoming. 26

[31] Davis, G. A. and Owens, B. (2003). Optimizing the level of renewable electric r&d expenditures using real options analysis. *Energy Policy*, 31(15):1589–1608. 28

[32] Dechezlepretre, A. and Sato, M. (2017). The impacts of environmental regulations on competitiveness. *Review of Environmental Economics and Policy*, 11(2):183–206. 59

[33] Dixit, A. (1989). Entry and exit decisions under uncertainty. *The Journal of Political Economy*, 97(3):620–38. 3, 28

[34] Dixit, A. (1992). Investment and hysteresis. *Journal of Economic Perspectives*, 6(1):107–32. 4

[35] Dixit, A. K. and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press, Princeton, NJ. 3, 9, 28, 32, 33, 101

[36] Duke Energy (2018). How energy works. https://www.duke-energy.com/energy-education/how-energy-works/electricity-from-coal. [Online; accessed 19-February-2018]. 67

[37] Dunne, T. and Roberts, M. J. (1991). Variation in producer turnover across u.s. manufacturing industries. In Geroski, P. A. and Schwalbach, J., editors, *Entry and Market Contestability*, pages 187–203. Oxford: Basil Blackwell. 27

[38] Finkelstein, A. (2007). The aggregate effects of health insurance: Evidence from the introduction of medicare. *Quarterly Journal of Economics*, 122(3):1–37. 76, 78

[39] Fleischman, L., Cleetus, R., Deyette, J., Clemmer, S., and Frenkel, S. (2013). Ripe for retirement: An economic analysis of the u.s. coal fleet. *The Electricity Journal*, 26(10):51–63. 26

[40] Folta, T. B., Johnson, D. R., and O'Brien, J. P. (2006). Uncertainty, irreversibility, and the likelihood of entry: An empirical assessment of the option to defer. *Journal of Economic Behavior & Organization*, 61(3):432–52. 3

[41] Forsyth, M. (2000). On estimating the option value of preserving a wilderness area. *Canadian Journal of Economics/ Revue canadienne d'Economique*, 33(2):413–434. 101, 105

[42] Foster, L., Haltiwanger, J., and Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425. 1

[43] Gearino, D. (2013). Aep to close coal-fired power plant instead of converting. http://www.dispatch.com/content/stories/business/2013/07/11/aep-to-close-coal-fired-power-plant-near-beverly.html. [Online; accessed 27-February-2017]. 42

[44] Geroski, P. (1995). What do we know about entry? *International Journal of Industrial Organization*, 13:421–40. 2

[45] Geske, R. (1979). The valuation of compound options. *Journal of Financial Economics*, 7(1):63–81. 4

[46] Ghosal, V. (2010). Quantifying the role played by sunk capital costs in real-options models. *Scottish Journal of Political Economy*, 57(3):343–358. 15

[47] Gray, W. B. and Shadbegian, R. J. (2003). Plant vintage, technology, and environmental regulation. *Journal of Envrionmental Economics and Management*, 46(1):384–402. 59

[48] Greenstone, M. (1998). The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the census of manufactures. Princeton University Working Paper #408. 59, 60

[49] Greenstone, M. (2002). The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the census of manufactures. *Journal of Political Economy*, 110(6):1175–1219. 59

[50] Greenstone, M., List, J. A., and Syverson, C. (2012). The effects of environmental regulation on competitiveness of u.s. manufacturing. 59

[51] Grenadier, S. R. (2002). Option exercise games: An application to the equilibrium investment strategies of firms. *Review of Financial Studies*, 15(3):691–721. 3

[52] Hefley, W. E. et al. (2011). The economic impact of the value chain of a marcellus shale well. *University of Pittsburgh*. 5, 6, 13, 14, 15

[53] Henderson, J. V. (1996). Effects of air quality regulation. *American Economic Review*, 86(4):789–813. 59

[54] Henson, L. (2004). Decommissioning Handbook for Coal-Fired Power Plants. Technical report. 42

[55] Hepbasli, A. (2008). A key review on exergetic analysis and assessment of renewable energy resources for a sustainable future. *Renewable and Sustainable Energy Reviews*, 12(3):593–661. 12, 36

[56] Hoag, J. H. and Wheeler, M. (1996). Oil price shocks and employment: The case of ohio coal mining. *Energy Economics*, 18(3):211–220. 26

[57] Holladay, J. S. and Soloway, S. (2016). The environmental impacts of fuel switching power plants. *The Energy Journal*, 37(4):187–213. 26

[58] Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60(5):1127–50. 1

[59] Insley, M. (2002). A real options approach to the valuation of a forestry investment. *Journal of Environmental Economics and Management*, 44(3):471–92. 4, 101, 105

[60] Insley, M. (2015). *On the Timing of Non-Renewable Resource Extraction with Regime Switching Prices: An Optimal Stochastic Control Approach*. Working Paper. 4

[61] Joskow, P. L. (2007). Competitive electricity markets and investment in new generating capacity. In Helm, D., editor, *The New Energy Paradigm*. Oxford University Press. 31

[62] Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica*, 50(3):649–70. 1

[63] Kaffine, D. T., McBee, B. J., and Lieskovsky, J. (2013). Emissions savings from wind power generation in texas. *The Energy Journal*, 34(2):155–75. 26

[64] Karaca-Mandic, P., Norton, E. C., and Down, B. (2012). Interaction terms in nonlinear models. *Health Services Research*, 47(1):255–74. 71

[65] Karlin, S. and Taylor, H. M. (1981). *A Second Course in Stochastic Processes*. Academic Press, New York. 30

[66] Kellogg, R. (2014). The effect of uncertainty on investment: Evidence from texas oil drilling. *American Economic Review*, 104(6):1698–1734. 4

[67] Knittel, C. R., Metaxoglou, K., and Trindale, A. (2017). Are we fracked? the impact of falling gas prices and the implications for coal-to-gas switching and carbon emissions. Oxford Review of Economic Policy. forthcoming. 26

[68] Kumbaroglu, G., Madlener, R., and Demirel, M. (2008). A real options evaluation model for the diffusion prospects of new renewable power generation technologies. *Energy Economics*, 30(4):1882–1908. 28

[69] Lanoie, P., Patry, M., and Lajeunesse, R. (2008). Environmental regulation and productivity: Testing the porter hypothesis. *Journal of Productivity Analysis*, 30(2):121–28. 59

[70] Li, E. (1998). Option value of harvesting: Theory and evidence. *Marine Resource Economics*, 13(2):135–42. 4

[71] List, J. A., Millimet, D. L., Fredriksson, P. G., and McHone, W. W. (2003). Effects of environmental regulations on manufacturing plant births: Evidence from a propensity score matching estimator. *Review of Economics and Statistics*, 85(4):944–52. 59

[72] List, J. A., Millimet, D. L., and McHone, W. (2004). The unintended disincentive in the clean air act. *B.E. Journal of Economic Analysis & Policy*, 4(2):Article 2. 60

[73] Loder, A. (2013). Fracking wells lack staying power. https://www.bloomberg.com/news/articles/2013-10-10/u-dot-s-dot-shale-oil-boom-may-not-last-as-fracking-wells-lack-staying-power. [Online; accessed 29-March-2017]. 5

[74] Lu, X., Salovaara, J., and McElroy, M. B. (2012). Implications of the recent reductions in natural gas prices for emissions of co2 from the u.s. power sector. *Environmental Science and Technology*, 46(5):3014–21. 26

[75] Marten, A. L. and Moore, C. C. (2011). An options based bioeconomic model for biological and chemical control of invasive species. *Ecological Economics*, 70(11):2050–2061. 11, 34

[76] Mason, C. F. (2001). Nonrenewable resources with switching costs. *Journal of Environmental Economics and Management*, 42(1):65–81. 4

[77] Mason, C. F. and Wilmot, N. A. (2014). Jump processes in natural gas markets. *Energy Economics*, 46:S69–S79. 25

[78] Massetti, E., Brown, M. A., Lapsa, M., Sharma, I., Bradbury, J., Cunliff, C., and Li, Y. (2017). Environmental quality and the u.s. power sector: Air quality, water quality, land use and environmental justice. Report ORNL/SPR-2016/772, Oak Ridge National Laboratory. 58

[79] McCarthy, J. E. (2007). Mercury emissions from electric power plants: States are setting stricter limits (crs order code rl33535. https://web.archive.org/web/20090210070452/http://ncseonline.org/NLE/CRSreports/07March/RL33535.pdf. [Online; accessed 19-February-2018]. 59, 61, 63, 66

[80] Millimet, D. L., Roy, S., and Sengupta, A. (2009). Environmental regulations and economic activity: Influence on market structure. *Annual Review of Resource Economics*, 1(1):99–118. 60

[81] Miranda, M. J. and Fackler, P. L. (2002). *Applied Computational Economics and Finance*. MIT Press, Cambridge, MA. 11, 33

[82] Muehlenbachs, L. (2015). A dynamic model of cleanup: Estimating sunk costs in oil and gas production. *International Economic Review*, 56(1):155–85. 4, 7

[83] Nostbakken, L. (2006). Regime switching in a fishery with stochastic stock and price. *Journal of Environmental Economics and Management*, 51(2):231–41. 4

[84] O'Brien, J. and Folta, T. (2009). Sunk costs, uncertainty and market exit: A real options perspective. *Industrial and Corporate Change*, 18(5):807–33. 28

[85] Ohio Department of Natural Resources (2005). Frequently asked questions. http://oilandgas.ohiodnr.gov/citizens/oil-gas-faq. [Online; accessed 13-December-2016]. 13

[86] Pachamanova, D. A. and Fabozzi, F. J. (2011). Modeling asset price dynamics. In Fabozzi, F. J. and Markowitz, H. M., editors, *The Theory and Practice of Investment*

Management: Asset Allocation, Valuation, Portfolio Construction, and Strategies, chapter 6, pages 125–58. John Wiley & Sons, Inc., 2 edition. 12, 13, 38, 103, 104, 107

[87] Paddock, J. L., Siegel, D. R., and Smith, J. L. (1988). Option valuation of claims on real assets: The case of offshore petroleum leases. *Quarterly Journal of Economics*, 103(3):479–508. 4

[88] Pakes, A. S. (1986). Patents as options: Some estimates of the value of holding european patent stocks. *Econometrica*, 54(4):755–84. 2

[89] Pettersson, F., Soderholm, P., and Lundmark, R. (2012). Fuel switching and climate and energy policies in the european power generation sector: A generalized leontief model. *Energy Economics*, 34(4):1064–73. 3, 26

[90] Pindyck, R. S. (1980). Uncertainty and exhaustible resource markets. *The Journal of Political Economy*, 88(6):1203–25. 6, 7

[91] Pindyck, R. S. (1991). Irreversibility, uncertainty, and investment. *Journal of Economic Literature*, 29:1110–48. 3

[92] Pindyck, R. S. (1998). *Sunk Costs and Sunk Benefits in Environmental Policy*. MIT Center for Energy and Environmental Policy Research, Working Paper 95-003. 4

[93] Pindyck, R. S. (1999). The long-run evolution of energy prices. *The Energy Journal*, 20(2):1–27. 8, 30

[94] Pindyck, R. S. (2000). Irreversibilities and the timing of environmental policy. *Resource and Energy Economics*, 22(3):223–59. 4

[95] Pindyck, R. S. (2002). Optimal timing problems in environmental economics. *Journal of Economic Dynamics and Control*, 26(9-10):1677–97. 4

[96] Pindyck, R. S. (2007). Uncertainty in environmental economics. *Review of Environmental Economics and Policy*, 1(1):45–65. 4

[97] Pindyck, R. S. and Rubinfeld, D. L. (1998). *Econometric Models and Economic Forecasts*. Irwin/McGraw-Hill, New York, NY. 101, 105

[98] Puhani, P. A. (2012). The treatment effect, the cross difference, and the interaction term in nonlinear "difference-in-differences" models. *Economics Letters*, 115(1):85–87. 71

[99] Reed, W. J. and Clarke, H. R. (1990). Harvest decisions and asset valuation for biological resources exhibiting size-dependent stochastic growth. *International Economic Review*, 31(1):147–69. 4

[100] Roberts, M. J. and Tybout, J. R. (1997). The decision to export in columbia: An empirical model of entry with sunk costs. *American Economic Review*, 87(4):545–64. 2

[101] Rosenbaum, D. I. and Lamort, F. (1992). Entry, barriers, exit, and sunk costs: An analysis. *Applied Economics*, 24(1):297–304. 27

[102] Rothwell, G. (2006). A real options approach to evaluating new nuclear power plants. *The Energy Journal*, 27(1):37–53. 28

[103] Rubashkina, Y., Galeotti, M., and Verdolini, E. (2015). Environmental regulation and competitiveness: Empirical evidence on the porter hypothesis from european manufacturing sectors. *Energy Policy*, 83(1):288–300. 59

[104] Shadbegian, R. J. and Gray, W. B. (2005). Pollution abatement expenditures and plant-level productivity: A production function approach. *Ecological Economics*, 54(2-3):196–208. 59

[105] Siegfried, J. J. and Evans, L. B. (1992). Entry and exit in united states manufacturing industries from 1977 to 1982. In Audretsch, D. W. and Siegfried, J. J., editors, *Empirical Studies in Industrial Organization: Essays in Honor of Leonard W. Weiss*. Dordrecht: Kluwer Academic Publishers. 27

[106] Siegfried, J. J. and Evans, L. B. (1994). Empirical studies of entry and exit: A survey of the evidence. *Review of Industrial Organization*, 9(1):121–55. 27

[107] Sims, C. and Finnoff, D. (2016). Opposing irreversibilities and tipping point uncertainty. *Journal of the Association of Environmental and Resource Economists*, 3(4):985–1022. 4

[108] Slade, M. E. (2001). Valuing managerial flexibility: An application of real-option theory to mining investments. *Journal of Environmental Economics Management*, 41(2):193–233. 4

[109] Snyder, L. D., Miller, N. H., and Stavins, R. N. (2003). The effects of environmental regulation on technology diffusion: The case of chlorine manufacturing. *American Economic Review: Papers and Proceedings*, 93(2):431–35. 60

[110] Spence, A. M. (1977). Entry, capacity, investment and oligopolistic pricing. *The Bell Journal of Economics*, 8(2):534–44. 1

[111] Taulli, T. (2011). *All About Commodities*. McGraw-Hill Education, 1 edition. 5

[112] Tsekrekos, A. E. (2010). The effect of mean reversion on entry and exit decisions under uncertainty. *Journal of Economic Dynamics & Control*, 34(1):725–42. 36

[113] U.S. Court of Appeals for the District of Columbia Circuit (2008). State of new jersey et al. vs. epa. https://www.cadc.uscourts.gov/internet/opinions.nsf/68822E72677ACBCD8525744000470736/$file/05-1097a.pdf. [Online; accessed 19-February-2018]. 64

[114] U.S. Energy Information Administration (2011). Age of electric power generators varies widely. https://www.eia.gov/todayinenergy/detail.php?id=1830. [Online; accessed 19-February-2018]. 61

[115] U.S. Energy Information Administration (2013). Annual report of natural and supplemental gas supply and disposition. https://www.eia.gov/dnav/ng/ng_pri_sum_dcu_nus_m.htm. [Online; accessed 29-September-2013]. 5

[116] U.S. Energy Information Administration (2014). Levelized cost and levelized avoided cost of new generation resources in the annual energy outlook 2014. https://www.eia.gov/outlooks/archive/aeo14/pdf/electricity_generation_2014.pdf. [Online; accessed 07-September-2016]. 42

[117] U.S. Energy Information Administration (2016a). Coal made up more than 80% of retired electricity generating capacity in 2015. https://www.eia.gov/todayinenergy/detail.php?id=25272. [Online; accessed 14-March-2017]. 28

[118] U.S. Energy Information Administration (2016b). Shale in the united states. https://www.eia.gov/energy_in_brief/article/shale_in_the_united_states.cfm. [Online; accessed 13-December-2016]. 13

[119] U.S. Energy Information Administration (2016c). U.s. natural gas gross withdrawals. http://www.eia.gov/dnav/ng/hist/n9010us2m.htm. [Online; accessed 15-December-2016]. 1

[120] U.S. Energy Information Administration (2017a). Form eia-860 detailed data. https://www.eia.gov/electricity/data/eia860/. [Online; accessed 19-February-2018]. 67

[121] U.S. Energy Information Administration (2017b). Net generation by energy source: Total - all sectors. https://www.eia.gov/electricity/data/browser/. [Online; accessed 30-April-2017]. 26

[122] U.S. Environmental Protection Agency (2011). Fact sheet mercury and air toxics standards for power plants. https://www.epa.gov/sites/production/files/2015-11/documents/20111221matssummaryfs.pdf. [Online; accessed 19-February-2018]. 59, 61, 63

[123] U.S. Environmental Protection Agency (2016). Coal plant decommissioning financing cleanup and redevelopment. https://nepis.epa.gov/Exe. [Online; accessed 26-April-2016]. 42

[124] U.S. Environmental Protection Agency (2017a). Basic information about mercury. https://www.epa.gov/mercury/basic-information-about-mercury. [Online; accessed 19-February-2018]. 64

[125] U.S. Environmental Protection Agency (2017b). Basic information about mercury and air toxics standards. https://www.epa.gov/mats/

basic-information-about-mercury-and-air-toxics-standards. [Online; accessed 19-February-2018]. 64

[126] U.S. Environmental Protection Agency (2017c). History of the mats regulation. https://www.epa.gov/mats/history-mats-regulation. [Online; accessed 19-February-2018]. 63

[127] U.S. Environmental Protection Agency (2017d). Summary of the clean air act. https://www.epa.gov/laws-regulations/summary-clean-air-act. [Online; accessed 19-February-2018]. 63

[128] U.S. Federal Energy Regulatory Commission (2017). Electric power markets: National overview. https://www.ferc.gov/market-oversight/mkt-electric/overview.asp. [Online; accessed 19-February-2018]. 67

[129] Vissing, A. (2017). *One-to-Many Matching with Complementary Preferences: An Empirical Study of Market Power in Natural Gas Leasing*. Under Review. 20

[130] Walker, W. R. (2011). Environmental regulation and labor reallocation: Evidence from the clean air act. *American Economic Review: Papers and Proceedings*, 101(3):442–47. 59

[131] Wang, Z. and Krupnick, A. (2013). A retrospective review of shale gas development in the united states. *Resources for the Future*. 1, 2, 12

[132] Wang, Z. and Xue, Q. (2014). The market structure of shale gas drilling in the united states. *Resources for the Future*. 1, 2, 12, 13, 15, 18, 19, 24

[133] Yang, M., Blyth, W., Bradley, R., Bunn, D., Clarke, C., and Wilson, T. (2008). Evaluating the power investment options with uncertainty in climate policy. *Energy Economics*, 30(4):1933–50. 28

[134] Yin, H., Kunreuther, H., and White, M. (2007). Do environmental regulations cause firms to exit the market? evidence from underground storage tank (ust) regulations. Risk Management and Decision Process Center, The Wharton School of the University of Pennsylvania. 60

[135] Zivot, E. and Andrews, D. W. K. (1992). Further evidence on the great crash, the oil price shock, and the unit-root hypothesis. *Journal of Business and Economic Statistics*, 10(3):251–270. 12

Appendices

A Appendix 1

A.1 Unit Root Tests

Given that real options results critically depend on choosing the correct stochastic process, we tested whether data were consistent with Brownian motion instead of our assumption of mean reversion. This is typically completed using an augmented Dickey Fuller test [23, 41, 59, 97]. Geometric Brownian motion (GBM) assumes P is log-normally distributed. The logged price level $p = \ln(P)$ is normally distributed and follows an arithmetic Brownian motion (ABM) $\mathrm{d}p = \mu \mathrm{d}t + s\mathrm{d}z$. If p is consistent with ABM, Ito's Lemma ensures P must be consistent with GBM. To test that p are consistent with ABM, we ran a restricted regression

$$(p_t - p_{t-1}) = \beta_0 + \beta_1(p_{t-1} - p_{t-2}) + \epsilon_t \tag{1}$$

and unrestricted regression

$$(p_t - p_{t-1}) = \beta_0 + \beta_1(p_{t-1} - p_{t-2}) + \beta_2 t + \beta_3 p_{t-1} + \epsilon_t \tag{2}$$

The null hypothesis that corresponds with p being ABM is $\mathrm{H}_0 : \beta_2 = \beta_3 = 0$.[1] This null hypothesis is rejected at the 1% significance level for natural gas wellhead prices (Table A.1). This is consistent with previous literature that has indicated that natural resource prices exhibited mean reversion tendencies [35]. It is not possible to perform a unit root test for natural gas proved reserves per well, since this study uses only data from 1989 to 1999. However, Figure A.1 appears to reject the GBM specification.

A.2 Geometric Mean Reversion Parameter Estimation

After determining that Brownian motion is inappropriate for modeling both natural gas prices and reserves, we turned to geometric mean reversion (GMR) $\mathrm{d}P = r_P(\bar{P} - P)P\mathrm{d}t + \sigma_P P \mathrm{d}z_P$. The geometric mean reversion model can be written as the following:

$$P_{t+1} = P_t + r_P(\bar{P} - P_t)P_t + \sigma_P P_t \epsilon_t \tag{3}$$

[1] This set-up is also true for testing if reserves R follow GBM.

Table A.1 : Unit Root Test for U.S. Natural Gas Wellhead Price Data, Monthly January 1976-December 1999

Unrestricted regression Coefficient	Estimate	Std. Error	t-statistic
β_0	0.0186	0.0086	2.17
β_1	0.276	0.0572	4.82
β_2	0.00004	0.00006	0.68
β_3	-0.0392	0.0125	-3.14
Restricted regression			
β_0	0.0035	0.0041	0.84
β_1	0.268	0.0579	4.63
$N = 286$			$F = 5.48$
Prob > F = 0.0046			

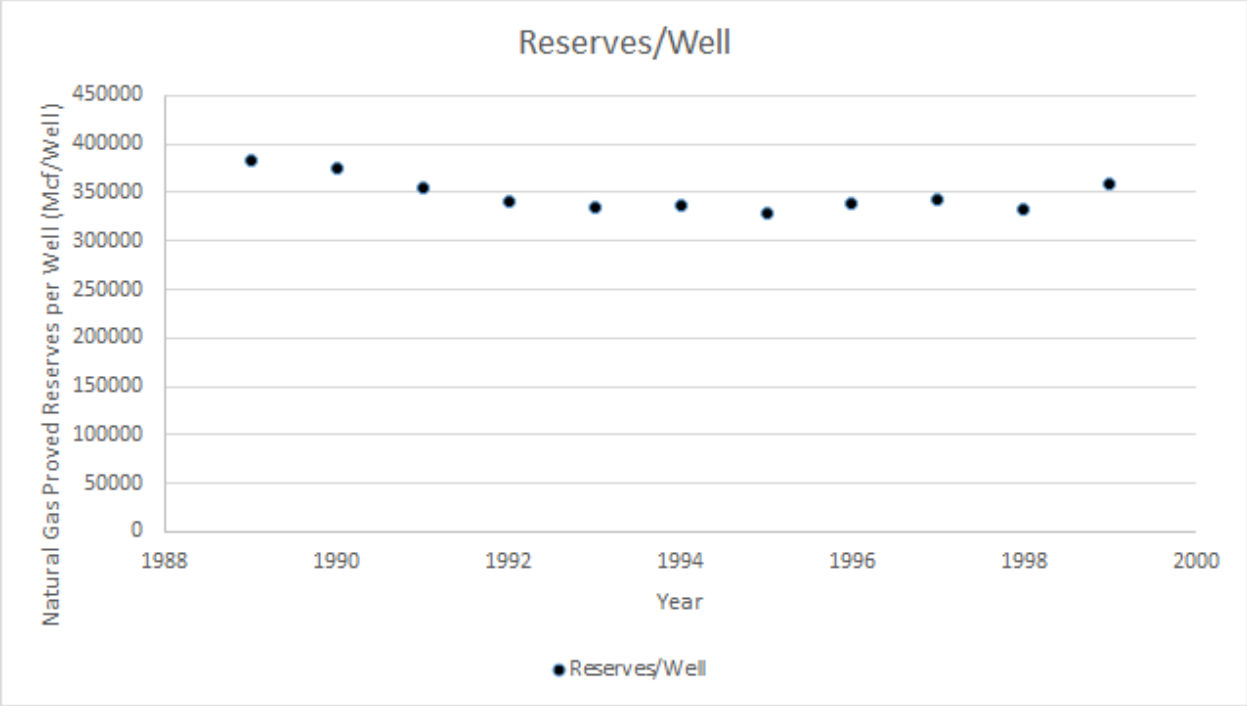

Figure A.1 : U.S. Natural Gas Proved Reserves per Well, 1989-1999

Here, ϵ_t is a standard normal random variable. To estimate the parameters r_P, $\bar{P}$, and σ_P, we used a series of 286 historical observations. Completing a Zivot-Andrews unit root test that allows for a single break in the intercept of the time series on data starting in January 1976, we determined that there was a structural break in January 2000 at the 1% significance level. Therefore, the 286 observations began in January 1976 and continued monthly to December 1999. This structural break is consistent with the natural gas (hydraulic fracturing) boom during which an influx of firms entered the market. It is with this pre-break data that we completed all estimations, including the previous augmented Dickey Fuller.

We assume that the parameters of the geometric mean reversion remained constant during the time period of estimation. Rewrite the equation for GMR as

$$\frac{P_{t+1} - P_t}{P_t} = r_P \bar{P} - r_P P_t + \sigma_P \epsilon_t \tag{4}$$

This equation bears characteristics of a linear regression model, with the percentage price change $\frac{P_{t+1}-P_t}{P_t}$ as the dependent variable and P_t as the explanatory variable.

According to [86], the estimate of r_P is obtained as the negative of the coefficient in front of P_t. One way to ascertain whether GMR is a consistent assumption for natural gas prices is to determine whether the coefficient in front of P_t is positive, since the rate of reversion r_P cannot be a negative number. The estimate for the long-run mean of prices $\bar{P}$ is obtained as the ratio of the intercept term estimated from the regression and the negative of the slope coefficient in front of P_t. The last estimate for volatility σ_P is obtained as the standard error of the regression.

Two other methods for determining whether GMR is suitable are the following: (a) the p-value for the coefficient in front of P_t should be small, preferably less than 0.05, and (b) the points in a scatter plot of P_t versus $\frac{P_{t+1}-P_t}{P_t}$ should vary around a straight line with no visible cyclical or other patterns. Table A.2 displays the results of the regressions that determine the parameter estimates for prices. Natural gas wellhead prices satisfy all three checks described by [86]. Our results were simulated under the assumption that natural gas wellhead prices follow GMR as previous studies have determined to be more fitting.

Table A.2 : Geometric Mean Reversion Estimates for U.S. Natural Gas Wellhead Price Data

Coefficient	Estimate	Std. Error	P-Value
$r_P\bar{P}$	0.0499	0.0136	0.000
r_P	-0.0235	0.0072	0.001
$N = 287$			$r_P = 0.0235$
$\bar{P} = \$2.12$			$\sigma_P = 0.0692$

Table A.3 : Geometric Mean Reversion Estimates for U.S. Natural Gas Recoverable Reserves per Well Data

Coefficient	Estimate	Std. Error	P-Value
$r_R\bar{R}$	0.443	0.210	0.068
r_R	-0.00000129	0.000000606	0.065
$N = 10$			$r_R = 0.00000129$
$\bar{R}$ =343,235			$\sigma_R = 0.0316$

It is not possible to perform a Zivot-Andrews unit root test allowing for a structural break in the intercept of the natural gas proved reserves per well time series due to limited observations. Simulations were completed with the time series from 1989 to 1999 following that of natural gas wellhead prices. Using [86], natural gas proved reserves per well satisfy two of the three checks for GMR (Table A.3). The p-value associated with the coefficient on R_t is slightly larger than the prescribed 0.05, but that coefficient is negative as it should be before applying -1 to find r_R. We checked the sensitivity of the results by modeling natural gas proved reserves per well as GBM and found no change.

B Appendix 2

B.1 Unit Root Tests

Given that real options results critically depend on choosing the correct stochastic process, we tested data for consistency with Brownian motion instead of the assumption of mean reversion. This is typically completed using an augmented Dickey Fuller test [23, 41, 59, 97]. Geometric Brownian motion (GBM) assumes P is log-normally distributed.[2] The logged price level $p = \ln(P)$ is normally distributed and follows an arithmetic Brownian motion (ABM) $\mathrm{d}p = \mu \mathrm{d}t + s\mathrm{d}z$. If p is consistent with ABM, Ito's Lemma ensures P must be consistent with GBM. To test that p are consistent with ABM, we ran a restricted regression

$$(p_t - p_{t-1}) = \beta_0 + \beta_1(p_{t-1} - p_{t-2}) + \epsilon_t \tag{5}$$

and unrestricted regression

$$(p_t - p_{t-1}) = \beta_0 + \beta_1(p_{t-1} - p_{t-2}) + \beta_2 t + \beta_3 p_{t-1} + \epsilon_t \tag{6}$$

The null hypothesis that corresponds with p being ABM is $\mathrm{H}_0 : \beta_2 = \beta_3 = 0$. This null hypothesis is rejected at the 1% or 5% level for all coal generators in the analysis. This is true for coal prices and electricity prices. Tables B.1 and B.2 provide an example of the augmented Dickey Fuller tests for a random retired coal generator.

B.2 Geometric Mean Reversion Parameter Estimation

After determining that Brownian motion is inappropriate for modeling both coal prices and electricity prices, we turned to geometric mean reversion (GMR) $\mathrm{d}P = r_P(\bar{P} - P)P\mathrm{d}t + \sigma_P P \mathrm{d}z_P$. The geometric mean reversion model can be written as the following:

$$P_{t+1} = P_t + r_P(\bar{P} - P_t)P_t + \sigma_P P_t \epsilon_t \tag{7}$$

[2] P is for either P_E or P_C

Table B.1 : Unit Root Test for Delivered Coal Prices at a Random Plant

Unrestricted regression		
Coefficient	Estimate	Std. Error
β_0	0.374	0.144
β_1	-0.363	0.0995
β_2	-0.000365	0.000222
β_3	-0.111	0.0558
Restricted regression		
β_0	0.00136	0.00607
β_1	-0.393	0.0985
$N = 90$		$F = 4.093$
Prob > F = 0.0200		

Table B.2 : Unit Root Test for Wholesale Electricity Prices at a Random Generator

Unrestricted regression		
Coefficient	Estimate	Std. Error
β_0	0.0314	0.0836
β_1	2.717	0.566
β_2	-0.401	0.0732
β_3	-0.00275	0.000679
Restricted regression		
β_0	-0.168	0.0823
β_1	-0.00417	0.0221
$N = 145$		$F = 15.06$
Prob > F = 0.00000019		

Here, ϵ_t is a standard normal random variable. To estimate the parameters r_P, $\bar{P}$, and σ_P, we used either FERC 714 data or PJM Zonal Prices. We assumed that the parameters of the geometric mean reversion remain constant during the time period of estimation to rewrite the equation for GMR as

$$\frac{P_{t+1} - P_t}{P_t} = r_P \bar{P} - r_P P_t + \sigma_P \epsilon_t \tag{8}$$

This equation bears characteristics of a linear regression model, with the percentage price change $\frac{P_{t+1}-P_t}{P_t}$ as the dependent variable and P_t as the explanatory variable.

According to [86], the estimate of r_P is obtained as the negative of the coefficient in front of P_t. One way to check whether GMR is a consistent assumption for prices is to determine whether the coefficient in front of P_t is positive since the rate of reversion r_P cannot be a negative number. The estimate for the long-run mean of prices $\bar{P}$ is obtained as the ratio of the intercept term estimated from the regression and the negative of the slope coefficient in front of P_t. The last estimate for volatility σ_P is obtained as the standard error of the regression.

Two other methods for determining if GMR is suitable are the following: (a) the p-value for the coefficient in front of P_t should be small, preferably less than 0.05, and (b) the points in a scatter plot of P_t versus $\frac{P_{t+1}-P_t}{P_t}$ should vary around a straight line with no visible cyclical or other patterns. Tables B.3 and B.4 display the results of the regressions that determine the parameter estimates for coal and electricity prices for a random coal generator in our analysis. Coal and electricity prices satisfied all three checks described by [86] for all generators. Our results were simulated under the assumption that coal and electricity prices follow GMR.

Table B.3 : Geometric Mean Reversion Estimates for Delivered Coal Prices at a Random Plant

Coefficient	Estimate	Std. Error
$r_P\bar{P}$	0.173	0.0598
r_P	-0.0460	0.0160
$N = 91$		$r_P = 0.0460$
$\bar{P} = \$3.76$		$\sigma_P = 0.0622$

Table B.4 : Geometric Mean Reversion Estimates for Wholesale Electricity Prices at a Random Generator

Coefficient	Estimate	Std. Error
$r_P\bar{P}$	0.226	0.0687
r_P	-0.0133	0.00436
$N = 146$		$r_P = 0.0133$
$\bar{P} = \$16.99$		$\sigma_P = 0.343$

Vita

Rebecca "Becky" Davis was born to Frank and Brenda in Richmond, Virginia on March 18, 1990. She was raised in Elkins, West Virginia and graduated from Elkins High School in 2008 as the valedictorian of her class. She attended West Virginia Wesleyan College and graduated summa cum laude in December 2011 with a Bachelor of Science in Economics. She moved to Knoxville, Tennessee in July 2012 to begin her graduate studies at the University of Tennessee. She received her Master of Arts degree in Economics in December 2013 and her Doctor of Philosophy degree in Economics in May 2018. Her work in energy and environmental economics has been supported by external funding from the Alfred P. Sloan Foundation through a Pre-Doctoral Fellowship Program on Energy Economics awarded by the National Bureau of Economic Research. She has presented her research at conferences across the United States, even as an invited presenter for the Southern Economic Association. She worked as a graduate research assistant for the Howard H. Baker Jr. Center for Public Policy for five years and contributed to white papers, reports, and trade publications for the Appalachian Regional Commission, the Tennessee Advanced Energy Business Council, and the Tennessee Department of Environment and Conservation. She was awarded the *Charles B. Garrison Award for Best Graduate Assistant* for her dedication to the center. She has taught undergraduate and graduate students at the University of Tennessee, and received the *Charles B. Garrison Award for Excellence in Teaching* from the Department of Economics. This fall, she will move to Nacogdoches, Texas, where she has accepted a position as Assistant Professor in Economics & Finance at Stephen F. Austin State University.