CAUSES AND CONSEQUENCES OF THE COAL MARKET DECLINE

by

CURTIS ALAN DLOUHY

A DISSERTATION

Presented to the Department of Economics
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

June 2020

DISSERTATION APPROVAL PAGE

Student: Curtis Alan Dlouhy

Title: The Causes and Consequences of the Coal Market Decline

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Economics by:

Wesley W. Wilson	Chairperson
Van Kolpin	Core Member
Keaton Miller	Core Member
Robert Ready	Institutional Representative

and

Kate Mondloch	Interim Vice Provost and Dean of the Graduate School

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded June 2020

DISSERTATION ABSTRACT

Curtis Alan Dlouhy

Doctor of Philosophy

Department of Economics

June 2020

Title: The Causes and Consequences of the Coal Market Decline

Fossil fuel markets have formed the backbone of commerce in the United States for the better part of the last century. Whether it be through the extraction of raw materials, the refinement for future use, their use in the transportation industry, or burning them for heat or electricity, fossil fuels have become a necessary resource in the post-industrial economy. While fossil fuels are indeed essential in many sectors, their roles have shifted due to changes in technology, public opinion and public policy.

An unfortunate byproduct of using fossil fuel use is a host of harmful pollutants in the form of sulfur, nitrogen oxides, and carbon. Due to their dirty nature, policy makers have tried to disincentivize fossil fuels or reduce their emissions. Starting in the 1970s, the United States began to reward reductions in fossil fuel use and the use of emissions-reduction technology through the Clean Air Act, its many amendments, and many other regional environmental policies. While the US is still very dependent on fossil fuels nearly 50 years after the institution of these original policies, the composition of fossil fuels used, and industries servicing users of fossil fuels have changed dramatically.

In this dissertation, I discuss my research investigating how the changing roles of fossil fuel have affected coal mining markets, electricity generation, and rail transportation. In my first chapter, I develop a model of sunk cost hysteresis in the coal

mining industry to discuss how the rise of natural gas and environmental regulations have affected coal mining operations. In chapter two, I discuss how a carbon tax shifts the dispatch order of fossil fuel electricity generators and the effect that this redispatching has on seasonal fossil fuel use. In chapter three, I discuss how the decline in coal mining and the diminished preference for coal-powered electricity affects rail rates and rail revenues from transporting coal.

CURRICULUM VITAE

NAME OF AUTHOR: Curtis Alan Dlouhy

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene
Nebraska Wesleyan University, Lincoln, NE

DEGREES AWARDED:

Doctor of Philosophy, Economics, 2020, University of Oregon
Master of Science, Economics, 2016, University of Oregon
Bachelor of Science, Mathematics & Economics, 2015, Nebraska Wesleyan University

AREAS OF SPECIAL INTEREST:

Industrial Organization
Energy Economics

GRANTS, AWARDS, AND HONORS:

Teaching Fellowship, University of Oregon, 2015-2020

Outstanding Undergraduate Teaching Award, University of Oregon, 2019

Kleinsorge Summer Research Fellowship, University of Oregon, 2019

Summer Teaching Fellowship, University of Oregon, 2017, 2018

Dale Underwood Outstanding 1st Year PhD Student, University of Oregon, 2016

TABLE OF CONTENTS

Chapter Page

Chapter Page

LIST OF FIGURES

LIST OF TABLES

CHAPTER I

INTRODUCTION

Fossil fuel markets have formed the backbone of commerce in the United States for the better part of the last century. Whether it be through the extraction of raw materials, the refinement for future use, their use in the transportation industry, or burning them for heat or electricity, fossil fuels have become a necessary resource in the post-industrial economy. While fossil fuels are indeed essential in many sectors, their roles have shifted due to changes in technology, public opinion and public policy.

An unfortunate byproduct of using fossil fuel us is a host of harmful pollutants in the form of sulfur, nitrogen oxides, and carbon. Due to their dirty nature, policy makers have tried to disincentivize fossil fuels or reduce their emissions. Starting in the 1970s, the United States began to reward reductions in fossil fuel use and the use of emissions-reduction technology through the Clean Air Act, its many amendments, and many other regional environmental policies. While the US is still very dependent on fossil fuels nearly 50 years after the institution of these original policies, the composition of fossil fuels used, and industries servicing users of fossil fuels have changed dramatically.

In this dissertation, I discuss my research investigating how the changing roles of fossil fuel have affected coal mining markets, electricity generation, and rail transportation. In my first chapter, I develop a model of sunk cost hysteresis in the coal mining industry to discuss how the rise of natural gas and environmental regulations have affected coal mining operations. In chapter two, I discuss how a carbon tax shifts the dispatch order of fossil fuel electricity generators and the effect that this redispatching has on seasonal fossil fuel use. In chapter three, I discuss how the decline in coal mining and the diminished preference for coal-powered electricity affects rail rates and rail revenues from transporting coal.

The first chapter of my prospectus focuses on the United States coal mining industry. The coal-mining industry in the United States has been in a steady decline since 2008, and in 2016, it posted its lowest aggregate production in the last 25 years. Because of the persistence of entry and re-entry despite the decline, a major platform in the 2016 election cycle was to return coal mining to its former prominence by adjusting environmental policy. In this chapter, I quantify the determinants of entry and exit into the coal-mining industry using a model of sunk-cost hysteresis. I find that the shale gas boom had a significant negative effect on a mine's propensity to participate in the coal mining industry, and that this effect was larger than any national environmental policy since 2000. Combined with the large fixed costs of entry, I find that returning coal mining to its former prominence is unlikely.

While I find that policy surrounding coal mining has had a relatively small effect on coal mine participation in my first chapter, I find that policy has had a much larger effect on power plants' behavior. The second chapter of my prospectus analyzes the carbon cap-and-trade policies. Cap-and-trade programs are a novel way to disincentivize fossil fuel use and decrease carbon emissions. I develop a theoretical model to demonstrate that carbon taxes encourage the use of natural gas plants to generate baseload electricity and coal to supplement seasonal changes in electricity demand, essentially switching their current roles. Using monthly US power plant data from the EIA and exploiting the seasonal nature of electricity demand, I test these theoretical conclusions by using the recently-created Regional Greenhouse Gas Initiative in the northeastern United States. I find that the reduction in coal use is largest in low-demand fall and spring months while the increase in natural gas use is uniform across months. Additionally, I find no evidence that the pollution haven hypothesis is relevant in coal electricity generation.

In the third chapter of this dissertation prospectus, I analyze how the sudden demise of coal has affected its closest complementary industry- railroads. As was discussed in the previous two chapters, natural gas, renewable energy, and environmental policies have rapidly displaced coal in electricity markets. The impact on the rail industry- which transports almost all coal in the US and derives around 40% of its traffice from transporting coal- has been profound, but little research has been done to quantify it. I study the effects of coal's demise on rail rates and coal quantities shipped and find that the as plants close down, rail firms raise their rates for plants that remain in the market. I find that the opposite result holds for coal mines, as rail firms appear to lower their rates on shipments when a mine in the area closes down. This indicates that rail firms aim to accommodate coal mines as the industry declines, but try to "shake out" any remaining profit from power plants as they switch away from coal.

CHAPTER II

RE-IGNITING THE COAL MINING INDUSTRY: IS IT POSSIBLE OR AR WE JUST BLOWING SMOKE?

II.1 Introduction

In May 2016, Donald Trump proudly claimed, "If I win, we're going to bring those miners back" at a West Virginia rally consisting largely of displaced coal miners.[1] Up until the turn of the millennium, coal was the undisputed backbone of the United States' energy grid. Due to a combination of environmental policies, technological advances and the rise of natural gas, coal's place as the king of electricity generation has slipped in the last decade, causing many mines to shut down and forcing many miners out of their jobs. This can be seen in Figure II.1. There is a sharp drop off in both the number of operating mines and employed miners around 2009 at the same time that the real natural gas price drops dramatically.

Throughout his campaign and the first year of his term, one of President Trump's policy platforms was to end the supposed war on coal that started under the Obama presidency, open new mines and deliver on his campaign promise to "bring those miners back." At best, many energy experts, journalists and political scientists are skeptical of his promises. Despite the skepticism, the Trump administration advocates for harmful changes to environmental policies to prop up the coal industry.[2]

[1] http://www.mcclatchydc.com/news/politics-government/article114378343.html

[2] https://www.npr.org/2018/02/23/586236738/coal-jobs-have-gone-up-under-trump-but-not-because-of-his-policies

Although many journalists have asserted that Donald Trump's goal to breathe life into the dying coal industry and open new mines is impossible due to the decline,[3] there has still been some recent entry and reentry into the industry. To my knowledge, no economic papers have empirically examined the claim that the coal mining industry can be restored through policy. In this chapter, I analyze the entry and exit decisions of coal mines to determine how much of the coal industry's decline is due to mine-level effects, how much is due to environmental policies, how much is due to the large drop in natural gas price using an empirical model of sunk-cost hysteresis.

Models of sunk-cost hysteresis are commonly used to analyze entry and exit patterns into industries with high fixed costs. Often they are used to analyze exporting decisions, sunk costs, and entry and exit levels (Roberts and Tybout (1997), Bernard and Jensen (2004) and Máñez et al. (2008)). I develop a model in the style of Bernard and Jensen (2004) to quantify the determinants of entry and exit into the coal-mining industry.

Figure II.1
Quarterly Operating Mines and Employment

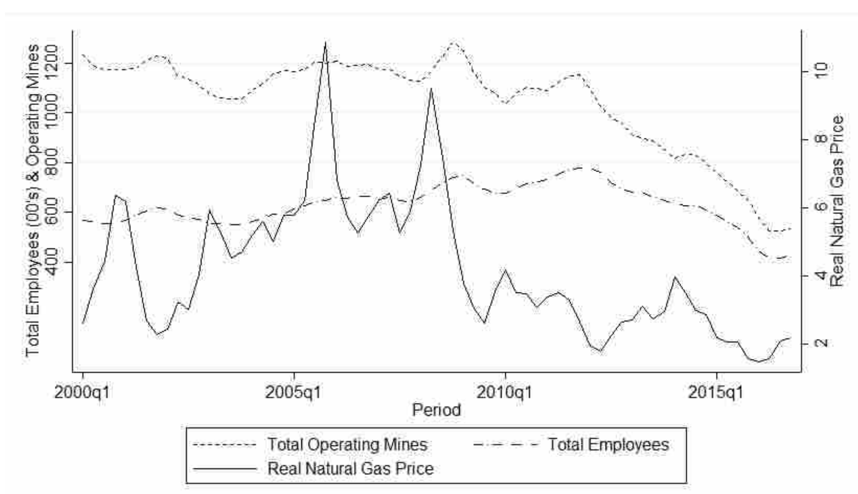

Note: Mine data are taken from the Mine Safety and Health Administration's quarterly report. Real Natural Gas Price is expressed in Jan 2000 $ per mmBTU

[3] http://www.independent.co.uk/news/world/americas/donald-trump-us-coal-industry-bleak-future-hiring-growth-prospects-energy-environment-rick-perry-epa-a8051886.html

My model allows for profitability to vary by mine, firm and macroeconomic events. I use it to discuss firm- and mine-level characteristics of mines that choose to exit or participate in the industry as well as the effects of two major policies and the drop in the natural gas price due to the shale boom. I find that the sudden drop in the natural gas price of around $8 per million BTU caused by the shale gas boom in 2008 led to approximately a 2.4 percentage point increase in the probability that a mine closes each quarter on average. This effect varies substantially across the five major coal basins in the United States. For example, the shale gas boom caused a 1.9 percentage point drop in the probability a mine from Northern Appalachia participates in the market each quarter and a 3.6 percentage point drop in the probability a mine from Central or Southern Appalachia participates in the market each quarter, but had little negative effect elsewhere. I also find that new environmental policy during the decline of coal has played an insignificant role in the decline of the coal industry, leading me to conclude that the shale gas boom rather than Obama-era environmental policy led to the decline of the United States' coal mining industry.

The rest of the chapter is structured as follows. Section 1.2 begins with a brief history of the coal industry and the environmental legislation relevant to it. I then discus the technological innovations in the energy sector and the rise in popularity of natural gas that occurred around the time of the decline of coal. Section 1.3 provides a review of the academic literature on coal mining, theoretical models of entry and exit and sunk-cost hysteresis, and empirical techniques used in the sunk-cost hysteresis literature. In Section 1.4, I develop a theoretical model of mine entry and exit based on the sunk-cost hysteresis literature, and in Section 1.5 I adapt the theoretical model into an empirical model of entry and exit based on the work of Bernard and Jensen (2004). I discuss the results of the empirical estimation in Section 1.6 and Section 1.7 provides some concluding remarks.

II.2 Background

Coal mining has traditionally been a critical industry throughout the history of the United States. In the 1800s coal was primarily used a heat source, but was eventually used for more productive purposes due to the Industrial Revolution. At the turn of the twentieth century, coal occupied three main functions in the US: electricity generation, steel manufacturing, and trade. As the century progressed, coal's main use shifted almost entirely to energy, with only a small percentage being devoted anything but domestic electricity production. However, the demand for coal as a power source dominated the demand for coal as a means to producer steel or a tradable commodity, and the coal industry continued to thrive through the 1900s.[4]

For the bulk of the twentieth century, fossil fuels formed the backbone of the United States electricity grid. Along with an incomplete understanding of the effects of fossil fuel emissions, fossil fuels' relative abundance and large power output made them an attractive power source for many local utilities. Being the most abundant, coal was the primary form of energy generation in spite of its reputation as a very dirty source of energy. Within the last ten years, natural gas prices have fallen dramatically due to the rise of hydraulic fracturing and horizontal drilling. This has substantially lowered the demand for coal.[5] Regulations on emissions have also made coal a much less attractive form of power than natural gas and renewable energy sources.

Figure II.2 shows the effects of this change since the year 2000, and demonstrates two things. First, the total net generation of electricity in the US from 2001 to 2016 has stayed relatively constant apart from a small recessionary dip in 2008, meaning that the rise in use of one fuel source will necessarily result in the substitution away from another fuel source. Second, while the total electricity consumption in the US

[4] This information comes from the American Coal Foundation. A brief timeline of US coal can be found at http://www.teachcoal.org/lessonplans/pdf/coal_timeline.pdf

[5] https://www.npr.org/2015/06/23/414926833/how-fracking-is-changing-the-nation-s-electrical-grid

Figure II.2
Net Generation for All Sectors, Yearly

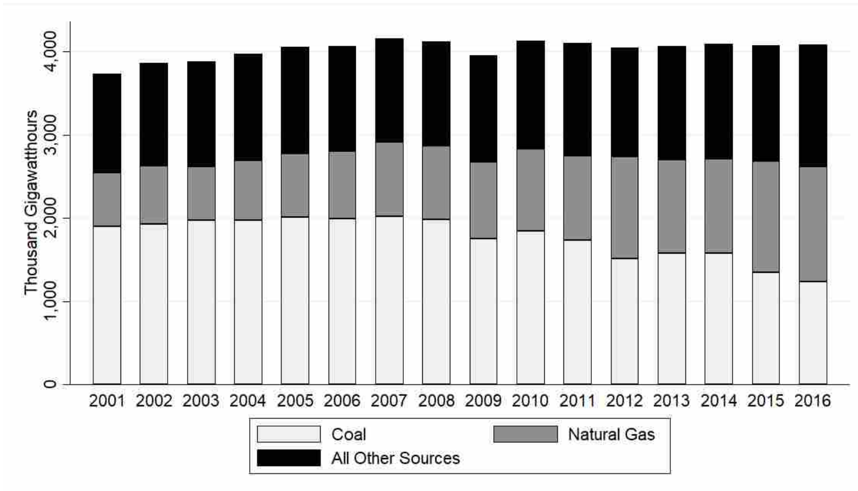

Note: These data are taken from EIA.gov's Total Energy Data Browser.

and the total generation by coal and gas combined have stayed relatively constant, there is a marked fall in the use coal and rise in the use of natural gas, indicating that there is some substitution between the two electricity sources. While the underlying determinants of this shift are not immediately apparent, this new preference of gas over coal has translated to a large decline in demand for US coal.

II.2.1 A Brief History of Environmental Legislation Relevant to the Coal Industry

In the 1960s, the United States became aware its substantial acid rain problem, which contributed to the deterioration of natural lake and forest habitats in addition to accelerating the erosion of buildings. As a response, Congress passed the Clean Air Act in 1963 to curb emissions of sulfur dioxide (SO_2), carbon, nitrogen dioxide (NOx) and other chemicals.[6] The original version of the Clean Air Act forced new power plants to adhere to low emissions standards but exempted older plants with

[6]For further detail, see the EPA's website for a brief history of the Clean Air Act https://www.epa.gov/clean-air-act-overview/evolution-clean-air-act

the expectation that these older plants would eventually be phased out and replaced with newer, lower-emissions plants. However, the Clean Air Act proved to be too lax, and the EPA found that many states had still not met their target levels of emissions after over a decade.[7]

As a response, many small amendments to the Clean Air Act and pieces of legislation were passed in the 1970s and 1980s that would affect the demand for coal to some degree, but perhaps the most important bit of legislation to the coal industry was the Clean Air Act of 1990. In this, Congress instituted a cap on SO_2 and NOx emissions by power plants, but cap was not initially binding. However, every year the allowable amount of pollution emissions was lowered, and eventually utilities had to respond to the now-binding cap in two ways: buying pollution credits or cutting back emissions.[8] The Clean Air Act of 1990 allowed pollution credits to be traded on the open market, allowing dirtier plants to buy pollution credits from cleaner plants. In order to avoid paying for costly pollution credits, many power plants voluntarily upgraded their facilities by installing flue-gas desulfurization (FGD) scrubbers to convert SO_2 emissions and other pollutants into a solid or liquid form that could be more easily disposed. Despite this setback, coal production still thrived and in fact led to a rapid development of the Powder River Basin due to its low-sulfur sub-bituminous coal.

Further limitations were put on the emissions over the next decade. The EPA issued the Clean Air Interstate Rule in 2005, which lowered the cap on SO_2 and NOx emissions in the eastern United States and limited the amount of mercury that a coal-fired power plant could emit. The Clean Air Interstate Rule further incentivized many power plants to install FGD scrubbers that eliminate SO_2 and catalyzers that

[7] https://www.sourcewatch.org/index.php/Sulfur_dioxide_and_coal

[8] There is an extensive body of literature in environmental economics detailing the effects of the Cap-and-Trade pollution credits program and its effect on both emissions and the portfolio of energy generation, starting with Schmalensee et al. (1998)

eliminate NOx from emissions, which is no small investment.[9]

Installing an FGD scrubber is no small investment. According to the EIA-860 Form on pollution abatement equipment, the mean nominal cost of installing an FGD scrubber since 1960 is approximately \$62 million, and has been over \$700 million.[10] This does not take into account any opportunity costs during installation or future maintenance costs, and many power plants estimate the present value of installing a scrubber to be well over a billion dollars. To give this value some context, the largest coal-powered plant in the country is the Scherer plant in Georgia, which generated around 16 million megawatt-hours of power in 2016 when the state's retail price for electricity was 9.62 cents per kilowatt-hour. This translates into over a billion dollars of yearly revenue created by the Scherer plant, meaning that a scrubber may cost as much as half of a plant's total yearly revenue.[11]

Congress passed the National Ambient Air Quality Standards (NAAQS) in 2007, which imposed a penalty on any county that was found to emit too much of various pollutants including SO_2 and NOx. The pollutants are monitored on three-hour, daily and yearly intervals and any state with a county that is out of attainment must submit a plan to limit their emissions or face harsh federal penalties. In 2010, the EPA revised the NAAQS standards on SO_2 to say that any county that was found to have over 75 parts per billion of each pollutant in each hour would be considered out of attainment.[12]

In 2009, ten states in the northeastern United States agreed to the Regional Greenhouse Gas Initiative (RGGI), the first program to impose a mandatory cap on carbon dioxide emissions from power plants. Much like the Clean Air Act of 1990 did

[9]https://www.eia.gov/todayinenergy/detail.php?id=10151

[10]The cost of installing an FGD scrubber varies greatly with plant scale. For a breakdown of projected costs, see http://www.powermag.com/whats-that-scrubber-going-to-cost/?pagenum=1

[11]These numbers come from the EIA-FERC 923 and state electricity retail rates found at https://www.eia.gov/electricity/state/

[12]https://www.epa.gov/so2-pollution/table-historical-sulfur-dioxide-national-ambient-air-quality-standards-naaqs

with SO_2 and NOx emissions, RGGI forced all power plants within the ten states to participate in a cap-and-trade program in the CO_2 market. It has been shown that this program was responsible for approximately half of the reduction of carbon emissions in the participating states Murray and Maniloff (2015). While all fossil fuels necessarily will emit CO_2 as a by-product, RGGI should have asymmetric effects on coal and natural gas because coal emits approximately twice as much CO_2 per unit of energy than natural gas.[13]

While it never went into full effect, the Clean Power Plan proposed during the Obama administration had many potential direct impacts on the coal industry. The Clean Power Plan was a comprehensive plan to upgrade the United States' energy grid and reduce airborne emissions by the year 2030, and it seemed to spell even more problems for the already declining coal industry. The EPA fact sheet on the Clean Power Plan stated that the first goal of the Clean Power Plan was to gradually shift all fossil fuel power generation to natural gas, and make the natural gas grid powerful enough to be able to generate 75% of the US's net summer capacity. Next, the Clean Power Plan outlines steps to shift electricity generation to entirely renewable resources with zero emissions in the long term. Because coal is a non-renewable resource that produces a lot of emissions, the Clean Power Plan essentially spelled out the eventual death of the United States coal industry. The Clean Power Plan never went into full effect due to many states' vehement opposition to it. 28 states filed lawsuits against the Clean Power Plan; all heavily relied on the fossil fuel industry in some form.[14]

While the Clean Power Plan may have sent a signal that the US was shifting away from coal, its reign was ultimately short lived. On March 28, 2017, President Trump signed an executive order to mandate that the head of the EPA review and potentially

[13]https://www.eia.gov/tools/faqs/faq.php?id=73&t=11

[14]At the beginning of his term, President Trump ordered the dismantling of many government programs and data sets on the internet, including documents relating to the Clean Power Plan. I include a link to an online screenshot of the EPA's fact sheet before it was officially removed from the EPA's website https://19january2017snapshot.epa.gov/cleanpowerplan/fact-sheet-overview-clean-power-plan_.html

suspend the Clean Power Plan as a part of his Energy Independence Executive Order, claiming that this would help return jobs to fossil fuel industries, especially the coal industry.[15] As has been the case since his campaign, many experts questioned the validity of Trumps promise to revitalize the coal industry, asserting instead that the decline of coal is due to market forces instead of any overly-stringent legislative action or environmental regulation.[16] On October 10, 2017, the EPA proposed a full repeal of the Clean Power Plan on the grounds that it incorrectly applied language from the Clean Air Act.[17]

II.2.2 The rise of natural gas; the decline of coal

A power plant that runs on coal can be converted to a plant that runs on another type of fossil fuel with relative ease but at a high cost. Most commonly, the boilers in coal plants are converted to use natural gas as its fuel due to the lower environmental impact of natural gas. In some cases, this conversion can cost upwards of $200 million.[18] Near the end of the 2000s, this conversion became more economically viable due to natural gas becoming very cheap. This can be seen by the Henry Hub Natural Gas Spot Price in Figure II.3. Apart from some occasional noise, the natural gas price seemed to be steadily trending upward until its peak in 2008. After 2008, the spot price of natural gas drastically fell, and has stayed consistently very low.

The sharp drop in price almost immediately translated in a substitution from coal to natural gas, which can be seen in yearly production of coal and natural gas. In Figure II.4, I show the yearly quantities of coal and natural gas extracted over

[15] https://www.whitehouse.gov/the-press-office/2017/03/28/presidential-executive-order-promoting-energy-independence-and-economi-1

[16] Many sources have written editorials on the president's claims. I include only one editorial that links to many others https://fivethirtyeight.com/features/trumps-plan-wont-reverse-coals-decline/

[17] https://www.epa.gov/sites/production/files/2017-10/documents/fs-proposed-repeal-cpp-final_oct10.pdf

[18] The cost of a conversion varies greatly based plant size, type of conversion (i.e. totally natural gas or natural gas-coal hybrid), and the current equipment the plant has.

time.[19] Both values stayed fairly steady until 2008, when the natural gas production increased dramatically and the coal production fell dramatically.

Apart from the large drop in the natural gas price, the shift away from coal in the electricity market can be partially attributed to its real negative externalities and those perceived by environmental advocacy groups. The Sierra Club is a national organization with over two million members that fights for environmental protection, and its Beyond Coal campaign advocates for the closure or conversion of coal-fired power plants. According to a 2015 news release by the Sierra Club, there were 523 power plants with coal-fired generators in 2010. Of these 523 coal-fired plants, over 250 were shut down or converted to natural gas plants by 2015.[20]

The transition away from coal has been especially hard on firms that have large

Figure II.3
Real Natural Gas Spot Price, Weekly

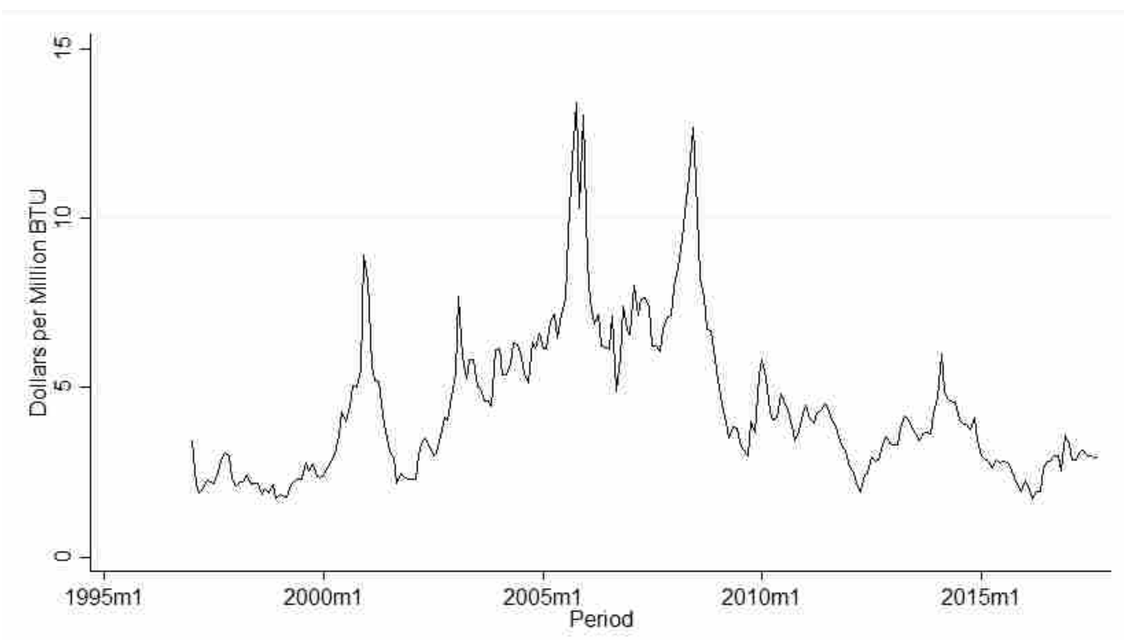

[19]It is worth noting that natural gas is used extensively in commercial and residential heating, and industrial production, and its consumption is highly cyclical.

[20]The URL links to the news releases by the Sierra Club are included here http://content.sierraclub.org/press-releases/2015/07/united-states-phases-out-200th-coal-plant-momentum-renewable-energy-grows and here http://content.sierraclub.org/press-releases/2017/03/milestone-250th-and-251st-american-coal-plants-announce-retirement.

Figure II.4
Coal v. Natual Gas Production, Yearly

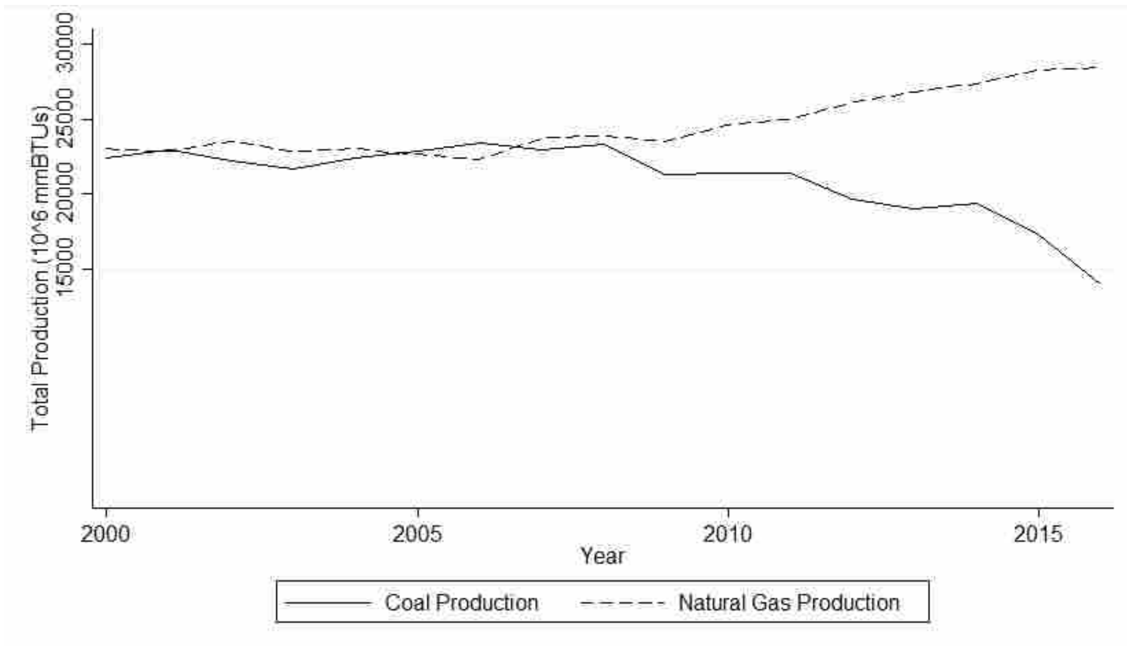

Coal production is calculated from MSHA Employment/Production Dataset. The natural gas production reflects all natural gas extracted reported by EIA Natural Gas Consumption End Use Data Series. Coal and gas production are converted to mmBTUs using the annual average heat content by fuel source from EIA's Electric Power Annual Reports.

stakes in the coal mining industry. On January 15, 2016, the former Secretary of Interior Sally Jewell announced a moratorium on permits to mine coal on federally controlled lands.[21] In 2015, the US coal industry also produced the lowest annual amount of coal in the last 30 years. The combination of the moratorium, the rise of natural gas, and the closure or conversion of so many coal-fired power plants led to Peabody Energy, Arch Coal, Walter Energy, and Alpha Natural Resources to file for Chapter 11 bankruptcy in 2016. While all four companies still participate in the coal-mining market, they all had to extensively restructure their companies to remain viable. Peabody Energy, Arch Coal and Alpha Natural Resources constituted 42% of the total quantity of coal mined in the United States in 2015.[22]

[21] This moratorium has since been lifted.
[22] http://fortune.com/2016/04/13/peabody-energy-chapter-11-protection/

II.3 Review of Literature

While most believe that coal mining is an industry in permanent decline, there is still entry. Given the presence of entry, President Trump believes that it can be saved. I present a review of the literature on declining industry models, entry-exit models, mining and exhaustible resources in general, and the coal industry in particular.

A very rich set of research to theoretically model the industry decline arose in the 1980s and 1990s. Ghemawat and Nalebuff (1985) set up a game-theoretic framework for a declining industry with exogenous plant capacities. They found that in an industry where each plant has a fixed capacity, as the industry declines, the larger firms within the industry will exit before the smaller firms in a subgame-perfect equilibrium due to an inability to recoup its higher fixed costs. Ghemawat and Nalebuff (1990) extend this research by allowing firms to adjust plant capacity every period, which factors in as a fixed cost paid during the operating period. By allowing plants to instead choose their capacity, they find that in equilibrium, as the industry shrinks, all large firms will lower their capacity until they are the same size of the smallest firm, and then all firms will lower their capacities at the same rate until the industry ceases to exist.

Reynolds (1988) and Whinston (1988) set up game-theoretic models in declining industries with multi-plant firms. Whinston extends Ghemawat and Nalebuff's 1985 model to allow for a firm to have multiple plants. Contrary to Ghemawat and Nalebuff, Whinston found no consistent rule that related firm size to exit order. Reynolds sets up a model that allows firms to own any number of different sized plants. He finds that firms will shut down high-cost plants first, and firms with many plants will begin shutting down plants before firms with few plants.

In addition to the theoretical literature on exit patterns in declining industries, there is a dense literature on entry and exit into an industry. In his seminal paper, Jovanovic (1982) sets up a model in which firms are unaware of their average production efficiency relative to the market's average production efficiency. This results in firms noisily entering and exiting, and a slow, steady concentration of the market. Ericson and Pakes (1995) create a similar, empirically estimable model in which firms make entry, exit and investment decisions in a game-theoretic environment. The model gives rise to a Markov-Perfect Nash Equilibrium.

Sunk-cost hysteresis is a broad class of entry-exit models where entrants are required to pay a sunk start-up cost to enter into an industry and a separate cost to exit. The sunk costs give rise to a market where entry does not occur until an industry becomes exceptionally profitable, and exit does not occur until an industry becomes exceptionally unprofitable. This leads to a range of profitability where no entry and no exit occurs called a "hysteresis band." Baldwin (1988) and Dixit (1989, 1992) provide theoretical models to motivate sunk-cost hysteresis and hysteresis bands. Baldwin and Krugman (1989) provide another model wherein large shocks in the exchange rate can cause firms to exit an export market due to low profitability and not re-enter due to never reaching a profit level above the hysteresis band. The various models of sunk-cost hysteresis all give rise to some common comparative statics: higher sunk costs lead to larger hysteresis bands, and hysteresis bands are sensitive to the persistence of individual shocks and the degree of market uncertainty.

Empirically, there has been some work done to model declining industries as well as entry-exit decisions by firms. Dunne et al. (1988) examine entry-exit decisions across markets. They aggregate data in various manufacturing markets and find two main trends: the entry rate and exit rate are highly correlated across industries, and within an industry, a spike in the entry one period is highly correlated with a spike in the exit in the following period. Lieberman (1990) examines plant-level exit decisions of chemical producers in declining chemical markets using a discrete-choice model. The main goal of the research was to determine whether small producers are shaken out of the declining industry by large producers, or if small producers stake out the market and force larger producers to scale down or exit earlier. The Lieberman results provide evidence in support of both the shake-out and stake-out theories. Harrison (1994) analyzes the welfare and trade implications of the declining US steel industry by setting up a trade model. He creates a structural model that matches some industry averages and widely accepted parameters, and then runs simulations over many hypothetical future industry outcomes. The simulations showed small welfare gains to subsidizing both large and small domestic steel producers. Dunne et al. (2005) examine the relationship between firm experience, entry characteristics and entry-exit decisions. They divide firms into different categories corresponding to their experience in different product lines, and find patterns of entry and exit that vary vastly by firm experience and sector. Miller and Wilson (2017) estimate entry-exit policy functions for non-profit and for-profit hospitals using a discrete-choice logit model, and find differences in the objectives of for-profit and non-profit hospitals.

Much work has been done to empirically study the determinants of plant closure for multi-plant firms. Blonigen et al. (2013) study the entry-exit decisions of US Steel producers in various product lines, Meyer and Taylor (2015) analyze oil refineries in the US, and Bichescu and Raturi (2015) investigate industry dynamics and plant closing announcements. Across all industries, there is a common result: firms tend to shut down smaller plants, and larger firms that own more plants are more likely to shut down any given plant than a firm that owns fewer plants.

Attempts to estimate entry and exit bands in sunk-cost hysteresis have been made in many trade models, starting with Roberts and Tybout (1997). They create an empirical model of sunk-cost hysteresis in the Colombian manufacturing market to model plant-level decisions to enter and exit the export market. They find that exporter experience is a significant determinant in a potential exporter's decision to enter the exporting market and that sunk entry costs are high, giving rise to a large hysteresis band. Máñez et al. (2008) estimate a model of hysteresis using data at the firm level rather than the plant level from Spanish manufacturing firms, and find that larger manufacturing firms face significantly lower sunk costs into the export market than smaller firms. Bernard and Jensen (2004) estimate determinants of entry and exit into various exporting industries using the same empirical hysteresis specification as Roberts and Tybout. Rather than estimate the model using a probit specification like Roberts and Tybout, Bernard and Jensen estimate a linear probability model. In doing so, they find upper and lower bounds on the importance of sunk costs that are both highly statistically significantly different than 0.

There is a related and extensive environmental economics literature on exhaustible resources. In their review of the subfield of exhaustible resources, Slade and Thille (2009) examine the Hotelling Model wherein a firm has a finite amount of a resource to use over the life of the firm, and the extraction of the resource becomes costlier as more of it is extracted. They, along with others, attempt to estimate the model in oil, gas, and metal-mining industries.

In a review of the technology in the mining industry, Hitzman (2002) notes that the decline of the mining industry as well as the closure of the U.S. Bureau of Mines has caused a notable lag in mining research and development. He points out that although there is not an absolute lack of innovation in the mining industry, there have been far fewer groundbreaking improvements that were once so common to the field. Although it was not applied to a rapidly shrinking industry, Moel and Tufano (2002) provide a relevant empirical examination in the gold mining industry. They calibrate a real options model to determine when firms optimally choose to temporarily or permanently close gold mines. To the best of my knowledge, Eyer and Kahn (2017) present the first paper to empirically study the effects of the decline of the American coal industry. They find that on the aggregate, the decline of the coal industry has positive welfare effects through improvements in the environment, but the costs of the decline are spatially concentrated and can lead to local negative welfare effects in areas that depend heavily on coal mining.

Some research has been devoted to analyzing the effects of the rise of natural gas in particular. Jenner and Lamadrid (2013) provide a direct comparison of the environmental impacts of coal and natural gas extraction and electricity generation on air, land and water. They find that natural gas has a smaller negative impact on air quality, uses less water, and uses less land than coal, but natural gas may contribute to drinking water contamination. Knittel et al. (2015) studies the effects of power plants' decisions to switch between coal and natural gas in the wake of the shale gas boom for plants that are fitted to use both coal and natural gas. They find that the sharp drop in price led to a 19% drop in emissions in traditional energy markets and a 33% drop in restructured energy markets. Fell and Kaffine (2018) finds that the decline in power plants' coal use is the effect to the interaction of the fracking boom and the rise of wind energy.

My research extends the literature by estimating a discrete-choice entry and exit model in the coal industry in the style of Bernard and Jensen (2004). The coal industry falls into the exhaustible resource literature, literature on exit in declining industries and the literature on sunk cost hysteresis. While sunk-cost hysteresis is used almost exclusively in the trade literature, it has a very convenient application to this particular problem. In both the theoretical and empirical hysteresis literature, firms are believed to make their participation decisions as a response to a change in a price of a close substitute, normally a similar good sold in another country. I model a mines' participation decision as a response to the price of natural gas, its closest substitute in the electricity market. So, an empirical hysteresis model provides a convenient framework to analyze the coal mining industry's response to the sharp drop in the natural gas price seen post-2008. To the best of my knowledge, this is the first research to estimate an entry-exit model in the coal mining industry.

II.4 A model of coal mine entry and exit

To motivate the empirical results of this chapter, I begin by outlining a simple model of mine openings, idlings, and shut downs in the coal market. This model draws inspiration from the empirical hysteresis work of Roberts and Tybout (1997), Bernard and Jensen (2004) and Máñez et al. (2008). Omitting any fixed costs, let the single-period profit of operating a mine i under firm f at time t be denoted by $\pi_{ift} = \pi_i(z_i, x_t, m_{ift})$. Here, x_t is a vector of macro- and market-level variables that the mine takes to be exogenous including the national spot price of natural gas and policies that may affect the demand for coal, z_i is a vector of time-invariant mine-level characteristics such as its location and mining method, and m_{ift} is a vector of state variables that is specific to each mine. Previous literature has shown that both plant- and firm-level characteristics affect profitability. While the specification of m_{ift} is flexible, in my specification it includes a mine's history of production and employment decisions as well as its history of safety violations, remaining reserves, and number of other mines owned by the same firm. The state variables may evolve endogenously in response to exogenous changes in market characteristics. When put together, π_{ift} represents the added profit to firm f by choosing to operate mine i in period t instead of choosing to idle it or shut it down.

Sunk costs of entry and exit costs are potentially very significant in the mining industry, but are absent from π_{ift}. Assume that if mine i was last active in the coal mining industry j periods ago and chooses to re-enter the mining industry, it faces a fixed re-entry cost F_i^j and earns profit $\pi_{ift} - F_j^i$ in its first period after re-entry. In a similar fashion, assume that a de novo mine entrant incurs a fixed cost F_i^0 upon entry into the coal market, and earns a profit $\pi_{ift} - F_0^i$ in its first period of operation. A mine that continues to operate at time t will earn profits π_{ift}. A mine i also faces significant exit costs or scrap value, denoted by X_i.

The i subscipt allows the entry and exit costs to vary by mine, and be a function of a mine's size, location, mining method, coal type, etc. The j superscripts on the re-entry costs allow the cost of re-entry to vary based on how long a mine has been inactive. The factors F_i^0 and F_i^j capture the inherent differences in costs of creating a new mine and reopening an idle mine.

To combine these profit inflows and outflows into a single expression, I define the indicator variable Y_{ift} to take on a value of 1 if the mine is producing in period t, and 0 otherwise. Additionally, denote the vector of the mine's operating history by $\mathbf{Y}^-_{\mathbf{ift}} = \{Y_{if,t-j} | j = 0, \ldots, J_i\}$ where J_i is the number of periods since the mine's de novo entry. With this, the added profits to firm f of operating mine i at time t are given by $R_{ift}(\mathbf{Y}^-_{\mathbf{ift}})$ and take the form:

$$R_{ift}(\mathbf{Y}^-_{\mathbf{ift}}) = Y_{ift}[\pi_{ift} - F_i^0(1 - Y_{if,t-1})] - \sum_{j=2}^{J_i}(F_i^j - F_i^0)\tilde{Y}_{i,t-j} - X_i Y_{if,t-1}(1 - Y_{ift})$$

where $\tilde{Y}_{i,t-j} = (Y_{i,t-j}\Pi_{k=1}^{j-1}(1 - Y_{i,t-k}))$. Put simply, the variable $\tilde{Y}_{i,t-j}$ creates an indicator variable for the length of time between a mine's last active period and the period they choose to re-enter the industry. The variable $\tilde{Y}_{i,t-j}$ takes on a value of 1 if the mine was last active j periods ago and 0 otherwise. So, if a mine i was last active 5 periods ago, then $\tilde{Y}_{i,t-5} = 1$ and $\tilde{Y}_{i,t-j} = 0$ for $i \neq 5$. I assume that any mine that has been closed for more than eight quarters faces the same re-entry cost. As in, $F_i^j = F_i^k$ for $j, k \geq 8$.

Previous attempts to estimate hysteresis bands, including Roberts and Tybout (1997), have assumed that the cost to re-enter the industry is less than the cost of initial entry, and approaches the cost of initial entry as the length of a firm's time out of the market increases. Like other hysteresis literature, I include a term to capture the costs associated with longer time away from the market, however I assume that the maximum cost of re-entry differs from the cost of de novo entry. The difference in these costs is captured by $F_i^0 - F_i^8$.

Owners of a mine at time t are assumed to make a decision over an endogenous finite horizon to maximize the expected present value of operating. Define the optimal sequence of finite horizon decisions as $\mathbf{Y_{ift}^T} = \{Y_{if,t+j} | T \geq j \geq 0\}$ where T denotes the length of the finite horizon, as chosen by the mine owner. The maximized payoff is:

$$V_{ift}(I_{ift}) = \max_{T, \mathbf{Y_{ift}^T}} E_t \left(\sum_{j=t}^{T} \beta^{j-t} R_{ift} | I_{ift} \right)$$

where β is a single-period discount rate and expectations on the transition of market and macro variables are conditional on each mine's information set I_{ift}. Using Bellman equation notation, plant i's current operating decision for $t < T$ can be represented as the Y_{ift} value that satisfies:

$$V_{ift}(I_{ift}) = \max_{Y_{ift}} (R_{ift}(\mathbf{Y_{ift}^-}) + \beta E_t\{V_{if,t+1}(I_{if,t+1}) | \mathbf{Y_{ift}^-}\})$$

where E_t denotes the expected value conditional on the information set I_{ift}. Through some algebraic manipulation, the previous equations can be rewritten to characterize the decision to participate in the coal market rather than idle or exit entirely. The participation condition for mine i in firm f at time t is:

$$\begin{aligned} (1) \quad & \pi_i(z_i, x_t, m_{ift}) + \beta\big[E_t(V_{if,t+1}(I_{ift}) | Y_{ift} = 1) - E_t(V_{if,t+1}(I_{ift}) | Y_{ift} = 0)\big] \\ & \geq F_i^0 - (F_i^0 + X_i) Y_{if,t-1} + \sum_{j=2}^{J_i} (F_i^j - F_i^0) \tilde{Y}_{i,t-j} \end{aligned}$$

where $-(F_i^0 + X_i)$ is the sum of sunk entry cost for a new mine and the exit cost for a mine, known as the "hysteresis band" in the theoretical sunk cost hysteresis literature. Its value represents the range of profitability for which mines will not exit if they're in the industry, and won't enter if they're out of the industry.

Equation (1) provides the conditions for de novo entry, re-entry, and exit that will be estimated in the empirical section. Some results are immediately apparent from equation (1): In the absence of any de novo entry or re-entry costs, the participation condition collapses down to $\pi_i(z_i, x_t, m_{ift}) \geq 0$. There is no distinction between the exit costs of idling a mine or permanently shutting down a mine, but this is not to say that the decision to exit or shut down are equivalent. The decision to permanently shut down is captured by the choice of the finite horizon T. The levels of each F_i^j can provide insight into the sunk costs that can be recovered by a mine that re-enters. We should expect $F_i^j > F_i^k$ for $j > k$. I will test for the presence of entry and exit costs and the decay of the recoverable sunk costs of re-entry.

II.5 Empirical Specification

I begin deriving my empirical specification with a mine's participation decision from equation (1). For notational simplicity, define:

$$\pi^*_{ift} = \pi_i(z_i, x_t, m_{ift}) + \beta\big[E_t(V_{if,t+1}(I_{ift})|Y_{ift} = 1) - E_t(V_{if,t+1}(I_{ift})|Y_{ift} = 0)\big].$$

This represents the present value of the total added profit of choosing to operate in the coal market today rather than not operate. Then, by rearranging equation (1), we can denote the participation decision Y_{ift} in the coal market at time t as the discrete-choice:

$$(2)\quad Y_{ift} = \begin{cases} 1 \text{ if } & \pi^*_{ift} - F_i^0 + \sum_{j=2}^{J_i}(F_i^0 - F_i^j)\tilde{Y}_{i,t-j} + (F_i^0 + X_i)Y_{if,t-1} \geq 0 \\ 0 & \textit{otherwise} \end{cases}$$

In equation (2), I approximate profit over fixed costs as a reduced-form expression of individual time-invariant mine effects (z_i), a vector of macro-level characteristics (X_t), a vector of firm- and mine-specific state variables (M_{ift}) and noise (ε_{ift}). Within the vector (X_t), I include dummy variables for policies, the national natural gas spot price, and a linear time trend with a structural break at the end of the Great Recession.[23] The vector (M_{ift}) contains variables for the average number of employees a mine uses during a period, the total coal extracted from the beginning of the data until period t, firm characteristics, and the history of violations by mines. Therefore, the profit condition in equation 2 can be re-expressed as:

$$(3) \quad \pi^*_{ift} - F^0_i = z_i + \beta_1 X_t + \beta_2 M_{ift} + \varepsilon_{ift}$$

I am not be able to separately identify the sunk cost of entry and the exit cost, but I am be able to estimate the size of sum of entry and exit costs, which is referred to as the "hysteresis band." In order to estimate the hysteresis band, I employ a few identifying assumptions. Firstly, I assume that the sunk cost of de novo entry and re-entry do not vary between periods or mines, and are given by F^0 and F^j for $1 \leq j \leq 8$. Second, I assume that exit costs do not differ between mines or time periods, and are denoted by X

Using the above identifying assumption, I redefine $F^0 + X = \gamma^0$ and $F^0 - F^j = \gamma^j$. Additionally, let the variable $Idle_{ift,j}$ take on a value of 1 if mine i in firm f at time t has been idle for the previous j periods. By substituting these definitions and equation (3) into equation (2), my estimating equation becomes:

$$(4) \quad Y_{ift} = \begin{cases} 1 \text{ if} & z_i + \beta_1 X_t + \beta_2 M_{ift} + \gamma^0 Y_{if,t-1} + \sum_{j=1}^{8} \gamma^j Idle_{ift,j} + \varepsilon_{ift} \geq 0 \\ 0 & otherwise \end{cases}$$

The model allows for a number of testable hypotheses. First, the shale gas boom anecdotally has been the largest reason for the decline of the coal industry. A positive

[23]Preliminary data work shows that there is a structural break in the linear trend and a mean shift at the end of the Great Recession during the second quarter of 2009.

value on the coefficient attached to the real natural gas price would indicate that a large drop in the price of natural gas would lead to much lower profitability in the coal industry, leading to higher exit rates. Second, one would expect that the cost of re-entry is weakly increasing with time. I test the hypothesis $\gamma_1 < \gamma_2 < \cdots < \gamma_8$ after I estimate my model. Third, as a mine extracts coal, it is necessarily taking away coal that it could extract in the future, but the mine may become more efficient due to "learning by doing." I include variables for the log of the total cumulative extraction by a mine as well as the square of the log of cumulative extraction in M_{ift}. A positive value on the total extraction and a negative value on the square of extraction would indicate that mines learn to be more efficient through working but also reduce future profitability through extraction.

Finally, the empirical exit literature has generally found that plants owned by larger firms are more likely to exit than plants owned by smaller firms. As is consistent with the exit literature, I include the total mines owned by a firm and an indicator variable for a mine owned by a multi-mine firm. A negative value on these coefficients would support the literature's findings. The empirical exit literature has generally found that smaller plants are more likely to exit than larger ones. As is consistent with the exit literature, I include the log of the average number of employees as proxy for the size of the mine. A positive value on this would support the literature's findings.

II.5.1 Estimation Issues

I can control for much of the heterogeneity between mines using the vector M_{ift}, which contains mine- and firm-level variables, including the number of mines a firm owns, the average number of employees the mine has, the total coal extracted, the time since the mine's de novo entry, the mine's history of safety violations and the fines paid as a result of the safety violations. Even though M_{ift} can control for much of the individual-specific factors of mine profitability, I expect there to be other unobservable factors that affect profitability including the mine's remaining reserves, differences in technology and differences in managerial ability. Unobservable factors relating to a mine's profitability will be accounted for using z_i, which will be estimated using random-effects and fixed-effects specifications of a linear probability model.

Heckman (1981) points out that discrete-choice models where an individual makes repeated choices are prone to bias from two separate sources: individual heterogeneity and state dependence. In my model, any unobserved individual heterogeneity will be picked up by the z_i parameters. Heckman notes that the presence of state dependence will bias errors even with the inclusion of a lagged dependent variable. While there have been many ways to address this issue, the sunk-cost hysteresis literature has taken two main approaches: impose standard error restrictions on a panel-probit model or estimate a panel linear model with serially-correlated errors. I employ the second approach by clustering my standard errors at the mine level. Bernard and Jensen (2004) shows that a linear model with AR(1) errors and individual-specific fixed effects leads to attenuation bias in the estimation of the hysteresis band γ^0. So, my estimates of γ^0 should be interpreted as lower bounds of the true value of the hysteresis band. As a robustness check, I re-estimate my model using AR(1) errors and present the results of the estimation in the appendix. The results are qualitatively the same as my main specification.

II.6 Data Sources and Variables

II.6.1 Data Preparation

My dependent variable Y_{ift} in the empirical analysis is a mine's decision to participate in the coal mining industry, and it takes on a value of 1 if the mine participates. As is done in the work of Roberts and Tybout (1997), Bernard and Jensen (2004) and Máñez et al. (2008), the hysteresis band will be estimated using the coefficient attached to the lagged values of the dependent variable. Other mine-level independent variables will be a measure of firm scale, mine scale, cumulative coal extraction, cumulative safety and health violations, and the length of a mine's time out of the market.

The primary data on coal mining operations comes from quarterly Employment/Production Data Set from the Mines Safety and Health Administration (MSHA). This includes data on the activities of plants in the coal industry including their production, labor use measured by the average number of employees on hand and labor hour use, primary type of coal mined, and mining methods. These data were merged with other data sets released by MSHA containing information on each plant's location, ownership history, and safety violations history using a unique ID number asigned by the MSHA. From this, I construct a panel of data with observations on the plant-quarter level ranging from the first quarter of 2000 until the fourth quarter of 2016.

The raw data often contain multiple observations per plant-quarter. Upon inspection, whenever there are more than a single observation per plant-quarter, all but one of the observations contain no coal production and a negligible amount of labor hours used. Within each plant-quarter, I keep only the observation with the highest coal production. I also omit all plants that are observed to produce no coal through the entire course of the data, and all observations that occur after a mine is officially listed as "Inactive" or "Abandoned" by the MSHA's Mines Data Set. It assumed that any gaps in the data occur because a mine is inactive, and I fill in all gaps with an observation where a mine produces no coal and uses no labor.

Plants in the coal-mining industry can be broken down into two separate, but not mutually exclusive categories: mines and facilities. Mines are plants that extract previously untouched coal from the ground, and are the object of my research. Facilities are any plants that perform auxiliary functions such as processing or disposing of coal.[24] Within the data, a plant that acts as both a mine and a facility is labeled as a facility. Filtering out all observations that produce no coal over the course of the data eliminates most facilities, but leaves 7,421 plant-quarter observations out of 111,700 total observations that are labeled as facilities. These come from 165 unique plants out of the 4,004 total plants that are ever present in the data and used in the empirical analysis. I cannot determine whether these plants act primarily as coal mines or facilities, so I conduct my analysis assuming that they function mainly as regular mines and then drop them from the data and re-estimate my model as a robustness check. The results of this estimation can also be found in the appendix. The results are qualitatively the same with or without the facilities in the data set.

After filtering out facilities, the data from the MSHA contains an unbalanced panel of mines at the mine-quarter level, where a mine only enters the data after its first operational period and leaves the data after it permanently shuts down. Knowing this, I construct indicator variables for the period when a mine is a *de novo* entrant in order to estimate the hysteresis band. Any mine that enters the data after quarter 1 of 2000 is assumed to be a newly opened mine, and any mine that drops out the data prior to quarter 4 of 2016 is assumed to have permanently exited the coal mining industry. I assume that a mine participates in the coal-mining industry if their production for the quarter is positive. Using this, I create indicator variables to denote when a mine participates and when it chooses to idle. I interpret idling as a decision to temporarily exit the industry.

The empirical hysteresis work assumes that the length of a firm's absence from an industry affects its sunk costs of re-entry. Using the indicator variables to denote idling mines, I create a running variable totalling up the number of consecutive quar-

[24]Some previously inactive mines are used as refuse sites to dispose of waste coal. Until recently, these sites were environmental hazards, but are now often re-mined as another potential electricity fuel source. These are also difficult to separately identify from other true facilities and are not the main concern of this research. The included link discusses refuse sites in greater depth. http://www.powermag.com/coal-refuse-dilemma-burning-coal-environmental-benefits/?pagenum=1

ters that a mine is idle. In keeping consistent with the hysteresis literature, I create indicator variables for the total number of quarters that a mine has idled, with all mines that have idled at least eight quarters receiving the same indicator variable. Of the 4,004 mines I observe, 3,388 have idled for at least one period over the 16 years of my data, indicating that temporary exit and re-entry are common behaviors.

Within each panel, I capture the scale of each firm by summing up all the mines owned by a single firm in a period. The estimate on this coefficient estimate represents any possible economies or diseconomies of scale at the firm level. Finally, I merge data on each mine's controller history to denote when a mine is sold and how long the mine has been operational using MSHA's Controller/Operator History Data Set.[25] The MSHA's Violations Data Set details each health and safety violation incurred by a mine, the date it occurred, the type of violation and the fines associated with the violation. I aggregate the violations to the quarter level, and then sum up the total violations a mine has incurred throughout the course of the data. I do the same thing for the dollar value of the fines paid as well.[26]

Natural gas prices are taken from the Henry Hub Natural Gas Spot Price weekly time series. Prices are averaged at the quarter level, and reflect the nominal average quarterly price of one mmBtu. Using January 2001 as a base period, I put all natural gas prices into real dollar terms. There were few large environmental policy changes between 2000 and 2016, so I include only two: the Clean Air Interstate Rule in March 2005 and the revision to the NAAQS Standards in 2010. All observations are given a dummy variable equal to 1 for the NAAQs revision in 2010. The Clean Air Interstate Rule only covers power plants in the eastern US, so mines in these states are given a

[25] All mines that opened prior to 1950 are listed as having opened in 1950 in the controller and operator history data set.

[26] I do not use the dollar amount of the violations as a regressor due to collinearity with the total number of violations and endogeneity. According MSHA's rules for assessing fines, the fines charged for a violation is a function of the type of violation, the number of employees a mine has and the amount of coal it produces. This makes it endogenous to two regressors in my model, and it can be shown that the fines paid is highly collinear to my measures of mine scale and total violations.

Figure II.5
Net Coal Mine Entry and Natural Gas Price

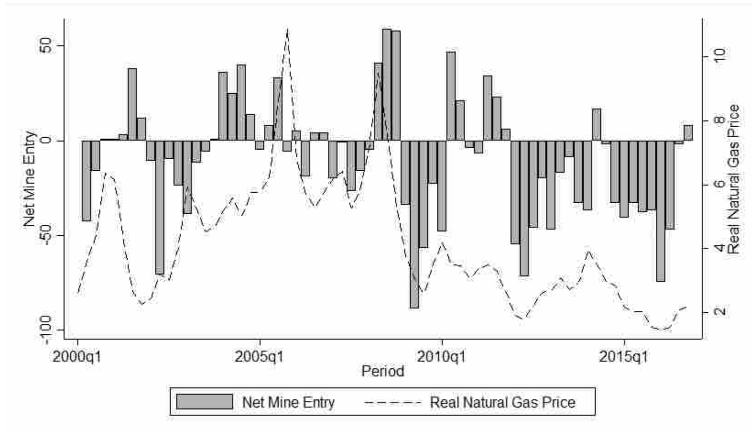

The real natural gas price is measured in January 2000 $'s per mmBtu.

dummy variable equal to 1 after March 2005.[27]

II.6.2 Summary Statistics

In Figure II.5, I show the net mine entry per quarter and the real natural gas price. Prior to 2010, coal mine openings generally outpaced closure, causing the number of mines in the market to grow. However, from 2010 until the end of the panel, exits outpaced entrances. Of greater interest, spikes in the natural gas price appear to be highly correlated with spikes in the net entry of coal mines.

I present the total number of entries, de novo entries and exits per quarter over the course of the data as well as the total number of operating mines per quarter in Figure II.6. The figure demonstrates that the switch in net entry is due in large part to de novo entry and re-entry into the market dwindling in final years of the data. This switch in entry and exit patterns occurs at the same time as the sharp drop off in the natural gas price seen in Figure II.5 and the drop off in aggregate coal production and the rise in natural gas production shown in Figure II.4. It is also worth noting

[27]A map of states covered by the Clean Air Interstate Rule is included in the link below https://archive.epa.gov/airmarkets/programs/cair/web/html/index.html

that the de novo entry is essentially 0 starting in 2015, which is a full year before the new federal coal mine moratorium imposed at the end of President Obama's tenure.

In Figure II.7 I display the ownership characteristics of mines across time. Mining becomes more concentrated with time, as firms that own a single mine become scarcer while multi-mine companies continue to add more mines. Although the total number of single-mine owners is essentially cut in half over the course of the data, only 51 mines owned by a single-mine owner were sold to multi-mine firms. Instead, the decline in single-mine owners primarily through single-mine owners choosing to open up a second mine or exit the mining industry entirely. As was the case in Figure II.6, Figure II.7 shows that the total number of mines owned by both single- and multi-mine owners declines after around 2010, with multi-mine owners shutting down far more mines than single-mine owners.

II.6.3 Exiters

Between the second quarter of 2000 and the fourth quarter of 2016, 2,096 mines of the 4,004 total mines that were ever active over the course of the data shut down and permanently exited the coal industry. Within that span, I observe 1,296 mines that are both *de novo* entrants and permanent exiters in the data. Under the hysteresis

Figure II.6
Mine Entry, Mine Exit, and Total Market Participants

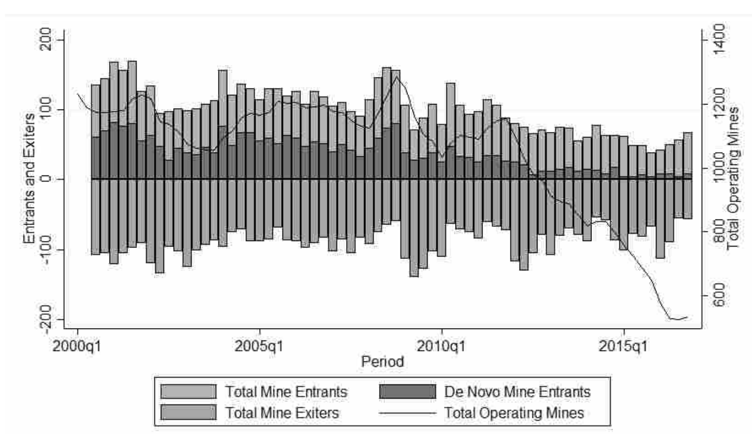

Figure II.7
Single-Mine Owners v. Multi-Mine Owners

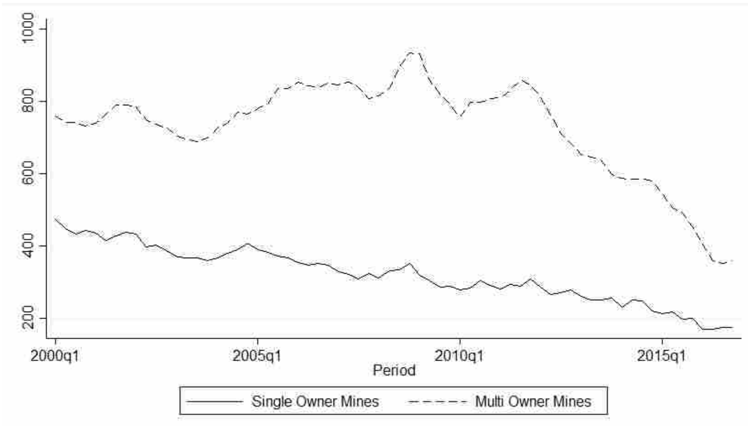

framework, there are only a handful of reasons to believe that a permanent exit occurred:

(1) A mine extracted all recoverable coal that was economically viable, making a firm believe that it is impossible to ever reach the entry hysteresis band.

(2) Natural gas became so inexpensive that all the mine's potential buyers substituted coal for natural gas, leaving no buyer for the mine; and

(3) Environmental policy or safety and health violations made a mine too expensive to profitably operate.

To address concern (1), I employ variables to indicate how much coal has been extracted over the course of the data and individual mine coefficients. Point (2) is the main parameter of interest, and is analyzed by merging in data on the quarterly average spot price of natural gas, as compiled from the Henry Hub data series.[28] The effects of environmental policy are of secondary interest and will be analyzed by adding in a dummy variable when and where the policies go into effect. All safety and health violations are totaled for each quarter for each mine, with the dollar amounts of the fines noted. These will be used as additional controls. Coal production is a

[28]This data series can be found at the weekly level at https://fred.stlouisfed.org/

decision that is made only after the decision to participate in the industry, so I do not use it as a regressor, but I include it for expository purposes. Table II.1 contains mine-level characteristics subdivided by each mine's status as a permanent exiter or a surviving market participant. I do not separate mines that have ever idled from the rest of the sample because a large majority of mines are observed idling at some point in the data. In general, mines that shut down come from larger firms, produce less coal, and have fewer employees.

II.6.4 Entrants

Because this research focuses on a mine's propensity to both enter and exit the coal industry, it may be important to note that the characteristics of a *de novo* entrant may vary inherently from its incumbent counterpart. Over the course of the data, 2,511 out of the 4,004 total mines mines are *de novo* entrants. Table II.2 provides a summary of entrants compared to the rest of the sample in a way that is analogous to

Table II.1
Summary Statistics: Non-exiters vs. Permanent Exiters

	Present in 2016q4		Shut Down	
	5/95%tile	mean	5/95%tile	mean
Coal Production (Short Tons)	1147.00	366227.07	1376.00	112922.61
	1524640.00		508779.00	
Avg # of Employees	2.00	70.45	2.00	45.32
	320.00		184.00	
Violations in Quarter	0.00	14.99	0.00	13.57
	71.00		62.00	
Fines Paid in Quarter ($)	0.00	7582.72	0.00	6792.22
	39692.00		29988.00	
Multi-Mine Firm (0/1)	0.00	0.69	0.00	0.81
	1.00		1.00	
Total Mines in Company	1.00	14.98	1.00	22.46
	62.00		112.00	
Observations	39106		32109	

All observations where a mine is idle are omitted.

Table II.2
Summary Statistics: De Novo Entrants vs. Incumbents

	Present in 2000q2		De Novo Entrant	
	5/95%tile	mean	5/95%tile	mean
Coal Production (Short Tons)	1330.00	413545.00	1146.00	90006.18
	1525022.00		297367.00	
Avg # of Employees	2.00	79.67	2.00	38.51
	336.00		134.00	
Violations in Quarter	0.00	16.28	0.00	12.42
	83.00		54.00	
Fines Paid in Quarter ($)	0.00	8495.11	0.00	5953.69
	44108.00		27539.00	
Multi-Mine Firm (0/1)	0.00	0.71	0.00	0.78
	1.00		1.00	
Total Mines in Company	1.00	18.36	1.00	18.34
	111.00		83.00	
Observations	35661		35554	

All observations where a mine is idle are omitted.

Table II.1. In general, the incumbents tend to be larger mines in both employee count and coal extracted, and tend to incur more violations. Apart from these dimensions, there does not appear to by any noticeable difference between a *de novo* entrant and a mine that began its life before the data.

Table II.3 directly compares the characteristics of a mine that entered the data after the second quarter of 2000 and a mine that shut down prior to the fourth quarter of 2016. Time-varying values are measured at the period of entry for de novo entrants and the period prior to exiting for mines that shut down. Over its lifetime, mines incur an average of almost $120,000 of fines prior to shutting down. It is worth noting that coal production and the average number of employees seem to change over the life of the mine. This can be seen by comparing Tables II.1, II.2, and II.3. A mine has fewer employees and produces less coal on average during its period of de novo entry than an average de novo entrant has during a typical quarter throughout the course of the data. The same can be seen with permanent exiters. An exiting mine has fewer employees and produces less coal in the period prior to exiting than the mine did during a typical quarter over the rest of its life in the data. This provides some weak evidence that coal mines have some life cycle of productivity.

The concentration and proportion of entering and exiting mines varies greatly between coal basins as well. In Table II.4, I present the total number of mines present at some point in the data as well as the number of mines that are de novo entrants or permanent exiters by basin. Immediately, it is evident that Appalachia has both the largest share of mines in the coal industry and the largest share of mine closures, particularly the Central and Souther parts of Appalachia. Despite playing a very prominent role in the United States' electricity market, the western Uinta and Powder River Basins have very few mines and a large share of mines that are always present in the data. This is mainly due to the scale of the mines and the high demand for their coal, as both basins are home to some of the biggest mining operations in the United States and produce low-sulfur sub bituminous coal.

Table II.3
Summary Statistics: De Novo Entrants v. Permanent Exiters

	Period of De Novo Entry		Period of Shut Down	
	5/95%tile	mean	5/95%tile	mean
Coal Production (Short Tons)	0.00	9245.33	0.00	6081.73
	44540.00		29568.00	
Avg # of Employees	0.00	7.36	0.00	5.71
	28.00		27.00	
Cumulative Violations Incurred by Mine	0.00	0.65	0.00	228.24
	4.00		982.00	
Cumulative Fines Paid by Mine	0.00	60.26	0.00	117051.78
	193.24		421245.53	
Multi-Mine Firm (0/1)	0.00	0.73	0.00	0.77
	1.00		1.00	
Total Mines in Company	1.00	13.80	1.00	17.52
	62.00		96.00	
Observations	2489		2074	

All observations where a mine is present for only a single period are omitted.

Table II.4
De Novo Entrants and Permanent Exiters By Basin 2001-2016

	Total Mines	Always Present	Entrant Only	Exiter Only	Entrant & Exiter
Central/So. Appalachia	2,929	358	915	560	1,096
No. Appalachia	782	240	225	170	147
Illinois	139	27	52	29	31
Powder River	34	23	2	6	3
Uinta	42	15	7	13	7
Other Coal Beds	78	26	18	26	8
Total	4,004	689	1,219	804	1,292

II.7 Empirical Results

II.7.1 Base Model Results

In Table II.5, I present estimates of my base model and perform a similar bounding exercise to Bernard and Jensen (2004). Column 1 contains estimates of the linearly probability model without mine-level random effects or fixed effects. Column 2 contains the same model with added mine-level random effects. Column 3 contains my preferred specification: mine-level fixed effects. Standard errors are clustered at the mine level across all three columns. The empirical models of hysteresis advocate using AR(1) errors with mine-level effects. As a robustness check, Appendix Table II.5A contains estimate of the same models in columns (2) and (3) with AR(1) standard errors. As a further robustness check, I re-estimate my model after dropping all mines that are labeled as a facility. These results are contained in Appendix Table II.5B. The coefficient estimates do not change significantly between Tables II.5, II.5A and II.5B. I perform a Hausman test on the random-effects model in column 2 and reject it in favor of the fixed-effects model in column 3.

It should be immediately apparent that mine-level heterogeneity heavily biases coefficient estimates in the models of columns 1. As pointed out by Bernard and

Jensen (2004), the omission of individual effects leads to extenuation bias in the estimate in the hysteresis band γ_0. In my preferred specification, I find that the estimate of the hysteresis band is $\gamma_0 = .223$, indicating that a mine that chose to participate in the coal mining industry in the previous quarter is is 22.3 percentage points more likely to participate in the coal mining industry this quarter. This value changes very little when I consider only non-facility mines in Table II.5B.

Table II.5
Base Linear Probability Model Results

VARIABLES	(1)	(2)	(3)
Real Nat. Gas Price	0.00401	0.00331	0.00312
Log(Avg # of Employees)	0.158	0.206	0.216
Log(Total Extraction)	0.0806	0.0901	0.0878
Log(Total Extraction)^2	-0.00500	-0.00595	-0.00581
Multi-Mine Firm (0/1)	-0.0505	-0.0301	-0.00587*
Total Mines in Company	-0.000253	-0.000135‡	0.000150†
Cumulative Violations Incurred by Mine	-2.55e-05	1.46e-05	2.47e-05
γ_0	0.373	0.239	0.223
γ_1	-0.323	-0.219	-0.195
γ_2	-0.377	-0.250	-0.222
γ_3	-0.406	-0.264	-0.233
γ_4	-0.421	-0.270	-0.237
γ_5	-0.438	-0.282	-0.249
γ_6	-0.435	-0.275	-0.240
γ_7	-0.440	-0.277	-0.242
γ_8	-0.501	-0.287	-0.244
Clean Air Interstate Rule (0/1)	-0.0197	-0.0256	-0.0246
NAAQS Revision (0/1)	0.0142‡	0.0108*	0.00876*
Time Trend	0.00210	0.00225	0.00191
Mean Shift for t>2009q1	0.332	0.345	0.288
Trend for t>2009q1	-0.00176	-0.00180	-0.00150
Observations	93,837	93,837	93,837
Mine Effects	N/A	RE	FE
Number of Mines Used		3,981	3,981

The real natural gas price is given in dollars per million BTU and total extraction is measured in short tons.
Standard Errors are clustered at the mine level.
* $p \not< 0.1$, † $p<0.1$, ‡ $p<0.05$

The coefficient estimates of γ_j for $j = 1, \ldots, 8$ represent the portion of the fixed costs that is unrecoverable by re-entering the coal mining industry after an absence of j periods. In my preferred specification, the absolute value of γ_j does not rise monotonically as j rises but does trend upward, providing evidence that re-entry becomes more costly as a mine spends more time idling. I test the hypotheses that any pair of γ_j is equal and fail to accept them all at the 1% level. In all three columns, I include dummy variables for when and where the Clean Air Interstate Rule went into effect and when the revision to NAAQS occurred. According to my preferred specification, the Clean Air Interstate Rule led to approximately a 2.46 percentage point decrease in the probability that a mine participates in the market, while the NAAQS revision led to a .8% percentage point *increase* in a mine's propensity to participate in the market. This provides some evidence that the effects of these policies are indeed statistically important, but are quite small.

II.7.2 Natural Gas Effects

The coefficient attached to the price of natural gas is of primary interest. According to column (3), every rise in the price of natural gas by \$1 leads to about a .312 percentage point drop in the probability that any mine would operate in the coal industry. After the shale gas boom, the price of natural gas dropped by approximately \$8, leading to approximately a 2.5 percentage point drop in the probability that a mine would participate in the coal mining industry every quarter compared to natural gas's peak price. While this may seem small, this large drop entirely offsets the effect of the revision to the NAAQS, and is approximately 11% of the size of the hysteresis band.

II.7.3 Mine Characteristics

At first glance, there are results that are consistent with the existing exit literature: The highly significant positive coefficient attached to the log of the average employee count suggests that large mines are more profitable than their smaller counterparts, thus less likely to close. This result is consistent with the empirical exit findings of Blonigen et al. (2013), Meyer and Taylor (2015), Miller and Wilson (2017) and Bichescu and Raturi (2015).

Across the first two specifications, the coefficients attached to total numbers of mines in the company and the indicator for a multi-mine firm are both negative and highly statistically significant. This indicates that there are diseconomies of scale associated with adding mines to a firm, meaning that as a firm adds more mines, the individual profitability of a single mine falls. This suggests that all else being equal, a mine in a large firm is more likely to exit the coal mining industry than a similar mine owned by a smaller firm. However, this statistical significance of the multi-mine firm indicator is lost in my preferred specification and the total number of mines becomes small and positive. I attribute this to very few mines being observed changing ownership from single-mine firms to multi-mine firms. This would cause the mine-level fixed effect to capture much of the effect of being owned by a large firm.

As was discussed in the empirical section, mines' profitability change ambiguously as they choose to extract coal for two competing reasons. First, as a mine extracts coal we should expect it to become more efficient through "learning-by-doing", a phenomenon first pointed out by Arrow (1962). This should lead to extraction having a positive effect on a mine's participation decision. Second, a mine only has a finite amount of recoverable coal. No matter how efficient a mine becomes through learning by doing, eventually coal extraction will become so prohibitively expensive that the mine will have to close. The coefficient on $Log(Extraction)$ is positive and the coefficient on its square is negative, indicating that initially the positive effect of "learning-by-doing" outweighs the negative effect of depleting a mine's reserves. The negative coefficient on $Log(Extraction)^2$ provides evidence that in the long run, the effect of depleting reserves outweighs the effects of learning by doing. By solving the polynomial, I find that the finite resource problem becomes more important after a mine has extracted approximately 3,700 short tons of coal. Most mines in the data produce more than 1,800 short tons of coal in a single period, so the effects of depletion significantly outweigh any potential learning by doing.

II.7.4 Policy Effects

A large talking point in the previous election cycle was the "war on coal," and removing stringent policies in order to induce re-entry into the coal-mining industry. I include dummy variables for when and where the Clean Air Interstate Rule (CAIR) went into effect, and when NAAQS was updated to allow much lower emissions. In my preferred specification, I find that the Clean Air Interstate Rule had a significant, negative effect on a mine's participation decision, and that the revision to NAAQS had a small but statistically significantly *positive* effect on a mine's participation decision. The Clean Air Interstate Rule went into effect in 2005, which was well before the intensely documented decline of coal mining, and the revision to NAAQS occurred in 2010, *during* the start of coal's decline. This provides some weak evidence that the recent large decline in the coal mining industry is not due to any recent environmental policies, but rather to firm- and mine-level characteristics and the sharp drop in the natural gas price.

The positive effect of the NAAQS revision is counterintuitive, but it may be a result of investment decisions on the part of power plants. The Clean Air Interstate Rule went into effect in 2005 while the price of natural gas was very high. To meet tighter air quality standards, fossil fuel power plants could make two decisions: switch over natural gas or install scrubbers on their current equipment. If a power plant were to install scrubbers, the plant would be even less likely to switch over to natural gas generation than before. So, if the price of natural gas is restrictively high, power plants may choose to install scrubbers, effectively reaffirming their commitment to using coal in the future and decreasing a power plant's probability of converting to natural gas.

A 2013 report filed by the EIA provides some evidence to support this conclusion.[29] According to the report, 91 gigawatts of coal-fired power capacity was retrofitted with scrubbers between 2005 and 2011. Additionally, it noted that many power plants put forth effort to further limit coal-related emissions by employing catalyzers to reduce NOx emissions during the generation process. By 2011, 67% of coal-fired power plants had some form of catalyzer installed. This led to a total reduction of both SO2 and NOx emissions over the period, despite a rise in the quantity of coal used by the electricity sector.

II.7.5 Heterogeneous Basin Effects

Table II.4 showed that mine entry and exit varied drastically between coal basins. So, one may expect that the sudden drop in the natural gas had a heterogeneous effect on mines in different basins. In this section, I re-estimate the models of Table II.5 with an interaction term between the basin and the real natural gas price in Table II.6. All other covariates in the regression are identical to the model estimated in Table II.5.

These results suggest that some regions are not very susceptible to the natural gas price. In particular, there is no significant effect on mines in the Illinois Basin in any specification. Additionally, the coefficient to the Powder River Basin is statistically significant and *negative*, indicating that natural gas's decline helped mines in the

[29] https://www.eia.gov/todayinenergy/detail.php?id=10151

Table II.6
Heterogeneous Basin Effects

Basin · Real Nat. Gas Price	(1)	(2)	(3)
Central & So. Appalachia	0.00205***	0.00232***	0.00452***
Northern Appalachia	0.0125***	0.00865***	0.00243*
Illinois	0.000542	-0.00224	-0.00162
Powder River	0.00558	-0.00200	-0.00549**
Uinta	-0.00967***	-0.00668**	7.22e-05
Other Coal Beds	-0.00186	-0.00374*	-0.00523**
Observations	93,837	93,837	93,837
Mine Effects	N/A	RE	FE
Number of Mines Used		3,981	3,981

The real natural gas price is given in dollars per million BTU.
Each column represents the coefficient of the regressions of Table II.5 with an interaction term between the real natural gas price and an indicator for the coal bed. All other regressors are identical to column (3) and don't change significantly
Standard Errors are clustered at the mine level.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Powder River Basin. Appalachia has anecdotally been the region hit hardest by the shale gas boom, and the positive coefficients attached to the interactions between Central & Southern Appalachia and Northern Appalachia with the real natural gas price provides evidence to support this. Based on column (3) of Table II.6, the drop in the natural gas price of approximately $8 in late 2008 would make a mine in Northern Appalachia approximately 1.9 percentage points more likely to close down each period, and a mine in Central or Southern Appalachia approximately 3.6 percentage points more likely to close each period.

II.8 Concluding Remarks

In this chapter, I study a coal mine's participation decision by using an empirical sunk-cost hysteresis framework with a linear probability model. After controlling for mine-level heterogeneity and serial correlation, I find that the sunk cost of de novo entry is still large and significant. As a result of the significant sunk costs of entry, a mine's probability of participation in the coal market is 22.3% higher if it participated in the coal mining industry in the previous period. After estimating my empirical model, I analyze the characteristics that lead to mine closure, and find that the characteristics are largely similar to the existing body of exit literature.

Of greater interest, I estimate the effects of environmental policy and the natural gas price on a mine's entry and exit decision. Surprisingly, environmental policies that were enacted during coal's decline had a positive effect on a coal mine's propensity to participate in the coal industry. On the other hand, the natural gas price has a large, significant effect on a mine's participation decision. On average, a mine was 2.5 percentage points less likely to participate in the coal mining industry each quarter after the shale gas boom in 2008. When I break down this effect by basin, I find that Appalachia is the most susceptible to the natural gas price, with the shale gas boom causing a 1.9 percentage point drop in the chance that a mine from Northern Appalachia participates and a 3.6 percentage point drop in the chance that a mine from Central or Southern Appalachia participates in the coal industry each period. These effects are much larger in magnitude than any policy intervention in my model, and are dwarfed by the significant sunk cost of entry. This provides strong evidence of the effects of sunk-cost hysteresis in many theoretical models: a shock to market forced many mines out out of the market, and the cost of re-entry is so prohibitively high that re-entering the market is not feasible. It also suggests that policy interventions of the same magnitude as the policies discussed in this research will not solve the decline in the coal industry. Rather, coal's decline is almost entirely a product of its closest substitute becoming cheap and moving the industry away from a profitability level high enough to induce large-scale entry.

CHAPTER III

EFFECTS OF A CARBON TAX ON FOSSIL FUEL GENERATOR DISPATCH ORDER: EVIDENCE FROM RGGI

III.1 Introduction

For the better part of the 21st century, fossil fuels, particularly coal, formed the backbone of the United States' electrical grid and many other developed countries despite increasingly stringent environmental regulations at the state, national and international level and the rise of renewable energy. Fossil fuels still remain the most heavily used fuel sources in electricity generation even though the portfolio of fossil fuels has changed. A notable environmental policy instituted in 2009 is the Regional Greenhouse Gas Initiative (RGGI), wherein ten states in the northeastern United States agreed to implement the United States' first ever Cap-and-Trade program on carbon emissions in the electricity sector.[1] While neither RGGI nor any other environmental policy appears to have changed the total use of fossil fuels in the US or even the states affected by RGGI, the mix of fossil fuels has drastically changed

[1] Cap-and-Trade programs, also referred to as emissions trading programs, are a type of environmental policy wherein some pollutant is controlled by a governing body that sets a binding cap on the quantity of the pollutant that can be emitted in a certain time and distributes the credits to emit the pollutant to polluting entities. Once the credits are distributed, holders of the credits are free to buy or sell the credits on a secondary market. The United States instituted a Cap-and-Trade program on sulfur in 1990 and on nitrogen oxide emissions in the 2000s. See Tietenberg (2010) for a summary of the theory and history of emissions trading programs.

since its implementation as well as the role of different types of fossil fuels. This can be seen in Figure III.1, where I provide a year-by-year area chart on the sources of all generated electricity since 2001.[2]

A few things are apparent from Figure III.1. First, apart from a dip in total electricity generation caused by the 2008 recession, the total amount of electricity generated has stayed relatively constant from year to year in the US as a whole, and the states in the RGGI region have only slightly lowered their net electricity generation. This provides some evidence that any change in the quantity of electricity generated by one source will be offset by a nearly-direct substitution to other sources. Second, coal and natural gas together generate the majority of the total electricity in the US and approximately half of the electricity in the RGGI region in any given year, while renewable energy sources are still not responsible for very much of the United States' overall portfolio of electricity generation. Third, while the amount of electricity generated by coal has substantially dropped from 2001 to 2016, the

Figure III.1
Yearly Aggregate Net Electricity Generation by Source

(a) US Total **(b) RGGI States**

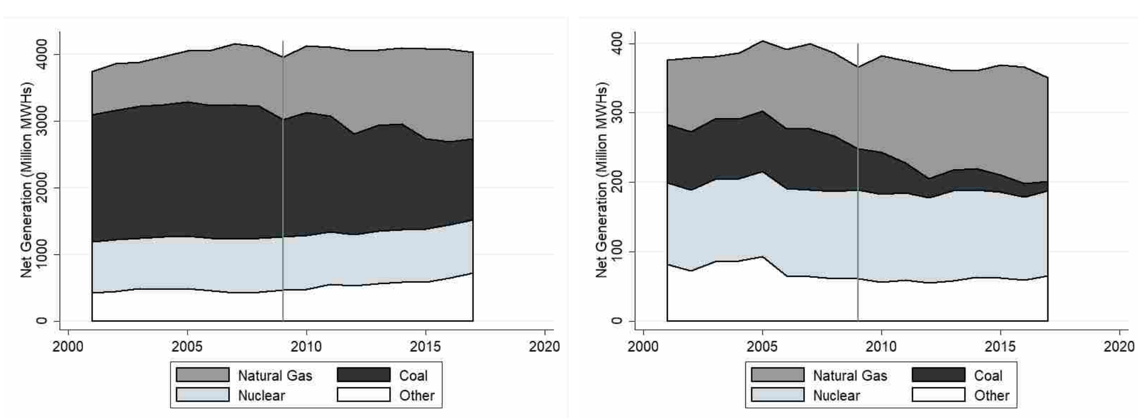

Note: New Jersey dropped out of RGGI in 2012 but is included in the right panel. Data are gathered from EIA-FERC Form 923.

[2]The ten states that agreed to the Regional Greenhouse Gas Initiative are Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Rhode Island and Vermont. New Jersey pulled out of the agreement in 2012.

total amount of electricity generated by natural gas and coal combined has stayed relatively constant. Finally, the transition away from coal to natural gas seems more pronounced in the RGGI region than in the rest of the United States post-RGGI.

The heterogeneity of carbon emissions from coal and natural gas is the mechanism behind the switch from coal to natural gas in RGGI states. In particular, coal emits approximately twice as much carbon per unit of electricity than natural gas (Energy Information Administration (2019)). Although all fossil fuels emit carbon, this difference causes a carbon tax to have an asymmetric effect on the post-tax price of coal and natural gas. In previous research Murray and Maniloff (2015) show that RGGI is responsible for approximately half of all carbon-emissions reductions in RGGI states. Using a synthetic control estimator, Kim and Kim (2016) provide evidence that RGGI caused its member states to change their coal-to-gas consumption ratio at a higher rate than states that are not affected by RGGI . Fell and Maniloff (2018) extends this to show that coal generator utilization falls in RGGI states while natural gas generator utilization rises in RGGI states and states that likely export electricity to RGGI regions.

Unlike the existing literature, I examine the effects of RGGI on seasonal power plant use. It is a well-known fact that electricity demand is higher in the summer and winter than it is in the spring and fall. Indeed, in Figure III.2 I disaggregate electricity production to the monthly level. At both the national level and the RGGI level, large spikes in electricity use occur in the winter, small spikes occur in the summer and troughs occur in the spring and fall. Of greater importance, Figure III.2 also demonstrates that the majority of these seasonal changes in electricity demand are supplied through coal and natural gas plants. In panel (b) of Figure III.2, it is apparent that after the institution of RGGI in 2009, coal electricity generation in RGGI states almost entirely disappears in low demand months, but is still present in higher demand months.

In contrast to previous papers researching carbon taxes, I analyze how a carbon tax affects the seasonality of the sources of electricity generation. I motivate my empirical analysis by creating a theoretical model of generator dispatch, and use it to demonstrate a carbon tax leads to natural gas generators moving up the dispatch order relative to coal generators. By moving up the dispatch order post-carbon tax, natural gas generators displace coal generators as a form of baseload electricity generation.[3] This leads to large drops in coal use in low electricity demand periods and smaller drops in higher demand periods, and increased natural gas use across all months. This provides evidence that a carbon tax incentivizes utilities to use natural gas generators as a form of baseload electricity generation and use coal generators to supplement spikes in demand.

In my empirical analysis, I use a fixed effect differences-in-differences estimator to determine the effects of RGGI on a power plant's decision to use coal and natural gas at the monthly level. This allows me to separate the change in power plant behavior

Figure III.2
Monthly Aggregate Net Electricity Generation by Source

(a) US Total **(b) RGGI States**

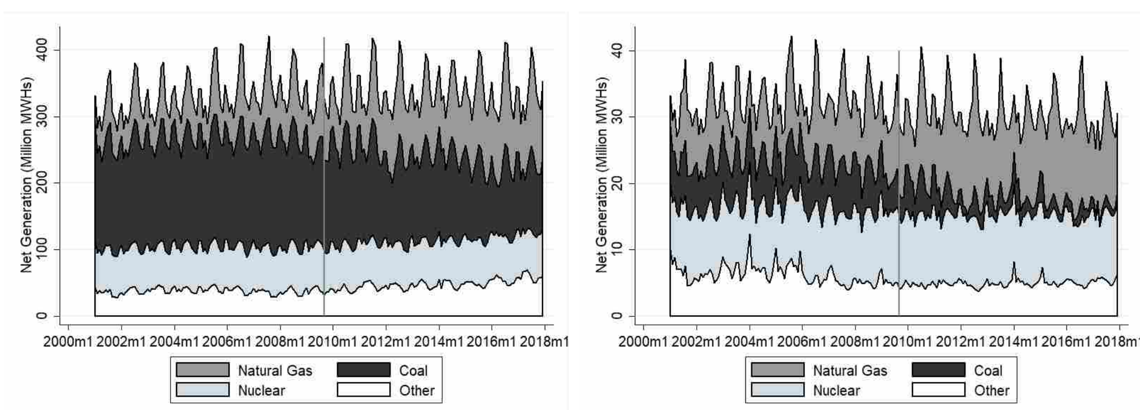

Note: New Jersey dropped out of RGGI in 2012 but is included in the right panel. Data are gathered from EIA-FERC Form 923.

[3]Baseload is defined as the minimum level of demand on an electrical grid over some span of time. In this paper, I consider the relevant span of time to be one month.

in periods of sustained high demand from periods of sustained low demand. I find that after RGGI, natural gas generators use more natural gas in all months, and the change in their month-to-month generation post-RGGI is consistent, providing evidence that natural gas generators moved up the dispatch order after RGGI and are being used more often for baseload electricity generation.

I find the opposite effect for coal generators. As is consistent with the findings of Fell and Maniloff (2018), coal electricity generation at the plant level fell in all months. However, I also find that the decline in coal use at the plant level is most pronounced in off-peak demand months. This suggests that coal is being used less often as a form of baseload electricity generation and instead being used to supplement seasonal changes in demand as a result of moving back in the dispatch order.[4]

As climate change becomes one of the largest global problems, novel programs to curb carbon emissions are seen as a policy priority. Carbon taxes are quickly becoming a popular way to disincentivize fossil fuel use, as evidenced by the new programs created by the European Union and California's Air and Resource Board. This paper contributes to the emerging field on the effectiveness and implications of carbon emissions trading schemes. The results can be used to inform policy makers on the expected outcomes of a new carbon tax and contributes to literature in the fields of industrial organization and environmental economics.

The rest of the chapter is structured as follows: Section 2.2 provides a brief overview of the RGGI program and other carbon emissions trading programs, and a background of the substitutability of coal and natural gas. Section 2.3 contains a review of the literature on the electricity industry, particularly on fossil fuels and generator dispatch choice. Section 2.4 provides a simple theory of generator dispatch order and its expected change after a carbon tax. In Section 2.5, I specify my main

[4] I conduct my analysis using data compiled from the Energy Information Administration (EIA) and the Federal Energy Regulatory Commission (FERC). The main data come from the EIA-FERC Forms 906, 920 and 923, the RGGI auction results and are supplemented by the EIA Form 860 and the Henry Hub natural gas spot price data series.

empirical model to be estimated as well as others that will be used as robustness checks. In Section 2.6, I discuss my data sources in greater depth, the process to assemble my data set and some summary statistics. I present the results of my empirical estimation, robustness checks, and potential policy spillovers in Section 2.7 and provide some discussion and concluding remarks in Section 2.8.

III.2 Fossil Fuel Energy Background

III.2.1 The Regional Greenhouse Gas Initiative

One of the primary drivers of the switch from coal to natural gas in the northeastern United States appears to be a new carbon tax that is not present in the rest of the United States. The Regional Greenhouse Gas Initiative (RGGI) was agreed upon by ten states in 2005 and went into effect in 2009, and was the first attempt to curb carbon emissions in the United States through a novel cap-and-trade program in the northeastern states. Unlike the emissions trading programs for NO_x and SO_2 emissions established by the Clean Air Act, the RGGI program does not directly give emissions credits to power plants. Instead, every quarter some quantity of pollution allowances is set by RGGI Inc. and entities bid on the right to own some number of pollution allowances in a sealed-bid auction with a price floor and a loose price ceiling. After the initial auction, owners of pollution credits can freely trade between each other like other cap-and-trade programs (RGGI-Inc. (2018)). I include a map of all coal and natural gas power plants present in 2016 that were affected by RGGI in Figure III.3.

Each auction has a minimum price for a pollution credit called the reservation price. The initial auction allows any bidder to submit a bid for both a quantity demanded of pollution credits and the price the bidder is willing to pay for each pollution credit. Bidders can submit any number of bids into the auction, allowing each

Figure III.3
Coal and Natural Gas Power Plants in RGGI Region

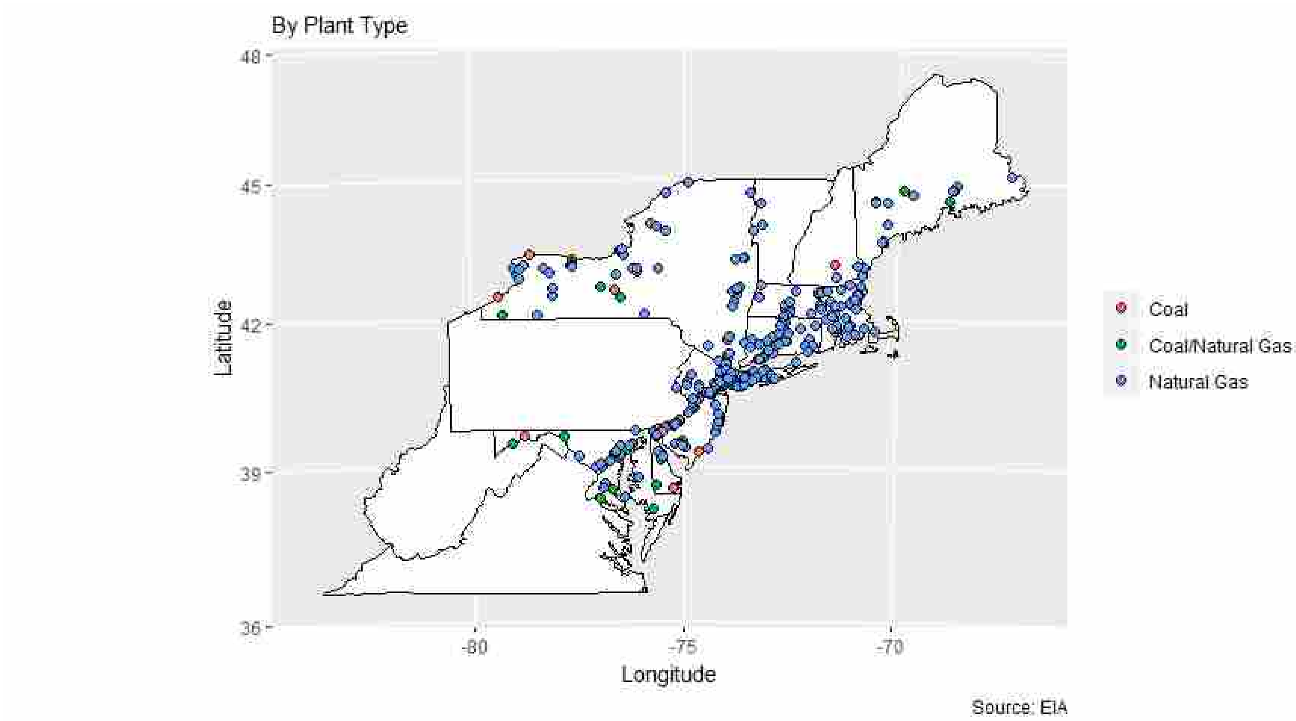

Note: This figure was generated using plants in the 2016 release of the EIA-FERC Form 860.

bidder to effectively create its own individual demand curve for carbon pollution credits. Once all the bids have been submitted, RGGI Inc. creates an aggregate demand curve, and uses the aggregate demand curve and the total allowance of carbon permits to determine an initial clearing price for a carbon permit. The initial clearing price is determined by the price at which the total quantity of pollution credits demanded by bidding entities meets or just exceeds the total number of pollution credits allotted by RGGI Inc. If the initial clearing price is higher than the reservation price, then the initial clearing price becomes the actual clearing price. If it is not above the reservation price, then the reservation price becomes the clearing price. All bidders who submitted a bid above the clearing price receive the quantity of pollution allowances they bid on at the clearing price, with marginal bids being randomly given to any bidders who submitted a bid at the clearing price.[5]

III.2.2 Other Carbon Pricing Programs

The Regional Greenhouse Gas Initiative is not the only market-based program to curb carbon emissions. Two other notable recent policies are the European Union Emissions Trading System (EU-ETS) and the recently-created California Air and Resource Board cap-and-trade carbon market. I provide a brief summary of these programs in this section.

In 2006, the EU-ETS became the first program to impose a limit on carbon emissions in the power sector. The program was divided into three phases which were meant to allow firms first to get accustomed to carbon pricing, and then to impose binding caps. In the first phase from 2005-2007, firms in electricity markets and other high-emission fields were allocated carbon permits and allowed to freely trade them. Due to an overallocation of permits and the inability to use the first phase permits past 2007, the price for a carbon permit quickly fell to zero.[6]

The second phase ran from 2008 until 2012, and was characterized by the introduction of permit auctions in some countries and a further lowering of the cap on carbon emissions. In the second phase, the penalty for non-compliance in the program also was increased from 40 Euros per ton emitted to 100 Euros per ton of carbon emitted and three new European countries joined the program. Due to the economic crisis of 2008, total economic activity and carbon emissions fell, leading to a sustained low carbon price. Throughout phase two, the total value of permits exchanged steadily

[5]Starting in 2014, RGGI instituted a loose price ceiling in their auctions called the Cost Containment Reserves (CCR) trigger price for the rare occasion of an abnormally high initial clearing price. If the initial clearing price is above the CCR trigger price and less than 10,000,000 CCR pollution credits have been released in the calendar year, the auction will release CCR pollution credits until the clearing price equals the CCR trigger price. If more than 10,000,000 CCR pollution credits have been released, then there will be no additional pollution credits released and the clearing price will remain above the CCR trigger price. In the 18 auctions since the start of the CCR program, the initial clearing price has only exceeded CCR trigger price only twice.

[6]See Ellerman and Buchner (2007) for an analysis of phase one of the EU ETS.

rose to over 56 billion Euros (European Commission (2019)).

Phase three of the EU-ETS runs from 2013 through 2020 and works to create a more uniform set of rules among all member countries. The most notable changes are a single EU-wide cap in carbon emissions that replace the disaggregated country-level emission caps. In phase three, auctions for carbon permits were made the standard allocation method instead of the free allocation method that was used more commonly in phases one and two.

In 2012, the California Air Resource Board instituted its own carbon cap-and-trade program with the goal of returning to the state's 1990 level of carbon emissions by 2020 and to further reduce emissions by 20% by the year 2030. The cap-and-trade program covered all entities emitting more than 25,000 metric tons of carbon annually starting in 2013. Unlike RGGI or the EU-ETS, California's program also covers all transportation or fuel distributors starting in 2015. Like the EU-ETS and RGGI, carbon-emitting entities are allowed to bank or trade pollution credits, but the cap on new carbon permits is reduced annually (California Environmental Protection Agency (2015)).

III.2.3 Price of Carbon

Figure III.4 shows the history of the RGGI clearing price. The cap on carbon permits was initially not binding, which led to the clearing price being set at the price floor from 2010 until 2013. After drastically lowering the cap on the number of carbon allowance permits, the nominal price of carbon steadily rose until its peak in 2013 of nearly $7 per ton of carbon emissions in 2015.

The sustained low price for a permit to emit carbon is a common occurrence in the early years of cap-and-trade programs. After having the price for a carbon permit peaking at 30 Euros in 2008, the EU ETS experienced prices under ten Euros from

Figure III.4
Nominal RGGI Carbon Prices

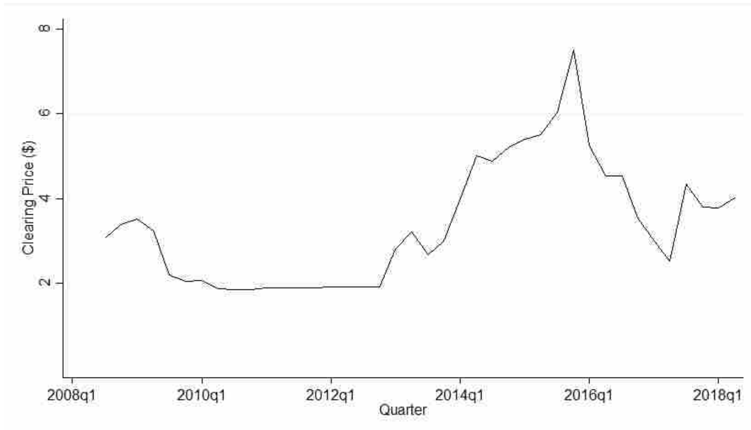

2012 until 2018.[7] California's cap-and-trade program has experienced a similar trend with carbon prices staying near $12 per ton emitted after starting around $20 per ton. As pointed out by Borenstein et al. (2015), theory models predict that cap-and-trade emissions auctions most often give rise to equilibrium prices determined by a price floor or a price ceiling due to the inelasticity of electricity demand curves.

III.2.4 The rise of natural gas; the decline of coal

A secondary reason for the switch away from coal and to natural gas is their substitutability and the rise of hydraulic fracturing. A power plant that runs on coal can be converted to a plant that runs on another type of fossil fuel with relative ease but at a high cost, and many power plants have both coal and natural gas generators. Most commonly, the generators in coal plants are converted to use natural gas as its main fuel due to the lower environmental impact of natural gas and to exploit the ability of natural gas to quickly adjust its power output. In some cases, this conversion can cost upwards of $200 million.[8] Near the end of the 2000s, this conversion became

[7]See Sandbag Climate Initiative (2019) and Climate Policy Initiative (2019) for full descriptions of the programs.
[8]The cost of a conversion varies greatly based plant size, type of conversion (i.e. totally natural gas or natural gas-coal hybrid), and the current equipment the plant has.

Figure III.5
Nominal Natural Gas Spot Price, Weekly

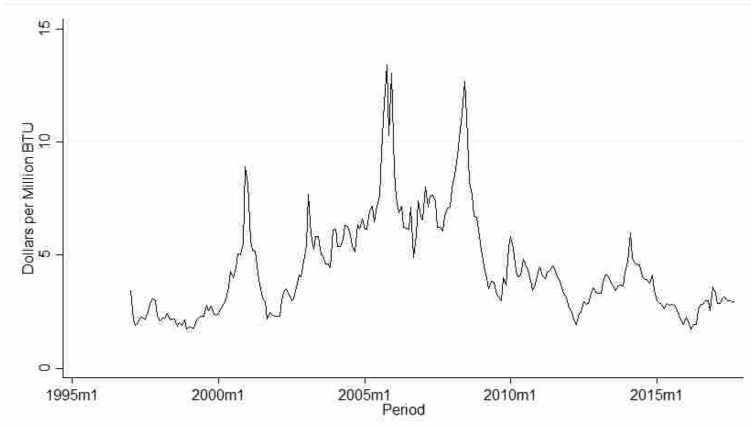

more economically viable as natural gas prices continued to fall and natural gas became more easily accessible due to the rise of fracking. This can be seen by the Henry Hub Natural Gas Spot Price time series in Figure III.5. Apart from some occasional noise, the natural gas spot price seemed to be steadily trending upward until its peak in 2008. After 2008, the spot price of natural gas drastically fell, and has stayed consistently very low since then.

The sharp drop in price almost immediately translated in a substitution from coal to natural gas, which can be seen in yearly production of coal and natural gas. In Figure III.6, I show the yearly quantities of coal and natural gas extracted over time. Both values stayed fairly steady until 2008, when the natural gas production began to steadily rise and the coal production fell dramatically.

III.3 Literature Review

Over the last 30 years, the energy market has evolved into a mix of regulated and unregulated utilities, independent power producers (IPPs) and investor-owned utilities (IOUs) that use an ever-changing portfolio of energy sources that vary due to changing economic conditions, climate change and environmental regulations. In this

Figure III.6
National Annual Coal and Natural Gas Production

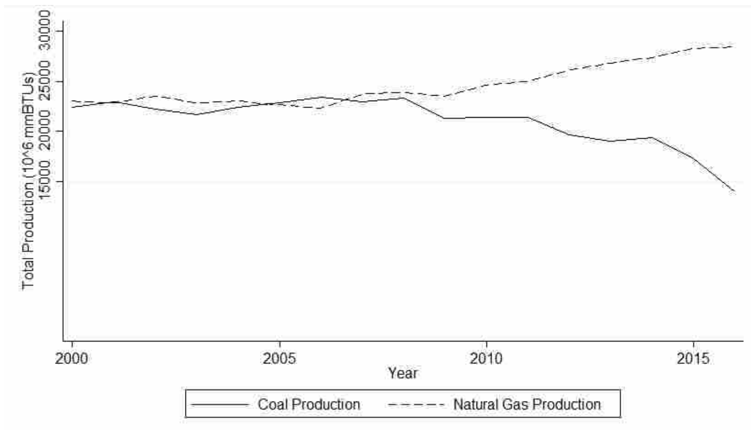

section, I will provide a review of the literature of electricity markets, and empirical studies of the effects of the fracking boom, RGGI and other carbon cap-and-trade programs on electricity markets.

Hirth et al. (2016) describe the heterogeneity of electricity markets and sources. While all electricity serves the same final purpose, it is incorrect to say that all electricity is the same for two main reasons. First, electricity is not a good that can easily be stored or transported across interconnections. Therefore, electricity generated today is sold in a different market than electricity generated in a week, and electricity cannot be transferred readily across large geographic distances. This creates a need for some energy product that can be quickly and easily dispatched locally. Second, the cost to produce electricity and the purpose of the electricity varies greatly between fuel sources, locations and even time of day. As pointed out by Covert et al. (2016), fossil fuels are much easier and less costly to dispatch than any renewable energy sources currently available even if they emit harmful pollutants. So, fossil fuel and renewable resources can broadly be considered two differentiated products in the energy market: one that has a negative environmental impact but is cheap and easy to adjust, and another that is clean but neither cost effective nor easily adjustable. The market can be further differentiated by types of renewable resources (i.e. wind, hydroelectric, solar, etc.) and types of non-renewable (nuclear, coal, natural gas, oil, etc.). Despite all the available technology, Covert et al. argue that fossil fuels will remain a dominant force into the foreseeable future.

Much research has been conducted to understand the environmental impacts of policies and different fuel sources. Jenner and Lamadrid (2013) provide a direct comparison of the environmental impacts of coal and natural gas extraction and electricity generation on air, land and water. They find that natural gas has a smaller negative impact on air quality, uses less water, and uses less land than coal, but natural gas may contribute to drinking water contamination and air pollution through leaky pipelines. Knittel et al. (2015) study the effects of power plants' decisions to switch between coal and natural gas in the wake of the shale gas boom for plants that have both coal and natural gas generators, and utilities that have both coal and gas power plants. They find that the sharp drop in spot price of natural gas in 2008 led to a 19% drop in emissions in traditional energy markets and a 33% drop in restructured energy markets. They also find that IOUs responded more extremely to the shale boom than IPPs. Jordan et al. (2018) study how costs and market forces affect a coal mine's profitability and closure probability, and find that the main reasons for coal mine closure are the rising cost of extracting Appalachian coal and the electricity sector's shift to natural gas use due to the fracking boom. Eyer and Wichman (2018) analyze how water levels affect an area's electricity generation portfolio, and find evidence that drought conditions lead to a substitution away from hydroelectric power and towards natural gas, but no evidence that drought conditions lead to higher coal use. Additionally, there is a large body of literature discussing how the emissions trading market and the Clean Air Act affected a power plant's decision to invest in pollution-abatement technologies (Popp (2003), Fowlie (2010), Bergek et al. (2014)).

Market-based carbon abatement policies are relatively new, however there exists some research analyzing the effects of carbon cap-and-trade programs on emissions and pollution-abatement incentives. Ellerman and Buchner (2007) analyzes firms' net short or long position on carbon permits in the EU-ETS to determine that although initially there may have been an overallocation of pollution credits in the EU-ETS, carbon abatement did occur. Using a matched diff-in-diff model that analyzes patents, Calel and Dechezlepretre (2016) shows that the EU-ETS induced firms to invest in pollution-abatement technology.

The Regional Greenhouse Gas Initiative has been the subject of intense research since its inception in 2009. Ruth et al. (2008) projected the effects of Maryland joining RGGI, and predicted that joining the RGGI would cause Maryland to sharply cut both its coal and natural gas generation, increase its electricity imports, and add no new coal or natural gas capacity to its grid. Murray and Maniloff (2015) analyze the causes of the reduction of carbon emissions in the RGGI region. They find that the RGGI program and the sharp drop in the natural gas price are both responsible for the reduction in carbon emissions in addition to other economic factors. In particular, they find that the RGGI program is responsible for half of the reduction in emissions in the RGGI region, which translated into a 24% drop in total carbon emissions. Kim and Kim (2016) use a synthetic control estimator to determine how much RGGI accelerated the trend of fossil fuel electricity generation switching from coal to gas. They find that the share of electricity generated by natural gas in the RGGI region is approximately 10-15% higher than the synthetic region. Using a differences-in-differences model, Fell and Maniloff (2018) estimate the change in plant-level generation decisions due to RGGI by plants in the policy region and plants in potential spillover states, and find that coal generation fell in RGGI states but natural gas generation rose in spillover states due to RGGI. None of these papers have attempted to discuss how a carbon tax affects seasonal changes in power plant-level generation decisions.

This research extends the literature by estimating how the roles of coal and natural gas electricity changed after RGGI. In particular, I estimate the changes in the seasonal fluctuations of coal and natural gas generation after RGGI. These changes provide evidence that a carbon tax shifts the role of natural gas towards a means of baseload generation while coal becomes more commonly used as a way to keep up with expected seasonal swings in electricity demand.

III.4 Conceptual Framework of Plant Behavior Under a Carbon Tax

Many studies have analyzed how an electricity market's supply curve changes due to price and policy changes. The framework presented here provides a simplified representation of these markets. For the ease of interpretation, the theory abstracts away from potentially relevant factors such as transmission constraints, start up time by generators, or strategic bidding by generators with the purpose of developing a set of empirically testable hypotheses on how the institution of a carbon tax can change the roles of carbon and natural gas. Figure III.7 presents a diagram to demonstrate how a tax with asymmetric effects on different fuel sources can change the dispatch order and quantity of electricity generated by each source. In effect, a carbon tax would be twice the size for a coal plant than a natural gas plant because coal emits approximately twice as much carbon than natural gas per unit of electricity produced as presented in Energy Information Administration (2019). Therefore, I model the carbon tax as a quantity tax that is twice as large for a coal generator than a natural gas generator, as is consistent with the emissions portfolio of typical coal and natural gas plants.

The demand curve, which represents the demand by consumers in a particular region during a particular season, is assumed to be incredibly inelastic if not perfectly

Figure III.7
Effects of a Carbon Tax on Dispatch Order

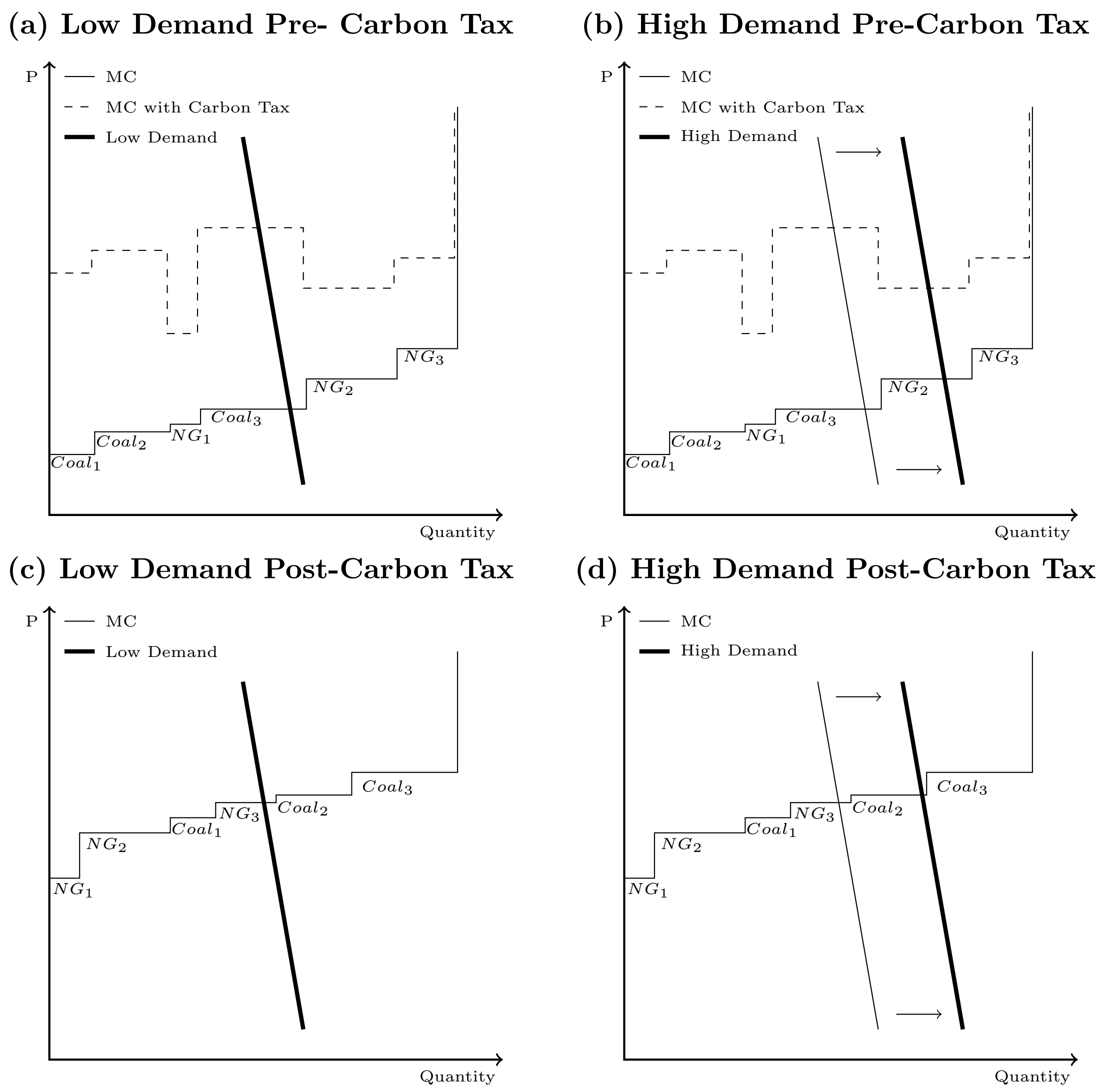

inelastic. The supply curve is block horizontal, where the height of each block represents the marginal cost of using the current-lowest cost generator to create electricity, and the width of each block represents the capacity of each generator. In panels (a) and (b) of Figure III.7, I present an electricity market with only three coal and three natural gas generators before a carbon tax is instituted. Each generator is labeled with its fuel source and a number to reflect the rank of its marginal cost relative to its fuel source.

Absent any carbon tax, coal is generally the lowest-marginal-cost method to generate easily-dispatched electricity, with natural gas being the second most. The spot price for natural gas can vary drastically and may give rise to temporary changes in the dispatch order. This is not to say that all coal generators are lower cost than all natural gas generators, but historically, the dispatch order is front-loaded with more coal than natural gas generators. In panels (a) and (b) of Figure III.7, I present a market with three coal generators and three natural gas generators, where the coal generators are generally lower-cost sources of electricity than the natural gas generators. This can be seen by the coal generators being the first-, second- and fourth-lowest cost generators out of the six generators in the dispatch order.

Low demand seasons are represented in panel (a) of Figure III.7. High-demand seasons will shift the demand curve to the right and require that more generators be used, and are represented in (b) of Figure III.7. Electricity demand generally peaks in the middles of summer and winter, and reaches a trough in the middles of spring and fall. In the low demand season presented here, two coal generators and one natural gas generators are running at full capacity. In the high demand season, a second natural gas generator is used.

According to the Energy Information Administration (2019), coal emits approximately twice as much carbon than natural gas per unit of electricity produced. Therefore, although a rise in the carbon tax would raise the marginal cost for both coal and natural gas generators, the rise in cost would be larger for coal generators than natural gas generators. This asymmetric effect is represented by the change from the solid line to the dashed line in panels (a) and (b) of Figure III.7, where the observed rise in marginal cost is twice as large for coal generators than natural gas generators.

After the institution of the carbon tax, the marginal cost curve is no longer monotonically increasing in panels (a) and (b), meaning that a utility would benefit by modifying the dispatch order of generators to bring the new lowest-marginal-cost generators to the front. Panels (c) and (d) depicts the utility's new marginal cost curve after modifying the dispatch order to allow a utility to use natural gas generators before coal generators. In the example provided, natural gas generators move up the dispatch order.

Some things should be apparent after instituting the carbon tax. First, more natural gas generators and fewer coal generators are being fully utilized in low demand periods. This can be seen in the shift from panel (a) to panel (c). This shows that under the model created, a carbon tax would lead to an increased use of natural gas generators as a form of baseload electricity generation. Second, regardless of the demand, less electricity is being generated by coal and more by natural gas, as is shown by the equilibrium demand of electricity in panels (c) and (d).

Third, the shift away from coal is much more dramatic in low electricity demand seasons than in high demand periods. This can be seen in the changes from panel (a) to (c) and panel (b) to (d). Only one coal generator is used in panel (c) while three are being used at or near full capacity in panel (a), indicating that coal is a much less attractive source of electricity in low demand seasons. However, in high demand periods, all three coal generators are used at full capacity in panel (b), and two coal generators are used near full capacity in panel (d). While the coal utilization did decrease in both high- and low-demand periods, the shift was smaller in magnitude in high demand periods than low demand periods. Therefore, a carbon tax would cause coal to be used more prominently as a means to generate more electricity in high demand seasons. In other words, the theoretical model predicts that a carbon tax causes coal's role to shift away from baseload electricity generation and shift towards supplementing seasonal changes in demand.

Finally, one would expect that natural gas use would rise across both periods of high and low demand. The mechanism is that a carbon tax moves natural gas generators up the dispatch order, making them more economically desirable across all seasons. This indicates that a carbon tax would cause the role of natural gas to shift towards baseload electricity generation.

These theoretical predictions lead to some testable hypotheses in my empirical estimation. First, the model predicts that a carbon tax would cause an increased utilization of natural gas and a decreased utilization of coal regardless of seasons. Second, the shift away from coal due to a carbon tax should be most pronounced in low demand periods for coal plants. Third, the shift towards natural gas should be positive across all months.

III.5 Data Sources and Variables

III.5.1 Data Sources and Preparation

All the data used in this study are publicly available. The main sources are the EIA Forms 906, 923 and 860 from 2001 to 2017. The 906 and 923 forms contain observations on all generation, fuel use, fuel characteristics and heat content at the plant-month-fuel-prime mover level. A plant that shuts down prior to December 2017 no longer appears in the data after exiting and a plant that opens after January 2001 appears in the data only after its opening date. As is common in the literature, I consider each plant-fuel source combination as a separate generator because I cannot separately observe how a plant chooses to utilize each individual generator. That is, I am unable to see how a plant chooses to dispatch its generators if it has multiple generators that use the same fuel, but I can observe the electricity generated by each fuel source at a plant. I aggregate these data to create an unbalanced panel of observations on the plant-month-fuel source level where I observe fuel use, fuel characteristics and net generation by source.

The 906 and 923 forms contain observations at a very disaggregated fuel level. For example, instead of listing that a power plant generated electricity using "Coal" or "Natural Gas" within a unit of observation, the raw data list that a power plant generated electricity using Bituminous Coal, Sub-Bituminous Coal, Lignite, Natural Gas, Blast Furnace Gas, etc. Using the EIA's descriptions of fuel types, I aggregate these up to the Coal, and Natural Gas Products level as defined by the EIA.[9] While the cap-and-trade program instituted by RGGI covers petroleum use, I omit all petroleum observations due to petroleum's very small role in the electricity industry.[10] At its peak during my sample, petroleum constituted only 5% of the total annual electricity generated in the US and 9% of the total electricity generated in the RGGI region. Petroleum use fell to 1% of electricity generated in both prior to the beginning of RGGI. My main data preparation and empirical analysis will be conducted using only two categories of fuel: coal and natural gas.

The EIA 860 form is an annual survey of power plants' equipment. Schedule 3 of this form lists each power plant's generators, the nameplate capacity of each generator, when the generator became operational, when it is scheduled to retire, and its fuel sources. A generator can list up to six potential fuel sources, where the first fuel source listed is the primary fuel source. When calculating a power plant's capacity for each fuel source, I use only the first fuel source listed, as is common in research that use these data.

Schedule 5 of the EIA 860 data contains information on pollution abatement equipment used by each power plant, the type of equipment, the date installation was complete, the cost of installation, and the date the equipment is scheduled to be retired. From this, I construct indicator variables for when a power plant has an operational scrubber, selective catalytic or non-catalytic reduction technology, and other pollution-abatement technology. These data are merged with the main panel created from the EIA 906 and 923 Forms at the plant-month level. So, observations

[9] Coal products include anthracite coal, bituminous coal, lignite coal, refined coal, recycled coal, coal-derived gas, subbituminous coal, and waste coal. Natural gas products include natural gas, blast furnace gas, and anything defined as "other gases" in the 923 forms.

[10] Petroleum products are generally only used to accommodate large daily or hourly fluctuations in electricity demand.

from a plant with coal and natural gas generators that installed a scrubber will all take on a value of 1, even though scrubbers are primarily used to clean coal emissions. I keep these controls in the natural gas analysis to control for a plant's possible preference away from natural gas if it has pollution abatement equipment that favors coal use. As an added control, I calculate the percentage of the total electricity generated by renewable resources in each state to control for any observable substitution away from fossil fuels towards renewable resources.

III.5.2 Summary Statistics

Table III.1 panel (a) contains a comparison of observable characteristics between power plants that use coal in RGGI states and non-RGGI states in the periods before RGGI went into place. Panel (b) of Table III.1 contains the same information as Panel (a) over the periods after RGGI went into place. I include both panels to demonstrate the change in power plant activity as well as the consistency of the control variables. Table III.2 contain information that is analogous to panels (a) and (b) of Table III.1, but for plants that generate electricity using natural gas instead of coal. A plant that is observed to use both coal and natural gas during a period will appear in both tables. I conduct a paired t-test to determine if observables are statistically significantly different from each other. I reject the null that observables are the same for all variables contained in Table III.1 and Table III.2. To account for these differences, I include all variables in Tables III.1-III.2 as controls in my differences-in-differences models.

Figure III.8 contains a plot of monthly means of the net generation by both coal and natural gas of plants inside the RGGI region and outside the RGGI region and the mean of capacity factors of operable plants that choose to operate within a month.[11]

[11]Capacity factor measures the utilization of a generator, and is defined as the amount of electricity generated in a time period divided by the total possible electricity generation in that time period. At the monthly level, capacity factor is given by: $\textit{Capacity Factor} =$

Energy consumption is highly seasonal and fossil fuel energies are the most easily dispatched forms of electricity, so the monthly swings are expected. Panel (a) of Figure III.8 shows a very stark drop in both the average net generation of active coal plants and the average capacity factor of all plants. It is also apparent that the drops are much more pronounced in spring and fall months. Additionally, it appears that after RGGI the average coal plant net generation did not change and the average capacity factor changed very little. In panel (b), it does not appear that RGGI had an easily observable effect on either measures of power plant use. In fact, it appears that average net generation in natural gas plants increased *more* in non-RGGI regions than in RGGI regions. This provides some evidence that a carbon tax is causing primarily a substitution away from coal, but not necessarily a substitution towards natural gas within RGGI states. This spillover is addressed in Fell and Maniloff (2018). As an

Table III.1
RGGI v. Non-RGGI Coal Plant Characteristics

(a) Pre-Policy | **(b) Post-Policy**

	Non-RGGI		RGGI		Non-RGGI		RGGI	
	p5/p95	mean	p5/p95	mean	p5/p95	mean	p5/p95	mean
Log(Net Generation)	7.11	11.07	7.38	10.98	6.51	10.97	6.55	10.24
	13.92		13.24		13.85		12.92	
Capacity Factor	0.07	0.55	0.11	0.56	0.04	0.48	0.02	0.30
	0.91		0.91		0.88		0.77	
State Renewable %	0.01	0.25	0.00	0.40	0.04	0.29	0.04	0.47
	0.52		0.61		0.57		0.69	
Log(Capacity)	2.14	5.30	2.01	5.18	2.01	5.45	2.13	5.34
	7.61		7.03		7.73		7.13	
Scrubber	0.00	0.25	0.00	0.12	0.00	0.45	0.00	0.39
	1.00		1.00		1.00		1.00	
SCR/SNR	0.00	0.12	0.00	0.27	0.00	0.31	0.00	0.51
	1.00		1.00		1.00		1.00	
Elec. Precipitator	0.00	0.57	0.00	0.59	0.00	0.58	0.00	0.68
	1.00		1.00		1.00		1.00	
Observations	59419		4799		46207		2693	

$\frac{Net\ Generation\ (in\ MWh)}{Capacity(in\ MW) \cdot 24 \cdot Days\ in\ Month}$

extension of my main model, I address how RGGI may affect the seasonal changes in electricity production in the "spillover" states defined by Fell and Maniloff (2018).

III.6 Empirical Modeling Approach

To determine how the roles of coal and natural gas changed after the institution of RGGI, I construct a generalized difference-in-difference model with observations at the power plant-month level. In the first subsection, I present my main empirical model and bring attention to the coefficients of interest and assumptions I employ in estimation. In the second subsection, I present an alternate model that I estimate as a robustness check.

Table III.2
RGGI v. Non-RGGI Gas Plant Characteristics

(a) Pre-Policy | **(b) Post-Policy**

	Non-RGGI		RGGI		Non-RGGI		RGGI	
	p5/p95	mean	p5/p95	mean	p5/p95	mean	p5/p95	mean
Log(Net Generation)	1.53	7.48	3.04	8.13	1.25	7.71	3.46	8.38
	12.33		12.35		12.70		12.66	
Capacity Factor	0.00	0.21	0.00	0.26	0.00	0.23	0.00	0.30
	0.84		0.83		0.83		0.86	
State Renewable %	0.02	0.28	0.02	0.42	0.07	0.32	0.03	0.47
	0.56		0.61		0.63		0.65	
Log(Capacity)	0.96	4.19	0.79	4.22	0.92	4.34	0.79	4.08
	7.01		6.94		7.09		6.88	
Scrubber	0.00	0.06	0.00	0.01	0.00	0.08	0.00	0.03
	1.00		0.00		1.00		0.00	
SCR/SNR	0.00	0.06	0.00	0.10	0.00	0.11	0.00	0.13
	1.00		1.00		1.00		1.00	
Electric Precipitator	0.00	0.12	0.00	0.10	0.00	0.13	0.00	0.10
	1.00		1.00		1.00		1.00	
Observations	153101		23129		159779		24416	

III.6.1 The Main Empirical Model

In my main empirical model, I construct fixed-effects differences-in-differences models to estimate the log levels of coal or natural gas used by a power plant in each month. The main econometric model is a linear fixed-effect regression that will be estimated separately for each month and fuel source. It takes the form:

$$y_{ist} = \alpha + \delta \cdot RGGI_{ist} + \rho \cdot Plant_i + \beta \cdot X_{ist} + T \cdot Period_t + \epsilon_{ist} \tag{III.1}$$

where i indexes the power plant, s denotes the state and t denotes the time period as measured in years. The dependent variable y_{ist} is the log of the net electricity generation at the plant-fuel source-month level. Only power plants with a positive net

Figure III.8
RGGI and Non-RGGI Monthly Trends By Source

(a) Coal **(b) Gas**

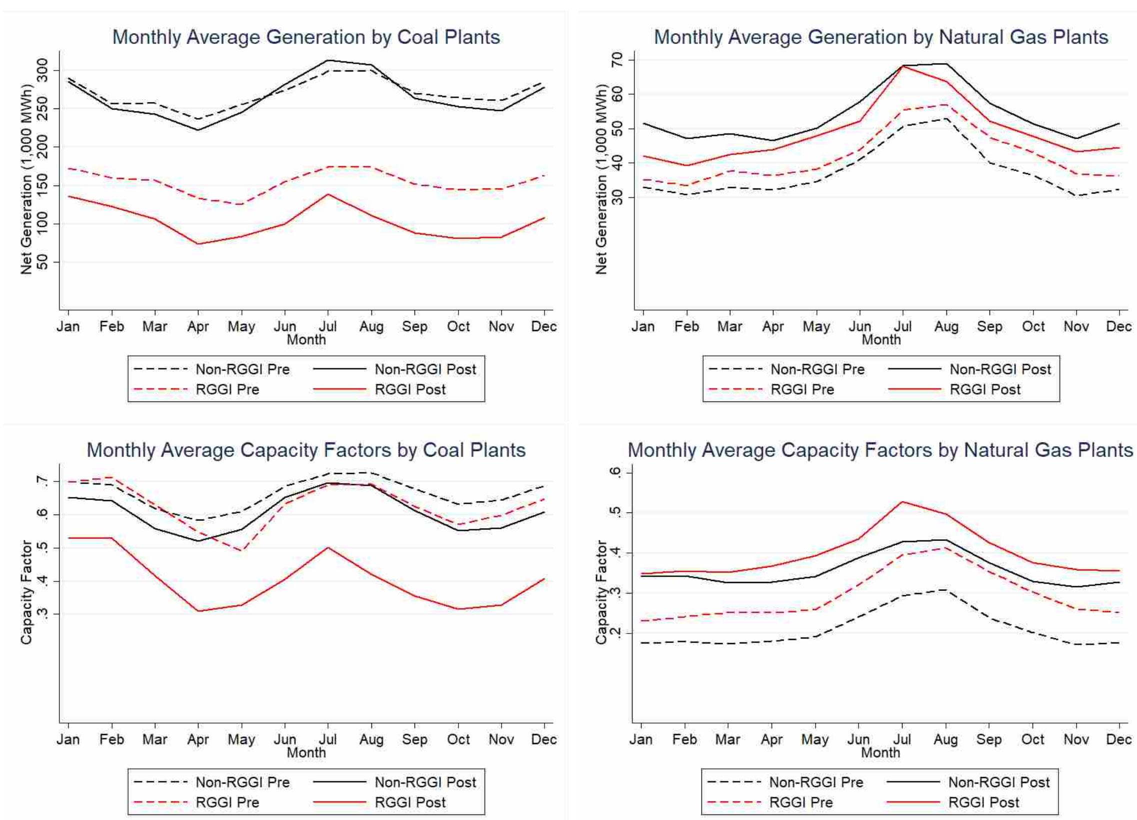

generation in a given month are included in the regression. The coefficient T is a linear time trend to capture any underlying national shift towards or away from each energy source. The vector X_{ist} controls for various state- and plant-level characteristics including the log of a plant's capacity, pollution-abatement equipment used by a plant, and the percentage of a state's electricity that is generated using renewable resources. Estimating a separate regression for each month allows the model to capture any potential heterogeneous effects of the control variables across months. For example, a plant's capacity may have more explanatory power in a high-demand month like January than in a low-demand month like April.

By inspection, it appears that the decision to use pollution-abatement equipment, particularly scrubbers in coal plants, may be affected by the institution of RGGI, which Angrist et al. (2013) points out can be a bad control. Appendix Figure A.III.1 demonstrates the apparent rise in scrubber adoption due to RGGI. I estimate the model with and without controls for pollution abatement equipment, but present only estimates of the model with pollution-abatement technology controls. In this case, including dummy variables for pollution-abatement equipment does not significantly alter any point estimates. All standard errors are clustered at the state level. I have also estimated the model with unclustered standard errors and standard errors clustered at the power plant level, and I find that state-level clusters give rise to the largest standard errors.

I abstract away from any strategic bidding for pollution credits and assume that RGGI affects all power plants symmetrically for two reasons. First, I cannot observe the bids submitted by each polluting entity or the credits that were ultimately awarded to each bidder. Second, the permit auctions end with all bidders receiving pollution credits at a common price with the option to resell credits in a secondary market. Therefore, each power plant should experience the same expected price for a pollution credit.

As a robustness checks, I estimate the same model where the dependent variable is replaced with a plant's capacity factor. Many of the results are still present under this alternative specification, and all are presented in later sections. I also estimate my main models without clustering, clustering at the state level, and clustering at the power plant level. Varying the levels of clustering only strengthens the statistical significance of the results presented in this paper, so only the state-level clusters are presented.

The variable $Plant_i$ is a plant-level fixed effect that is used to control for any unobservable differences between any plants. The variable $RGGI_{ist}$ is the treatment variable that takes on a value of 1 for any plant in an RGGI state after 2009 and a 0 otherwise, and the coefficient δ is the main coefficient of interest.[12] The coefficient δ will capture the change in the log of net generation by each power source by merely being in an RGGI state after RGGI went into effect.

The theoretical framework predicts that the coal use should decrease most extremely in low-demand periods, the fall and spring, and that natural gas use should increase relatively uniformly across months. So, one would expect that the estimate of δ would be negative and larger in magnitude in the winter and summer for coal plants, but positive and more consistent in magnitude for natural gas plants. In the results section, I present a plot of the estimates of δ across months for each fuel source. To summarize, in this model I estimate a separate regression and examine one coefficient of interest, δ, for each month, for a total of twelve regressions for each fuel source.

[12]New Jersey announced its intention to drop out of RGGI on November 29, 2011, with an effective date of January 1, 2012. I return all New Jersey plants to the untreated group in 2012. Plants in all other states are considered treated for the entire post-RGGI period.

III.6.2 Alternative Model

As a robustness check, I estimate a separate but similar model to characterize the seasonal swings in electricity generation at the power plant level. The alternative model is also a single fixed-effect differences-in-differences regression with month-level policy effects that will be estimated for each fuel source, but will instead be estimated as a single regression across months with an interaction between each month and the RGGI policy variable. It takes the form:

$$y_{ist} = \sum_{m=1}^{12} \alpha_m \cdot month_m + \sum_{m=1}^{12} \delta_m \cdot month_m \cdot RGGI_{ist} + \rho \cdot Plant_i + \beta \cdot X_{ist} + T \cdot Period_t + \epsilon_{ist} \tag{III.2}$$

In this alternate model, I pool all months into a single regression and estimate a separate effect of RGGI for each month. These estimates will be captured in the coefficients δ_m for $m = 1, \ldots, 12$. In place of a single constant term, I include a constant term for each month that I call α_m to capture the natural seasonal swings in power plant-level electricity generation. The final change in this alternate model is that I now measure time $Period_t$ at the month level instead of the year level. Apart from these changes, all other controls and predictions are identical to the main empirical model. This alternative model assumes that a power plant's response to any of the control variables does not inherently change across months. Unlike my main empirical specification, in this alternative model I estimate a single regression for each fuel source and examine twelve coefficients of interest, namely δ_m for $m = 1, \ldots, 12$.

III.6.3 Parallel Trends Assumption

In order to motivate the parallel trends assumption required to run a differences-in-differences model, I include a graph of the dependent variables aggregated up to the annual level for both the RGGI region and non-RGGI region for each fuel source in Figure III.9. As has been discussed in previous sections, fossil fuel electricity generations is highly seasonal so I include only the annual trends for the sake of readability.

In coal power plants, the average log of net generation and capacity factors appear to remain unchanged prior to the institution of RGGI for both plants in and out of the RGGI region. However, post-RGGI one can see clearly that the average log of net generation of coal plants drops significantly in the RGGI region after the institution of RGGI. Likewise, it is also apparent that the drop in the capacity factor is larger for RGGI plants than non-RGGI plants after the introduction of the policy. The trends do not appear as clear for natural gas plants. I address this through controlling for observables that were outlined in previous sections.

Figure III.9
RGGI and Non-RGGI Monthly Trends By Source

(a) Coal **(b) Gas**

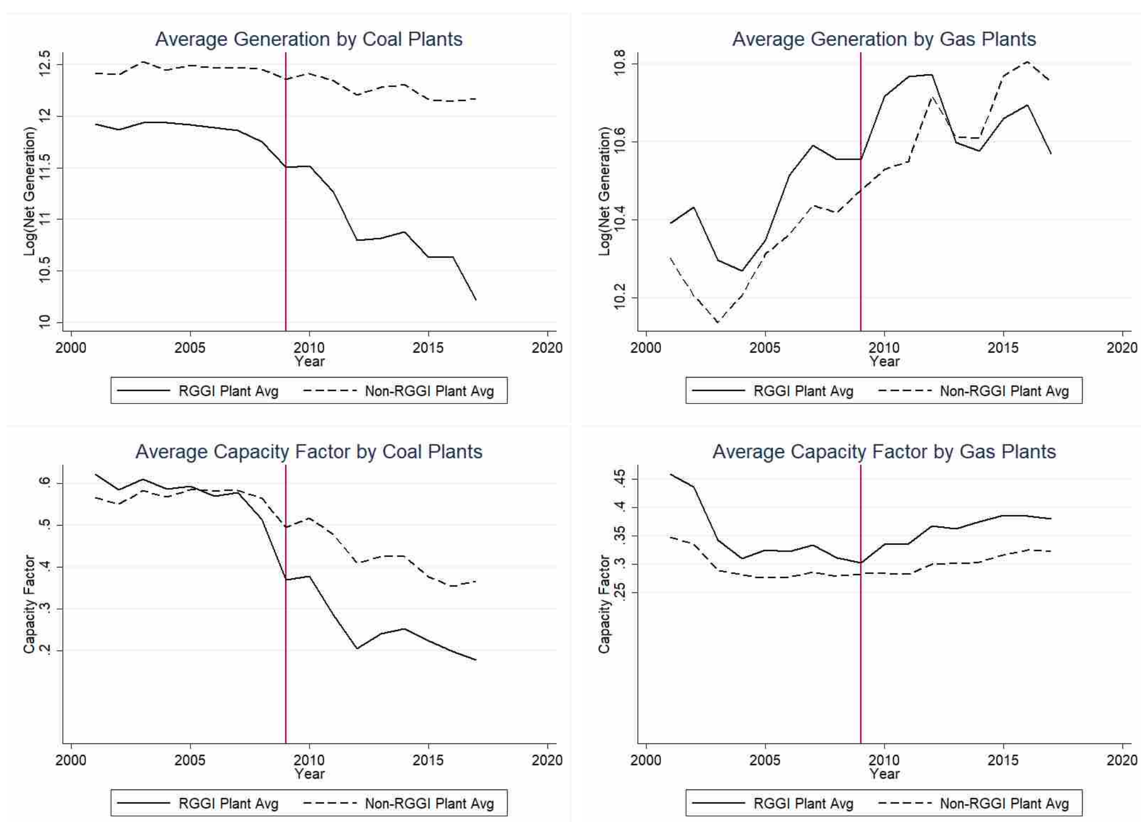

Figure III.10
Coal Monthly Difference-In-Differences

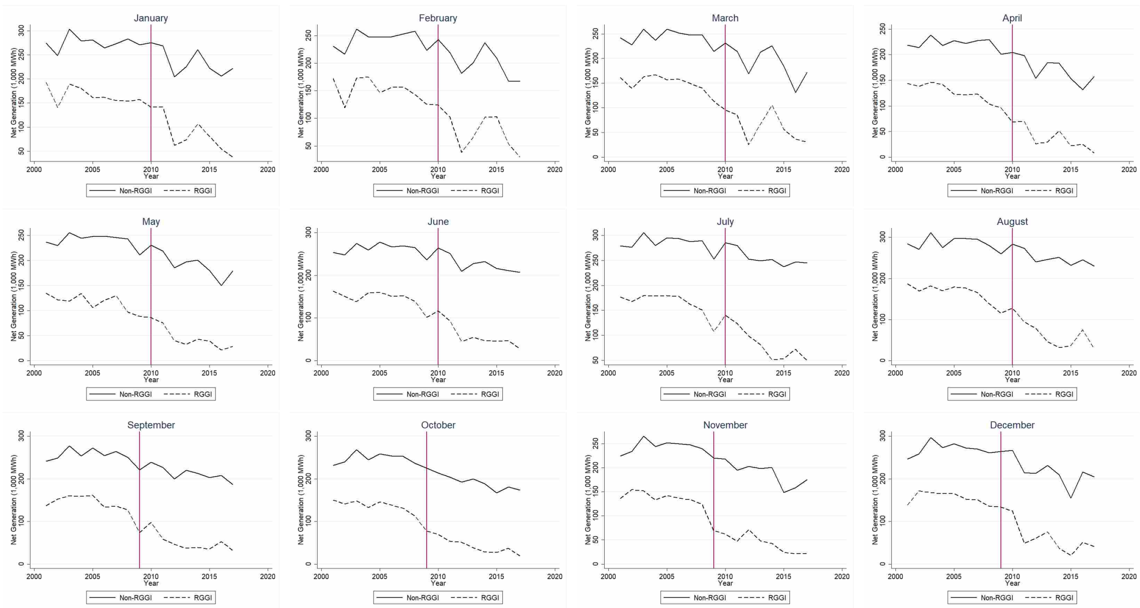

I not only discuss the shift away from coal and towards natural gas, but I also discuss how that shift varies seasonally. To demonstrate the seasonal heterogeneity, I present a separate difference-in-differences graph for each month for both coal and natural gas in Figures III.10 and III.11. In each graph, the vertical axis displays the monthly average net generation for power plants in the United States. The dashed lines denote the average of plants in the RGGI region while the solid lines denote the average of plants outside the RGGI region.

In Figure III.10, one can clearly see a few things. First, coal plants in the RGGI region produce electricity at a smaller scale than other coal plants in the United States. Second, across all months there is a clear drop in electricity production from coal. Finally, the largest visual drops in coal-powered electricity come in the spring and fall months.

Figure III.11 displays monthly difference-in-differences graphs for natural gas electricity production in the same manner as Figure III.10. Unlike the coal graphs in Figure III.10, there does not appear to be as clear of a transition towards natural

Figure III.11
Natural Gas Monthly Difference-In-Differences

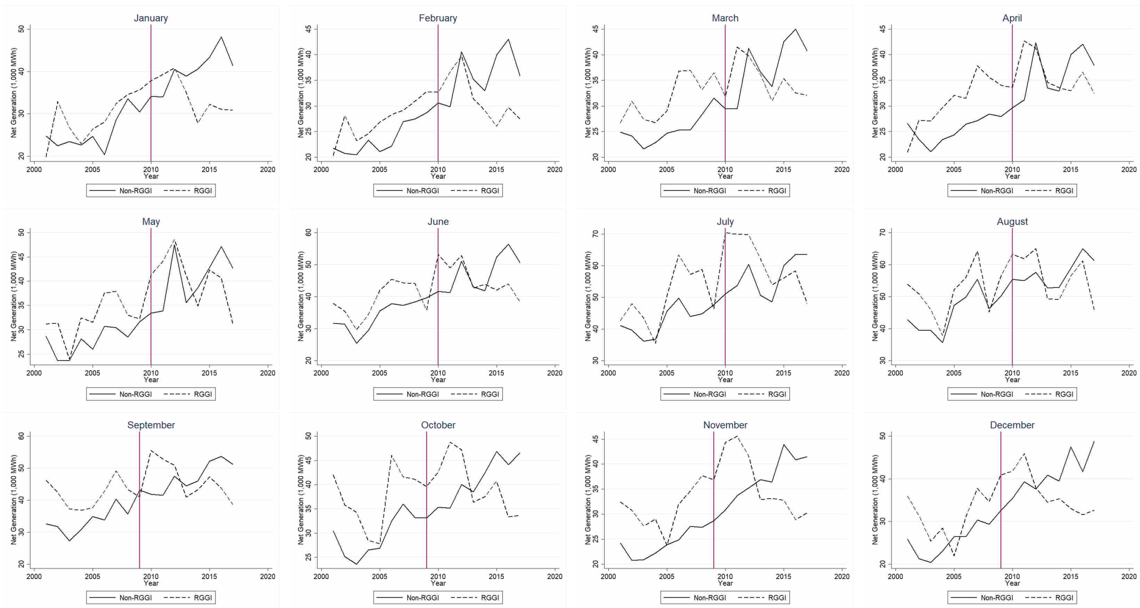

gas. However, it is clear that natural gas use is noisily increasing across all months. After controlling for confounding factors, my empirical results still show that there is a shift towards natural gas post-RGGI in the RGGI region.

III.7 Empirical Results

III.7.1 Main Empirical Results

Figure III.12 contains the estimated post-RGGI changes in the log of net generation at the plant level from equation (III.1) estimated separately for each month for coal and natural gas plants. Each point is a coefficient estimate from a separate model with 95% confidence interval bands drawn around it. Panel (a) contains the coefficient estimates from coal regressions and panel (b) contains the coefficient estimates from natural gas regressions. Tables A.III.1 and A.III.2 in the appendix contain all coefficient estimates of equation (III.1) for coal and natural gas models, respectively.

Some things should be immediately apparent. First, across all months the point estimates for the post-RGGI effect on coal electricity generation are uniformly negative and significant at the 95% level. Conversely, the point estimates for the post-RGGI effect on natural gas electricity generation are uniformly positive and all are significant at the 95% level except for December. This finding is consistent with the findings of Fell and Maniloff (2018).

Second, the post-RGGI change in coal power plant's electricity generation varies dramatically between months. The drop in electricity generation after the institution of RGGI in the winter months of January and February and the summer month of July have the smallest relative magnitudes, with estimated effects of -.37, -.49 and -.46 respectively. The largest post-RGGI effects appear to come in spring and fall, with estimated effects of April, September, October and November all having a magnitude larger than -.9. Furthermore, the January and February coefficient estimates of RGGI are statistically significantly smaller in magnitude than the estimates of the fall and spring months of April and November at the 95% level. These larger negative effects on coal use in low-demand months and smaller effects in high-demand months match

Figure III.12
Monthly RGGI Effects: Monthly Models

(a) Coal **(b) Natural Gas**

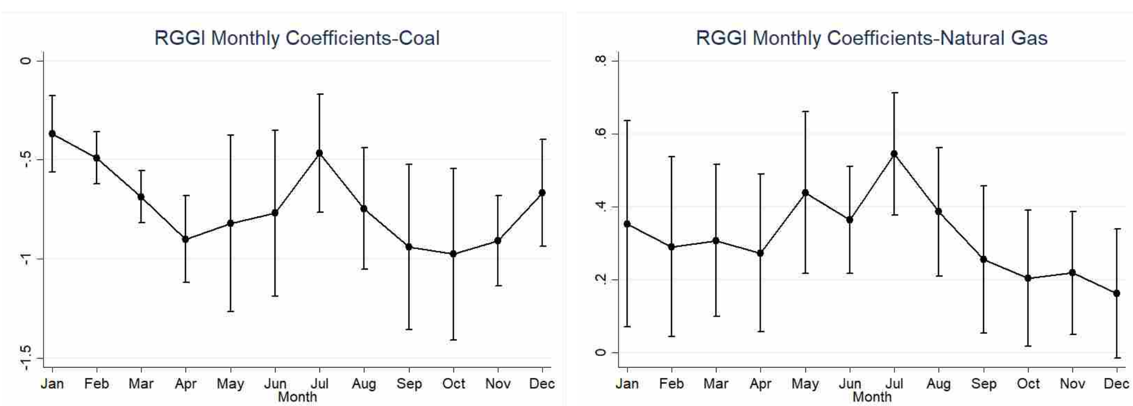

Note: This graph is generated by estimating equation (III.1) with the log of net generation as the dependent variable for each separate month and plotting the coefficient estimate for δ with the associated confidence interval. Standard errors are clustered at the state level. Unclustered errors and power-plant level clusters result in smaller standard error estimates.

up to the theoretical predictions outlined in section 4.

Third, the post-RGGI change in natural gas power plant's electricity generation is relatively steady across months. Only two months have significantly different changes in natural gas power plant behavior after RGGI, July and December. This can be explained by the base level of electricity generation in each month. Due to the seasonality of electricity demand, the estimated constant effect for December is larger than July, and both months are peak electricity-use months. So, the increased RGGI effect in July compared to December effectively normalized electricity generation between the two peak months. Additionally, one would expect that the effect of RGGI is largest for natural gas in the summer, because summer months experience the largest day-to-day swings in electricity demand and natural gas is especially easily dispatchable, and became relatively cheaper post-RGGI. These swings are mainly due to air conditioning expenses from exceptionally hot days.

III.7.2 Corroborating Empirical Results

As robustness check, I present the coefficient estimates from an alternative specification of my model in Figure III.13. Figure III.13 contains the estimated post-RGGI changes in the log of net generation at the plant level from equation (III.2) estimated separately for coal and natural gas plants. Each point in the graph is a coefficient estimate corresponding to the coefficient δ_m from equation (III.2) with its corresponding 95% confidence interval bands drawn around it. Panel (a) contains the coefficient estimates from coal regressions and panel (b) contains the coefficient estimates from natural gas regressions. Appendix Tables A.III.3 and A.III.4 contains coefficient estimates of the full equation (III.2) model for both coal and natural gas with various sets of controls. The coefficient estimates for each month change very little with each set of controls. All figures and tables presented in the rest of this section are created using the specification in column (4).

Many of the same trends are present in the estimation of equation (III.2) as equation (III.1). The most noticeable similarities are the distinct peaks in January and July, indicating that the change in generation due to a carbon tax is not as strong in the winter and summer as it is in the spring and fall. The effect of RGGI on January generation is statistically significantly smaller in magnitude than the effects on April, May, October and November, which again supports the theory outlined in section 2.4. As was the case in the model estimated in equation (III.1), all coefficient estimates in the natural gas model are positive and most are statistically significant, but there does not appear to be any differences between any pair of coefficient estimates in the natural gas model.

To further demonstrate the differences in seasonal effects of RGGI, I include Table III.3. In Table III.3, each entry (a, b) is the difference in the coefficient estimate of monthly RGGI effect in column a and the monthly effect in column b as estimated by equation (III.2) for coal plants. I test the hypothesis that differences of the monthly estimates of each pair is statistically different from zero, and bold every entry that

Figure III.13
Monthly RGGI Effects: Pooled Model

(a) Coal

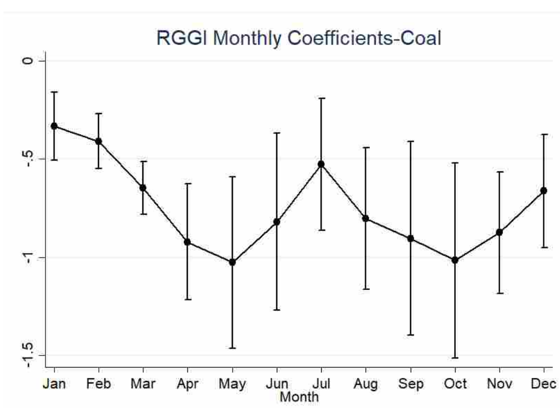

(b) Natural Gas

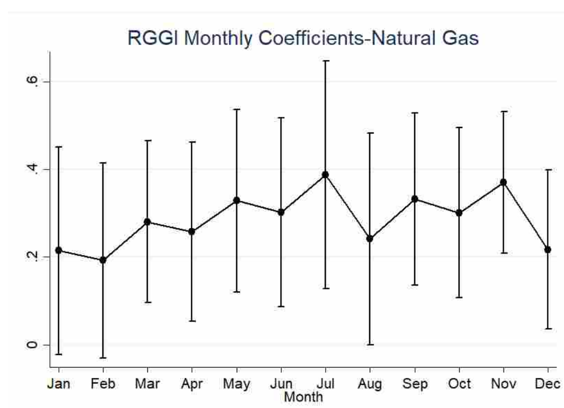

Note: This graph is generated by estimating equation (III.2) with the log of net generation as the dependent variable for both coal and natural gas and plotting the estimates of δ_m for $m = 1 \ldots 12$. Standard errors are clustered at the state level. Unclustered errors and power-plant level clusters result in smaller standard error estimates.

Table III.3
Difference in Monthly RGGI Effects-Coal Log Models

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Jan	0	.063	**.312*** **	**.579*** **	**.679*** **	**.476***	.166	**.453** **	**.55** **	**.665** **	**.534*** **	**.318*** **
Feb	-.063	0	**.249*** **	**.517*** **	**.616** **	.413	.103	**.39** **	**.487***	**.602** **	**.471*** **	**.255***
Mar	**-.312*** **	**-.249*** **	0	**.268** **	.367	.164	-.146	.141	.238	.353	**.222***	.006
Apr	**-.579*** **	**-.517*** **	**-.268** **	0	.099	-.104	**-.414*** **	-.126	-.029	.085	-.045	**-.261** **
May	**-.679*** **	**-.616** **	-.367	-.099	0	**-.203** **	**-.513*** **	**-.226** **	-.129	-.014	-.145	-.361
Jun	**-.476***	-.413	-.164	.104	**.203** **	0	**-.31*** **	-.023	.074	**.189*** **	.058	-.158
Jul	-.166	-.103	.146	**.414*** **	**.513*** **	**.31*** **	0	**.287*** **	**.385*** **	**.499*** **	**.369*** **	.153
Aug	**-.453** **	**-.39** **	-.141	.126	**.226** **	.023	**-.287*** **	0	.097	**.212** **	.081	-.135
Sep	**-.55** **	**-.487***	-.238	.029	.129	-.074	**-.385*** **	-.097	0	**.114*** **	-.016	-.232
Oct	**-.665** **	**-.602** **	-.353	-.085	.014	**-.189*** **	**-.499*** **	**-.212** **	**-.114*** **	0	-.13	-.346
Nov	**-.534*** **	**-.471*** **	**-.222***	.045	.145	-.058	**-.369*** **	-.081	.016	.13	0	**-.216*** **
Dec	**-.318*** **	**-.255***	-.006	**.261** **	.361	.158	-.153	.135	.232	.346	**.216*** **	0

The entry i, j represents $\delta_i - \delta_j$ in column (4) of the main results table
*** $p < .01$, ** $p < .05$, * $p < .1$

has a p-value of at most .1.

The same trends can be seen in this tables as in Figures III.12 and III.13. First, it is apparent that the the shift away from coal in January is the smallest of all months, and coal plants experienced a smaller drop in electricity generation in winter months than in spring and fall months. One can also see that coal-powered electricity experiences a larger drop in fall and spring months than in the winter and summer months. These changes are statistically significant.

In Table III.4, I present the differences between monthly coefficient estimates for natural gas plants. The arrangement of coefficients is identical to Table III.3. While there are some months in which natural gas generation varies significantly relative to other months, that the shift towards natural gas largely does not vary between months.

III.7.3 Further Robustness Checks

As a further robustness check, I run the models from equations (III.1) and (III.2) and replace the dependent variable with a plant's capacity factor in place of the log of net generation. Figure III.14 contains a plot of the regression coefficients from equation (III.1) and Figure III.14 contains the regression coefficients from equation (III.2). While the results are not as clear as the main models from the previous subsections, many of the same results are still apparent. In both models, a coal power plant's capacity factor decreased significantly after RGGI for all months, while a natural gas power plant's capacity factor rose statistically significantly after the institution of RGGI in summer months and had positive point estimates for most months. Additionally, the magnitude of the effect on coal plants remain the smallest in high-electricity-demand winter months.

Table III.4
Difference in Monthly RGGI Effects-Natural Gas Log Models

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Jan	0	.031	-.07	-.037	-.097	-.072	-.147	-.006	-.141	-.124	**-.2****	-.044
Feb	-.031	0	**-.101***	-.068	-.128	-.103	-.178	-.037	-.172	-.155	**-.231*****	-.075
Mar	.07	**.101***	0	.032	-.027	-.003	-.077	.064	-.071	-.054	**-.13****	.025
Apr	.037	.068	-.032	0	-.059	-.035	-.11	.031	**-.104***	-.087	**-.163***	-.007
May	.097	.128	.027	.059	0	.024	-.05	**.091***	-.044	-.027	-.103	.052
Jun	.072	.103	.003	.035	-.024	0	**-.075****	**.066****	-.069	-.051	-.128	.028
Jul	.147	.178	.077	.11	.05	**.075****	0	**.141*****	.006	.023	-.053	.103
Aug	.006	.037	-.064	-.031	**-.091***	**-.066****	**-.141*****	0	**-.135****	**-.118***	**-.194***	-.038
Sep	.141	.172	.071	**.104***	.044	.069	-.006	**.135****	0	.017	-.059	.097
Oct	.124	.155	.054	.087	.027	.051	-.023	**.118***	-.017	0	-.076	.079
Nov	**.2****	**.231*****	**.13****	**.163***	.103	.128	.053	**.194***	.059	.076	0	**.156****
Dec	.044	.075	-.025	.007	-.052	-.028	-.103	.038	-.097	-.079	**-.156****	0

The entry i, j represents $\delta_i - \delta_j$ in column (4) of the main results table
*** $p < .01$, ** $p < .05$, * $p < .1$

Figure III.14
Robustness Check: Monthly Models

(a) Coal

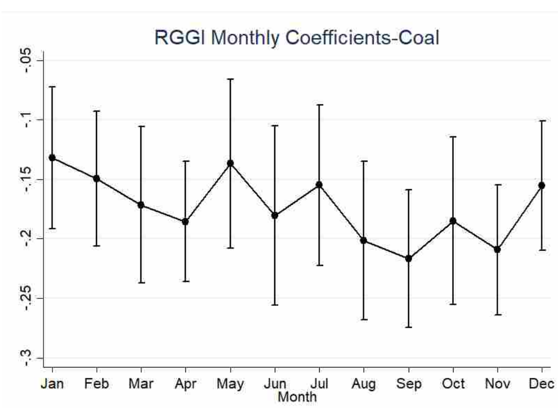

(b) Natural Gas

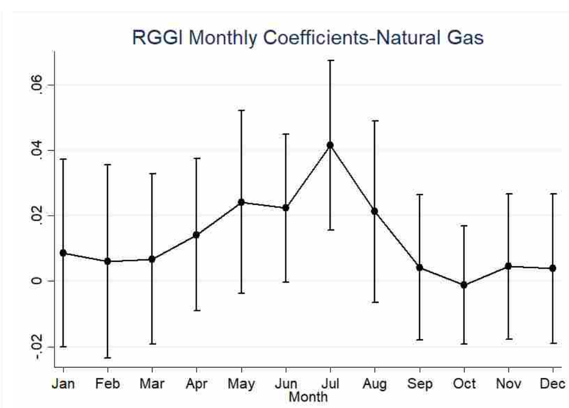

Note: This graph is generated by estimating equation (III.1) with capacity factor as the dependent variable for each separate month and plotting the coefficient estimate for δ with the associated confidence interval. Standard errors are clustered at the state level. Unclustered errors and power-plant level clusters result in smaller standard error estimates.

III.7.4 Spillover Effects

A common concern of any subnational environmental policies is the presence of any potential spillover effects due to the pollution haven hypothesis. The pollution haven hypothesis states that any environmental policy that does not affect an entire region uniformly will cause polluting industries to migrate to regions where the policy is not in effect. This is a large concern in electricity markets where energy can be easily transmitted across interconnections.[13]

To address the issue of any seasonal change in spillovers, I employ a similar methodology to Fell and Maniloff (2018). In their research, they note that the two most likely locations for any spillover effects are Ohio and Pennsylvania because they are both in the same interconnection as many RGGI states, are geographically very close to the regions covered by RGGI and rely heavily on fossil fuels. I omit all RGGI states from my analysis and re-estimate equations (III.1) and (III.2) with a

Figure III.15
Robustness Check: Pooled Model

(a) Coal

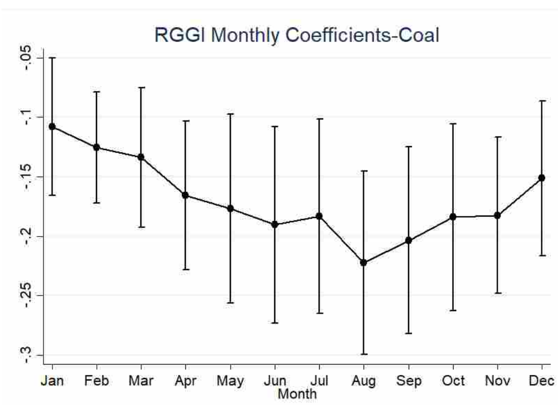

(b) Natural Gas

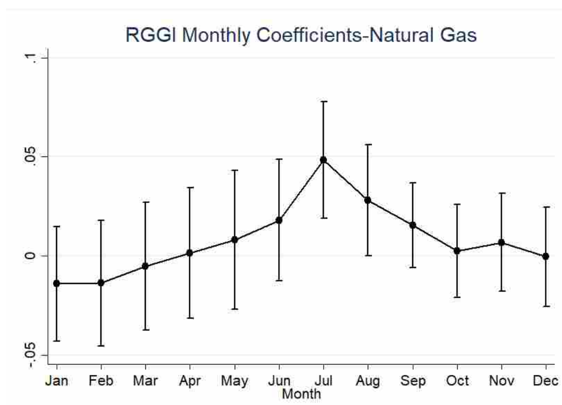

Note: This graph is generated by estimating equation (III.2) with capacity factor as the dependent variable for both coal and natural gas and plotting the estimates of δ_m for $m = 1 \dots 12$. Standard errors are clustered at the state level. Unclustered errors and power-plant level clusters result in smaller standard error estimates.

[13]For a full discussion of the pollution haven hypothesis, see Harrison and Eskeland (1997) and Cole (2004).

post-policy binary variable given to plants in Ohio and Pennsylvania instead of power plants in RGGI regions. The figures below contain the results of my estimation run with Log(Net Generation) as the dependent variable.

In general, the results of my estimation are consistent with those of Fell and Maniloff (2018). Namely, there does not appear to be any consistent positive effect on coal electricity generation in spillover effects in Figures III.14 and III.15, and there is a statistically significant *negative* effect in September and October. This provides evidence against the pollution haven hypothesis for electricity generated by coal.

Across both specifications, the spillover effects from RGGI are positive and statistically significant at the 95% level, indicating that some electricity generation is relocating to natural gas plants in non-RGGI states. However, there does not appear to be any noticeable change in the seasonality of plant use in the spillover states. This provides evidence that any seasonal changes in electricity generation due to a carbon tax are borne by plants directly affected by the policy. Therefore, the roles of

Figure III.16
Monthly RGGI Spillover Effects: Monthly Models

(a) Coal **(b) Natural Gas**

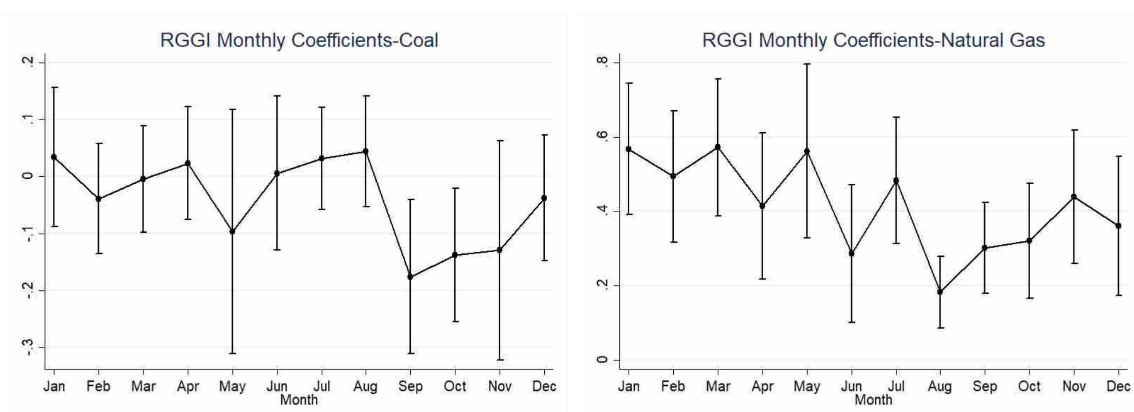

Note: This graph is generated by estimating equation (III.1) with the log of net generation as the dependent variable for each separate month and plotting the coefficient estimate for δ with the associated confidence interval. In these spillover models, I omit all RGGI states and estimate δ for only potential spillover states. Spillover states are assumed to be Ohio and Pennsylvania. Standard errors are clustered at the state level. Unclustered errors and power-plant level clusters result in smaller standard error estimates.

Figure III.17
Monthly RGGI Spillover Effects: Pooled Model

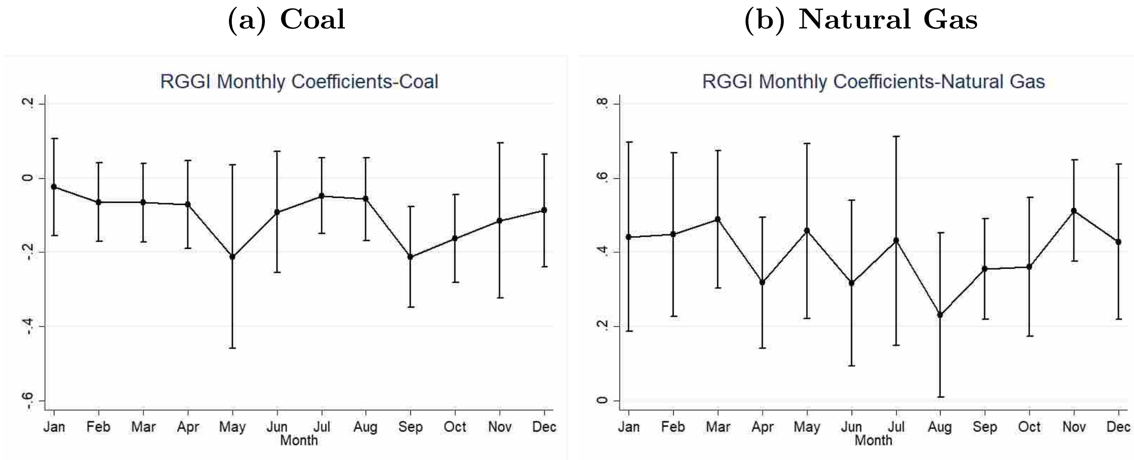

Note: This graph is generated by estimating equation (III.2) with the log of net generation as the dependent variable for both coal and natural gas and plotting the estimates of δ_m for $m = 1 \ldots 12$. In these spillover models, I omit all RGGI states and estimate δ_m for only potential spillover states. Spillover states are assumed to be Ohio and Pennsylvania. Standard errors are clustered at the state level. Unclustered errors and power-plant level clusters result in smaller standard error estimates.

natural gas and coal plants shifted only for plants in RGGI region, but not in regions that likely export electricity to the RGGI region.

III.8 Conclusion

In this chapter, I investigate how a carbon tax changes the role of coal and natural gas power plants in the dispatch order. To do this, I use the relatively new Regional Greenhouse Gas Initiative (RGGI) in the northeastern United States. Previous research has shown that after the institution of RGGI, carbon emissions decreased, coal use decreased dramatically while natural gas use rose. I build off of the research of Fell and Maniloff (2018), Murray and Maniloff (2015), and Kim and Kim (2016) to show that the roles of coal and natural gas changed after the carbon tax. Unlike these studies, I focus on the effects of RGGI on the seasonal changes in fossil fuel electricity generation. I find that a carbon tax leads to a consistent increase in natural gas electricity generation across seasons, a small decrease in coal electricity generation in high demand months, and much larger decrease in coal electricity generation in low demand months. This suggests that a carbon tax encourages the use of natural gas plants as a form of baseload electricity generation while relegating coal plants to acting as a supplemental form of electricity.

I begin by building up a theoretical model of dispatch order with both coal and natural gas plants. Because coal emits approximately twice as much carbon as natural gas, a quantity carbon tax would be approximately twice as costly for coal plants than natural gas plants. The theoretical model predicts that this would cause natural gas plants to move up the dispatch order and cause their utilization to increase consistently across high and low demand months. Conversely, this would cause coal plants to move back the dispatch order, causing large decreases in coal use in low electricity periods and relatively smaller decreases in high electricity demand periods.

To test this change, I investigate how power plant-level coal and natural gas use changes post-RGGI in RGGI states. I do this by constructing various difference-in-difference models that analyze coal and natural gas use in various months. High demand months generally occur in the summer and winter and low demand months are in the fall and spring. I find that natural gas use rises in a relatively consistent manner across months, while coal use decreases across all months and has the largest decrease in winter and summer months. This provides support for the theoretical conclusions, and the results hold across various specifications. As a final test, I check to see if there is a similar change in the role of coal and natural gas plants in states near to RGGI but not affected by RGGI. I do not find a similar change in seasonality in these states, providing evidence that the change in coal's and natural gas's role is only in areas directly affected by carbon taxes.

All these findings provide evidence that coal is less useful than natural gas in electricity generation when a carbon tax is in place. Both forms of energy are easily dispatchable and becomes much more expensive post tax. As carbon taxes become more commonplace, my findings show that electricity authorities and energy managers should invest more in natural gas plants and divest from coal plants.

CHAPTER IV

COAL'S DEMISE AND SHIPPING PATTERNS

IV.1 Introduction

The landscape of electricity generation in the United States is changing rapidly. Within the last twenty years, natural gas has displaced coal as the dominant source of electricity in the United States, a combination of policy and technological innovation has led to the development of renewable resources, and new long-term goals to reduce emissions promise to further diminish the role of all fossil fuels in the United States. Much work has been done to analyze these changes within mining and electricity markets (Yin and Powers (2010); Knittel et al. (2016); Holladay and LaRiviere (2017); Jordan et al. (2018); Linn and McCormack (2019)), but there are still many markets affected by this change that have been unexplored.

A topic that is often overlooked in the fossil fuel conversation is how sources of energy travel from their initial extraction points to the power plants where they are converted to usable energy. It is well established that natural gas is transported primarily via pipelines and coal is transported via rail and barge, but the scope of these movements has changed drastically as power plants substitute away from coal and instead to natural gas. One would expect that rail firms would change their pricing and operating strategies in response to the sudden substitution away from coal after the rise of fracking and institution of new environmental problems. To the best of my knowledge little work has been done to analyze how railroads have changed their shipping and pricing behavior in response to the collapse of coal. While the rise of natural gas and new environmental policies have led to an overall leftward shift of the demand for coal, this response likely is not uniform across coal basins. In this chapter, I investigate how have rail rates and quantities of coal shipments have changed as a result of environmental policy and the rise of natural gas. The results of this research particularly contribute to the literature on the rapidly changing fossil fuel industry as well as the broader literatures of transportation economics, industrial organization and environmental economics. I apply these results to discuss how the heterogeneous declines of coal mining in Appalachia and the Powder River Basin have affected the rail industry.

Coal and transportation via railroad have traditionally been close complements, with around 70% of coal in the United States being shipped in part by railroad and approximately 40% of all rail traffic in the US coming from coal (EIA (2019)). The inter-dependency of the these two industries makes collapse of the coal is a topic of interest for rail firms, electricity utilities and industry regulators. During the middle of the twentieth century, the railroad market was plagued by overregulation that led to the collapse of many major railroads and profits that were consistently negative. Through consolidations and reclassification, the number of Class-1 railroads fell from 40 firms to just seven. After the partial deregulation of railroads due to the Staggers Rail Act of 1980, the railroad industry came to be dominated by four major Class-I carriers that are able to charge lower rates than they were pre-deregulation. This was due to regulators allowing rail companies to shed costly, unnecessary equipment and high-cost or low-traffic routes.

Despite saving the railroad industry from the brink of collapse and improving the remaining firms' financial stability, the Staggers Act still imposes regulations on railroads that some experts argue are too strict. Industry insiders claim these regulations still leave railroads revenue inadequate (Winston (2005)) and economists researching the industry claim that the current railroad costing model doesn't properly address issues of competition among railroad firms (Wilson and Wolak (2018)). Wilson and Wolak (2018) create a new costing model to better address competition between firms and more accurately measure shipping costs. Due to this, the collapse of the coal industry- one of the primary commodities shipped by rail- could be catastrophic to railroad firms.

The electricity market's shift towards natural gas has had a major negative impact on coal markets, and these effects have been spilled over into rail markets. However, the effects are heterogeneous across coal basins. The US has two major coal basins, Appalachia and the Powder River Basin, and two minor coal basins, the Uinta and Illinois basin, that vary in coal quality, market composition and rail coverage. Rail firms have different pricing strategies in each of these basins that will evolve as coal production in these basins contract in different ways. Quantifying the effects of the shift towards renewable energies and natural gas can help gauge the magnitude of railroad firms' lost profits and any possible heterogeneous response in shipping patterns between coal basins. The results of this research can be used to inform policy-making decisions on the optimal level of railroad regulation.

The results of this chapter's research also contribute to the literature of firms' dynamic pricing strategies in declining markets. In the coal market, many contracts between mining companies and power plants can run for many years, while shipping decisions are made with a shorter time horizon with confidential contracts (Energy Information Administration (1991); Moyer (1965)). Therefore, a rail firm will be able to dynamically respond to a change in coal markets without risking a substitution away from rail.

In this chapter, I continue to explore the ever-changing fossil fuel electricity market, with an extension to the effects on rail markets. I use confidential waybill data provided by the Surface Transportation Bureau (STB) and combine them with mine and power plant data from the Mine Safety and Health Administration (MSHA) and the Energy Information Administration (EIA) to create a model of railroad pricing and quantity decisions that are a function of both coal demand near the destination of the shipment and coal supply near the origin of the shipment. It is entirely intuitive to expect that the quantity of goods sold in the market shrinks as the market declines. However, the pricing strategies are not as clear. In response to a shrinking market, it may be the case that rail firms choose to lower their prices in order to keep coal buyers and sellers in the market as long as possible. However, it may also be the case that a rail firm recognizes the impending sunset to the market, and chooses instead to raise the price to extract as much profit from captive agents in the market before they can make any future participation decision.

I find that the rate that a railroad charges rises as the demand market size for coal around the destination of the rail shipment falls, indicating that rail companies charge more to ship coal when the total coal market demand is small. This indicates that rail firms view a decline in demand as having immediate consequences and try to "shake out" any remaining profits from power plants before the industry declines any further. A similar relationship exists between the size of the mining market near the origin of the shipment and the price of the coal shipment. Interestingly, rail firms appear to charge lower rates as the number of individual mines near the origin or plants near the terminus of a shipment falls, rail firms take into account both the market size and the total number of agents in the market.

To account for heterogeneous market compositions and coal products, I run my main model separately for coal shipments originating from each of the four US coal basins. This exposes some previously unseen patterns across basins. In the Powder River Basin, I find that coal transportation rates respond only to the size of the local coal mining market changing while rail rates in Appalachia also rise as the number of individual mines rises.

The rest of this paper is organized as follows. In section 3.2, I provide a brief historical background of the rail and coal industries and discuss their current relationships to each other. In section 3.3, I provide a review of the literature on fossil fuel electricity, the collapse of coal, railroad pricing models. In section 3.4, I detail the data I use for my empirical analysis and how I address shortcomings of the primary data sets. I discuss the empirical techniques I use to estimate the effects of the fall of coal on railroad pricing and the size of railroad shipments in section 3.5 and provide summary statistics in section 3.6. I present the results of my preliminary empirical estimation in section 3.7 and provide an analysis of them. In section 3.8, I conclude my research.

IV.2 Railroad and Coal Background

Since the industrial revolution, railroads have been a key part to the economy of the United States. Markets that were previously too far away for large-scale trade of input goods and final goods could now sustain trade with each other and utilize complementarities that were once infeasible (Bain (2000); Dilts (1996); Moody (1921)). Over the next century one of the largest benefactors of these newly-integrated markets was the coal industry, with coal now being able to be shipped from remote mines to larger cities with power plants and factories that utilize coal. However in the last half century, coal markets have been the recipient of environmental regulation and competition from natural gas and renewable energy sources, and railroad companies dealt with ever-changing regulation that led many firms to bankruptcy and led to large-scale industry concentration. In this section, I provide a brief history of the important agents and regulatory environment of the modern US railroad market, the current state of the United States' coal market, and these two markets' interdependence.

IV.2.1 US Rail Markets

Throughout the middle part of the 20th century, rail firms in the United States were highly regulated in order to ensure the safety of rail users and to ensure that competition exists in various submarkets after it was observed that railroads of all sizes acted in an oligopolistic manner (National Archives (2020)).[1] Shipments often use multiple railroads or modes of transportation in their journey from the origin to the final destination, and for many years rail roads were the only effective way to ship large quantities of cargo across the country. As a result, railroads of all sizes were profitable. According to the STB waybill sample, around 14.5% of all sampled coal shipments from 2001-2016 used at least two railroads to transport coal.

As the 20th century progressed, highways and airlines became more widely used in the US for long-distance transportation, which led to a fall in railroad profitability for firms of all size. In addition to the rise of these competing transportation options, a commonly-cited factor for the fall in rail firms' profitability was overregulation MacDonald (1989). This led to many class-1 railroads abandoning or selling their tracks to short-line firms in order to stay afloat Fischer et al. (2001). The United States government believed this to be a large problem and responded in the late 1970s and early 1980s by passing new acts that reduced railroad oversight without completely eliminating it, starting with the Regional Rail Reorganization Act of 1973 (the 3-R Act) and the Railroad Revitalization and Regulatory Reform Act (the 4-R Act) in 1976. The most recent legislation was the Staggers Act, which has served as the primary regulatory document for railroads since October 14, 1980.

In general, the Staggers Act decreased regulation in three main areas: railroad pricing practices, route coverage and the type of equipment that a train must carry

[1] The Interstate Commerce Commission (ICC) classifies railroad firms by their sizes and total revenues, with the national railroad firms typically being designated as Class 1 railroads, regional railroads being designated as Class 2 and the smallest local railroads being classified as Class 3. Class 3 railroads are often also called short-line railroads.

on its journey. Prior to the Staggers Act, railroads were forced to set prices using a system established by the ICC that was criticized for being overly complex and restrictive and often forced railroads to set unprofitable transportation rates Stover (2008); Boyer (1981). After the institution of the Staggers Act, a rail firm was allowed greater freedom to set a rate that they chose and only received regulatory oversight if it was deemed that there was no effective competition on a certain route. In practice, the ICC rarely has to regulate rates.[2]

While deregulation should substantially increase the rates charged by rail carriers due to increased pricing power, rail rates declined for actually declined for many years. This was due to the Staggers Act incentivizing more multiple-car shipments rather than single car shipments (MacDonald (1989); Boyer (1981)). Additionally, railroads were given more freedom to choose optimal routes, effectively shedding more unnecessary costs. The total effect of this deregulation can be seen in a 2019 study by the Association of American Railroads. They show that over a decade, rail firms completely reversed their losses in total rail traffic while also eliminating costs and halving the average shipping rate (Association of American Railroads (2019)).

To get to this point, US rail firms went through a period of mergers, buyouts and closings that led to drops in rail coverage and the number of operating firms. As it stands now, the United States is serviced by seven Class 1 railroads and a host of shortline railroads that often partner with class-1 firms on longer shipments. To get an idea of where these railroads operate, I plot the current rail coverage in the continental United States by Class 1 firms in the left panel of Figure IV.1. One can clearly see that the east coast, Appalachian and midwest regions have many railroads whose lines overlap while many areas west of the Rocky Mountains only have a single Class 1 rail firm servicing a large area. This leads to far more potential pricing power

[2]The ICC currently mandates that a railroad will not receive any regulatory oversight for any rate that is less than 180% of the variable cost required to carry the shipment. Any rate that is observed to exceed 180% of the variable cost is handled on a case-by-case basis.

in the western United States. In the right panel, I plot all operational mines in 2016 in the United States scaled up by their annual production. It can be seen that Wyoming has only a few very larger mines while the Appalachian and Illinois Basins are characterized by a smattering of smaller mines.

IV.2.2 US Coal Markets

Much like the rail markets, coal had a very important role in the post-industrial history of the United States. In the late 19th and early 20th centuries, coal was primarily used to produce steel. As electricity became more commonplace as the 20th century progressed, the primary role of coal shifted to supplying a energy source for power plants. The majority of this coal came from basins that have had histories of strong rail coverage, namely the Illinois and Appalachian basins, and the coal was then transported to power plants via railroads or barges.

Throughout the bulk of the 20th century coal formed the majority of the United States' electricity grid, and the most of the coal used came from the numerous small mines that dotted Appalachia. However in the 1970's, the United States became acutely aware of the negative environmental effects caused by burning coal. Chief among these were nitrogen oxides and sulfur dioxide emissions that led to acid rain and negative human health outcomes.[3] As a result, the United States passed the

Figure IV.1
Railroad Coverage and Mine Locations

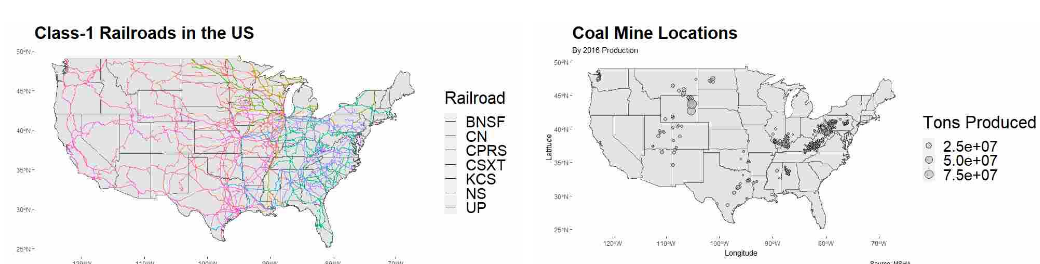

[3]See Ackerman and Hassler (1980) for a summary of the negative human health outcomes of coal and the US's response.

Clean Air Act in 1970 and a series of amendments that led to the Clean Air Act of 1990. The Clean Air Act of 1990 established cap-and-trade programs on sulfur emissions that increased the cost to generate electricity using the high-sulfur coal found in Appalachia. The electricity industry responded by shipping low-sulfur coal from the Powder River basin in Wyoming and Montana that was now much more cost-effective due to decreased rail rates (Schmalensee et al. (1998)). In the following years the Powder River basin rapidly developed a few large-scale mines, and the majority of the coal used in the United States is shipped from the Powder River basin.

Throughout the 2000s, regulations on nitrogen oxides and sulfur dioxide emissions became even more binding and talks of regional carbon taxes threatened to diminish the role of coal in the United States. Perhaps more important than increased regulation, coal faced new competition from the newly inexpensive natural gas in the 2000s. Like coal, natural gas can be burned to produce electricity and can be easily dispatched to match electricity demand, however it was used far less frequently than coal due to its relatively higher marginal cost. This began to change in the the 2000s when hydraulic fracturing became a more widely used method to extract petroleum and natural gas from wells in the United States. Hydraulic fracturing became such a successful means to extract natural gas that the average natural gas price experienced a sudden and permanent drop in 2008.[4] Power plants responded almost immediately by using more natural gas and less coal as well as converting coal generators to natural gas generators. As an additional benefit, natural gas does not produce sulfur dioxide or nitrogen oxide emissions when it is used.

The coal industry is very dependent on the rail industry to transport its products to its customers, with almost 70% of coal reaches its final destination via railroad according to EIA (2019). The rail industry is equally reliant on coal as well, with the largest plurality of rail traffic coming from shipping coal. Figure IV.2 panel (a)

[4]Hydraulic fractured wells produce so much natural gas that it is often times more preferred to burn off the excess gas rather than store it.

Figure IV.2
Annual US Rail Traffic

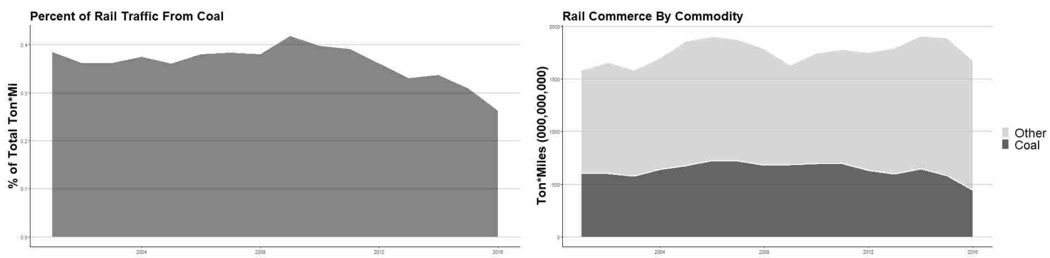

These figures were generated using the STB Confidential Waybill Sample.
Rail traffic is measured in $Tons \cdot Miles$.

presents the percentage of the total annual rail traffic in the United States that comes from coal from 2001 to 2016.[5] From 2001 to 2012, coal never made up less than 35% of the total rail traffic in the United States but then sharply declined 2013-2016 to less than 30%. Panel (b) present the total rail traffic of coal and other commodities shipped in the United States. The decline in the percentage of rail traffic by coal is due in part to a rise in the total rail traffic, but one can also see a sustained decline in coal shipped starting in 2011.

This drop in rail traffic has translated into a tangible loss in the total revenue derived from coal shipments and the importance of coal in the portfolio of commodities shipped. In Figure IV.3, I produce similar graphs to Figure IV.2 but replace the vertical axis with percent and total rail revenue derived by coal. In panel (a), the percent of revenue rail firms derive from coal shipments falls from approximately 24% to 14%. In panel (b), one can see that this is driven by an increase in revenue from shipping other commodities as well as a decrease in revenue from shipping coal.

IV.3 Review of Literature

In this research, I contribute to the economic literature on the fall of coal, the effects of energy policy and the general transportation economics literature. I begin

[5]Rail traffic is measured by $ton \cdot miles$

this section by providing a brief review of the literature of these distinct fields and conclude by providing contribution to the fields.

The demise of coal can be partially attributed to a change in electricity preferences due to fuel prices. Using hourly fuel-use data and a quasiexperimental design that leverages the fracking boom, Holladay and LaRiviere (2017) demonstrates that low natural gas prices lead to natural gas generators displacing coal generators at the bottom of the dispatch order. Knittel et al. (2016) analyze the impact of the fracking boom on power plants' decisions to use coal or natural gas. They find that the fracking boom caused electricity generation to shift away from using coal in favor of using natural gas, and this shift caused a large reduction in CO2 emissions. Linn and McCormack (2019) shows the same results hold by estimating a generator commitment model, which better captures the need for a power plant to ramp up its capacity, rather than a dispatch order model. Jordan et al. (2018) create an expected profits model and primarily attributes the large uptick in coal mine closures in Appalachia to a combination of rising marginal costs due to mine age and decreased demand for coal due to the fracking boom. Fell and Maniloff (2018) also attribute the decline of coal use for electricity generation to a combination of the rise of wind energy and the rise of natural gas. Previous research has also demonstrated that the welfare effects of shifting away from coal are positive overall, but there still exist negative welfare effects for communities that rely heavily on the coal industry (Eyer and Kahn

Figure IV.3
Annual US Rail Revenue

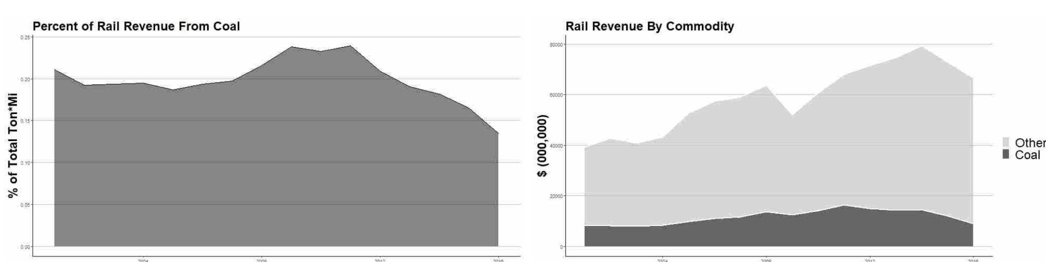

These figures were generated using the STB Confidential Waybill Sample.

(2017)). These welfare gains mainly come through cheaper electricity and decreased emissions, as outlined in Jenner and Lamadrid (2013).

The change in the portfolio of fossil fuels used in electricity is not only due to the rise of natural gas, but also due to new environmental policies. Chief among them was the Clean Air Act of 1990 and its subsequent amendments, which initiated a cap-and-trade program on sulfur dioxide and other pollutants. These policies have led to the rapid development of the Powder River Basin in northern Wyoming and a rise in the adoption of pollution-abatement technology at the power-plant level (Schmalensee et al. (1998), Popp (2003)). The Regional Greenhouse Gas Initiative (RGGI) is an auction-based cap-and-trade program in ten northeastern states that began in 2009. It has been found that RGGI has led to a large decrease in carbon emissions and a shift away from coal to natural gas in both RGGI states and potential spillover state (Fell and Maniloff (2018), Kim and Kim (2016), Murray and Maniloff (2015)). A similar policy in Europe has been shown to cause an increase in the adoption of low-carbon technology (Calel and Dechezlepretre (2016)). Renewable portfolio standards are a separate class of policies relevant to coal-powered electricity. In establishing these policies, states set long-term goals to shift some portion of their electricity generation to renewable sources and away from fossil fuels. While the details of these policies vary by state, Yin and Powers (2010) shows that they are indeed effective in reducing emissions.

There is a wide body of literature in the transportation economics field that attempts to model railroad freight rates. Post-Staggers Act, railroads are partially regulated but will not receive any regulatory oversight until their ratio of revenues to variable costs exceeds 180%. Schmalensee and Wilson (2016) argue that while the Staggers Act was successful in revitalizing the United States' railroad industry, its methods for judging the reasonableness of rates are outdated. Wilson and Wolak (2016) argue that the current measurement methods using the Universal Rail Costing

System are inadequate, and advocate that the STB instead adopt a price benchmarking approach. They present a model of this approach in Wilson and Wolak (2018). While using only data on delivered prices of coal, Dennis (1999) uses a spatial equilibrium model to predict rail rates for coal, and finds that rail deregulation led to a large decrease in the cost to transport coal from mine to power plant.

With the research conducted in this chapter, I contribute to the three broad literatures outlined above (fossil fuel choice, energy policy and transportation). I create a model to predict how changes in demand for coal and changes in potential suppliers of coal affect rail rates and the quantity of coal shipped. Often, the transportation prices are not taken into account in electricity pricing or policy evaluation models or are modelled implicitly. Unlike previous studies, I use data containing transportation rates and shipment characteristics to determine the effects that the switch away from coal has on rail rates and rail road profitability. This paper also contributes to the literature on firm pricing strategies in declining upstream markets.

IV.4 Data

The objective of this chapter is to analyze how changes in the composition of coal mine and coal-powered electric plants affect rail rates and rail traffic. To do so I use data from three main sources. While my data come at various frequencies, I aggregate all data up to the annual level. The first data source is the Mine Safety and Health Administration (MSHA). From it, I can observe the location of every mine in the United States, their production status and scale, ownership characteristics and proxies for mine size. Production status and origin-level supply change as mines open, close, temporarily idle, scale up or scale down. These data come at the quarterly level and run 2000-2017, and I use these to construct measures of origin-level coal supply as described in the previous section.

The second main data source is the Energy Information Administration (EIA) and Federal Energy Regulatory Commission's (FERC) 923 and 860 forms. These forms provide monthly data for all power plants with at least 50 MW of generating capacity in the US and run 2001-2017. The data allow me to observe power plants' locations, energy production, generator and equipment characteristics, electricity sources, characteristics of fuel, and characteristics of contracts to deliver the fuel. Coal capacity and operating status vary as power plants are built, are retired or convert their generators to run on natural gas rather than coal. From these, I create annual-level measures of demand for coal at the terminus of a shipment using a power plant's coal capacity. I further describe how I construct my coal demand variables in a later section.

The primary data source is the Surface Transportation Board's (STB) confidential waybill data. These data are collected yearly from 1984 to 2017 and allow me to observe characteristics of a subsample of all shipments by rail for all commodities. I can observe the rate charged, the weight and distance of the shipment, all rail lines used in a shipment, the commodity shipped, the route taken, and many other items. From this sample, I consider only coal shipments taken from 2001 to 2016. I use the total revenue charged for a shipment, additional fuel and miscellaneous surcharges, the weight of the shipment and the distance the shipment travels to construct the revenue per ton-miles of shipment. This is a common measure in the transportation literature and is the main measure of rail prices used.

IV.4.1 Data Preparation

The three principal data sets used for this research have issues that I address before conducting my empirical analysis. In this section, I outline the issues with the data and how I address them before conducting my analysis.

The first issue that I encounter comes from the MSHA mine data. In the MSHA's employment and production report, many mines are listed each year that produce no coal within the year due to the mine sitting idle temporarily or being shut down permanently but remaining in the data set. Although the MSHA data set does list whether a firm intends that the mine be temporarily idled or permanently shut down, this is often unreliable and does not indicate if the mine would have worked if there were sufficient demand. When I observe a mine that produces no coal in a calendar year, I consider it as being out of the market in that year and omit it when constructing any variables to measure the coal supply in a given geographic area.

The next issues arise from the EIA-FERC data on electricity generation. Many power plants can generate electricity from many different sources. While the EIA-FERC 860 Form lists all possible fuels that a generator within a power plant can use to produce electricity, the econometrician can only observe the total electricity produced by each source at the plant level. This makes identifying the extent to which each generator is used for each source implausible. I address this concern by calculating a power plant's coal electricity capacity by using only the generators that list coal as the primary fuel source in the EIA-FERC 860 data, as is common in the literature. For example, suppose a power plant with a total capacity of 1000 megawatts of capacity has coal listed as the primary source of fuel for 450 megawatts of the plant's generators but also has coal listed as the secondary fuel source for an additional 550 megawatts of capacity. In this case, I assume that this power plant has 450 megawatts of coal capacity.

As was previously discussed, the STB Waybill is a stratified sample of all rail shipments in the United States. Each observation also contains the number of cars that were present on the sampled shipment as well as an estimate of the total number of cars that are represented by similar but unobserved shipments. Following previous literature, I construct an expansion factor by dividing the total estimated cars by the total observed cars on the particular shipment. I discuss how I will use this expansion factor when I introduce my econometric model.

Within the STB waybill data, I observe the loading point and unloading point of each shipment that is sampled. Each point is called a Standard Point Location Code (SPLC) and allows me to identify the starting and ending GPS coordinates for each shipment. Unfortunately, many shipments stop temporarily in the middle of a rail shipment and then are entered into the dataset as two separate shipments when they resume. This is called rebilling and leads to many anomalies where coal shipments are observed starting in states that have never had an operational coal mine. Unfortunately I am unable to address this concern for all shipments surveyed due to the size of the data set and computational constraints. In order to address this concern as best as I can, I look at all states without a coal mine that are observed to have a waybill shipment originating from the state over the time frame of my panel, which run 2001-2016.[6] For each state, I investigate all shipments that either originate or terminate in the state. If I observe a pair of shipments A and B where shipment A terminates at the same SPLC that shipment B originates on the same day, I consider shipments A and B to be a single shipment. I then create a new observation that originates from the starting SPLC of shipment A and terminates at the end SPLC of shipment B, and delete the waybills attached to shipments A and B.

While this method eliminates many of the problematic observations described above, it does not eliminate them all. In my primary analysis, I keep all observations that I was unable to reconcile in my empirical analysis. As a robustness check, I drop all remaining shipments that I observe starting from a state without a mine and run the same models. Excluding the problematic observations does not dramatically alter my results.

[6]The seven states that meet these criteria are Nebraska, Iowa, Minnesota, Florida, South Carolina, Washington and Maine

IV.5 Methodology

Following previous literature, I model the natural log rail rates as a linear function of shipment characteristics but also add in my own measures of coal mine density at the origin of the shipment and power plant coal demand at the destination of the shipment. I develop a linear model meant to capture the effect of the declining coal industry on rail rates, quantity of coal shipped, shipped over a rail line, and the revenue earned on a shipment. Both models take the same form of:

$$y_{sodt} = \alpha_1 \cdot MineCnt_{ot} + \alpha_2 \cdot MinePrd_{ot} + \alpha_3 \cdot PlantCnt_{dt} + \alpha_4 \cdot PlantCap_{dt} + \gamma_t + \beta \cdot X_{sodt} + \varepsilon_{sodt} \tag{IV.1}$$

where s denotes each shipment, o denotes the origin of each railroad shipment, d denotes the destination of each shipment and t denote the time period as measured in years. The dependent variable y_{sodt} measures either the log of the rail rate, the log of the total weight of coal shipped on a certain shipment, or the log of the revenue earned on the shipment.[7] The regression is weighted by each observation's expansion factor.

I include fixed effects for each railroad firm and time period measured at the yearly level. The vector X_{sodt} contains controls for shipment size, distance traveled, the number of railroads used for a shipment, competing railroads, indicator variables for the primary railroad used, and fixed effects for the origin and destination states, fixed effects for intermodal shipments. I use the natural log of all non-indicator variables in the model. Following previous literature, I measure railroad competition by counting the number of Class-1 railroads within 50 miles of the origin or terminus of a shipment. Shipment size is measured by the number of cars used in a shipment. β is a vector of coefficients that will be estimated.

[7]The rail rate is measured as $Rate == \frac{Revenue}{Tons \cdot Miles}$

The coefficients α_i for $i = 1 \ldots 4$ are the principal objects of interest in this research. I quantify origin coal supply and destination coal demand using two variables each: one to capture the number of relevant firms in the area and another to capture the size of the market in the area. I take into account all coal mines and power plants in the country and measure their relevance to a coal shipment by their distance to an SPLC. In doing so, I define the coal demand (supply) around an SPLC as the sum of the number of plants (mines) divided by the distance from the destination (origin) of the shipment in the $MineCnt_{ot}$ ($PlantCnt_{dt}$). I sum up the annual production of all coal mines divided by their distances to an SPLC to produce my $MineProd_{ot}$ variable. To create my $PlantCap_{dt}$ variables, I replace annual coal mine production with annual power plant capacity.[8] Identification of these coefficients is coming through yearly variation in the locations and scales of mines and plants due to openings, closings, increases in scale, or decreases in scale of mines or plants.

The continuous versions of the variables $MineCnt_{ot}$ and $PlantCnt_{dt}$ can be interpreted as the denominators of Shepard's formula for inverse-distance weighting functions with a power parameter of 1. Equivalently, the continuous versions of the variables $MineProd_{ot}$ and $PlantCap_{dt}$ can be interpreted as the numerator. Although inverse-distance weighting is commonly used to generate unknown values of area coverage, it is not an appropriate measure here. Inverse-distance weighting is used to interpolate values when the space has a uniform coverage of observations. Mine coverage in the United States is clustered among the four basins, and interpolating coal supply between these basins would lead to areas areas with no mines for hundreds of miles being observed as having high coal supply. Further, creating a single measure

[8]That is, coal mine count near a given origin SPLC o at time t is defined as $MineCnt_{ot} = \sum_{n=1}^{n} \frac{Mine_{nt}}{Distance_{on}}$ where N denotes the total number of mines in the United States, and $Mine_{nt}$ takes on a value of 1 if mine n is operating in year t, and takes on a value of 0 otherwise. Coal mine production near a given origin SPLC o at time t is defined as $MineProd_{ot} = \sum_{n=1}^{n} \frac{Prod_{nt}}{Distance_{on}}$ where N denotes the total number of mines operating in the US, and $Prod_{nt}$ is the total coal extracted by mine n in year t. In both specifications, $Distance_{on}$ is the distance between a mine n and the relevant origin SPLC o.

of coal mine or power plant coverage would mitigate my ability to measure both the quantity of coal produced and the number of competitors in an area.

As robustness checks, I employ alternative definitions of mine and power plant counts and mine and power plant sizes. Instead of using all power plants or mines in the continental US, I instead measure the demand for coal near the destination as the quantity of power plants with coal generators within some radius of the endpoint of the shipment, and use an analogous method to measure the supply of coal near the origin of the shipment using coal mines. I present results using radii of 10 miles, 20 miles, 50 miles, 100 miles and 200 miles. I select which mines and power plants to include in this calculations using the criteria outlined in the data section. My alternative measures of plant or mine scale are calculated by summing all the coal production or power plant capacity within a relevant radius.

As a final robustness check, I also estimate versions of this model with all exogenous and endogenous variables aggregated up to the origin-destination-railroad-year level. I use my constructed expansion factors to properly scale up and weight all variables. I present the results of the estimation of my annual-level models in the appendix.

IV.6 Descriptive Graphs and Statistics

In this section, I present graphs and figures in the form of time series plots of variables of interest, and maps demonstrating characteristics of railroad shipments.

In Figure IV.4, I plot all SPLCs in the United States that have been used to load or unload coal in the United States. In general, the Appalachian region is populated more densely with SPLCs that are used to load coal anywhere else in the country. Notably, there are very few origin SPLCs in around the Powder River Basin. Almost all states in the country have a terminal SPLC in them, but they are still not uniformly

Figure IV.4
SPLC Locations

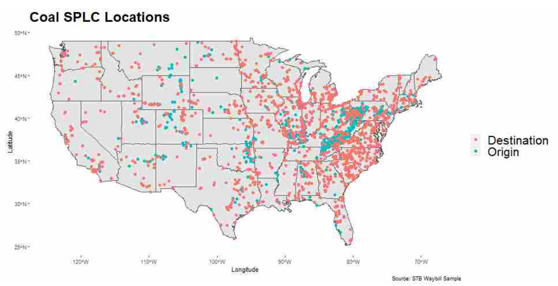

spread across the continental United States. There are far more terminal locations in the east coast of the United States than in other regions.

In addition to mapping out the SPLCs, in Figure IV.5 I also map out the locations of coal mines and power plants that have the ability to generate electricity using coal. The left panel of Figure IV.5 contains a map of all mines that were active in 2016 with each mine scaled up by the tons of coal that the mine produced in 2016. A few things should be very apparent in the left panel. First, the eastern half of the United States' coal mining industry is characterized by clusters of smaller mines in three distinct basins, namely the Northern Appalachian basin, the central Appalachian Basin and the Illinois Basin. The western half of the United States is largely dominated by just a few mines that are concentrated in Wyoming's Powder River Basin.

In the right panel of on Figure IV.5, I map all fossil fuel energy power plants in the United States. I color code the plants to denote which plants can use both coal and natural gas, and scale the plants to reflect each plant's coal electricity generation capacity. The size of a plant varies a lot throughout the United States and many plants appear to have the ability to use both coal and natural gas to generate electricity. Perhaps the most important takeaway from this graph is that most fossil fuel power plants are located in the eastern half of the United States, so one would expect that most coal shipped within the United States would reach it final destination in the east. These graphs match up well with the map of SPLCs used to load or unload

Figure IV.5
2016 Coal Mine and Power Plant Locations

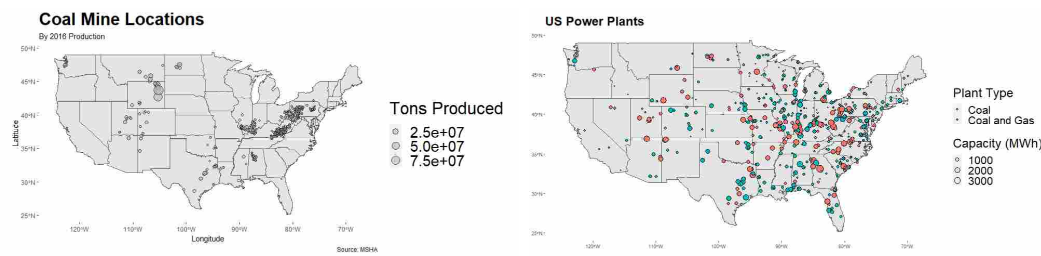

coal in the United States in Figure IV.4. This provides some initial evidence that coal mines are not willing to transport coal very far before loading onto a train, and power plants are not willing to transport coal very far after offloading it.

Based on the clustering of coal mines in the four main basins and the density of power plants in the east coast it's expected that most coal shipments originate in four main basins described above, and most coal shipments terminate in the eastern part of the United States. To demonstrate that this is indeed the case, in Figure IV.6, I create two maps to show the patterns of the origins and destinations of coal shipments in 2016. In the left panel, I present the weight of all coal shipments originating from each state and in the right panel I show the weight of all coal shipments that terminate in each state. I weight the areas in each graph by an observed shipment's expansion factor before aggregating in both graphs. Within the left panel of this graph, I want to bring the reader's attention to one key point: most coal in the United States is coming from Wyoming, Illinois and Appalachia, and this should be no surprise given the locations and scales of mines shown in Figure IV.5.

In the right panel of Figure IV.6, I map out the destinations of all coal shipments in 2016. Much like the right panel of this graph, the presented quantities of coal shipped are likely biased by some shipments temporarily stopping in a state before reaching their final destination. The eastern half of the country appears to be the largest destinations of coal shipments, with Texas, Missouri, Illinois and Indiana seeming to be the largest recipients of coal shipments. These trends again seem to match well with the map provided in the left panel of Figure IV.5, as the eastern half of the United States is home to far more fossil fuel power plants than the western half.

IV.6.1 Summary Statistics

Table IV.1 contains summary statistics for my shipment-level regressions. All variables have a wide distribution of values, particularly the shipment weight and the rail rate. Although I use the log of the rail rate as my primary dependent variable and use the log of the shipment weight as an explanatory variable, I present them here in their level forms. One can clearly see that there is a very large spread in the density of coal mines around an SPLC. This variance is driven by spatial differences in SPLC location, differences in the types of mines in each basin, and annual variation in the number of operating mines.

A similar spread can also be seen in the summary statistics for plants near the end of a coal shipment. While the number of plants is not nearly as large as the number of coal mines near any SPLC, there is still a large relative variance in the density of coal-fired power plants within the data. Much like the coal mine data, the variation is coming primarily through spatial differences and coal generators being retired or converted to natural gas generators over time.

Finally, in Figure IV.7 I generate four annual time series graphs of variables of interest from the years 2001 to 2016. In all panels, the variables of interest are

Figure IV.6
2016 Coal Shipments

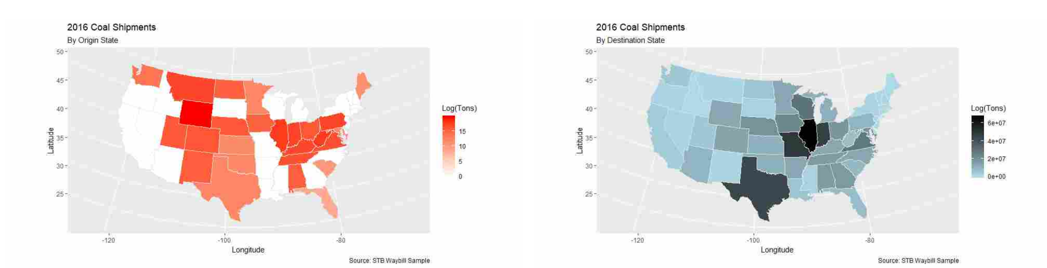

weighted by the expansion factor before aggregation. In panel (a), one can see that the the total weight of coal shipped by rail began around 2008 after many years of sustained activity. This coincides with the initial fracking boom and the start of the Great Recession. The large drop can also be seen in the total revenue generated by coal shipments in panel (c). Panel (b) shows a relatively steady decrease in the number of coal shipments over the course of the data, and panel (d) shows a coal rail rate that appears to be trending upward until 2008, then experiencing noisy changes in the ensuing years.[9] This may indicate that rail companies had to undergo a period of learning after the Great Recession and the initial natural gas boom in 2008.

IV.6.2 Basin Heterogeneity

Anecdotally, the shift away from coal in the United States has not been uniform across basins. To investigate this further, I provide a separate analysis for each basin. In Table IV.2 I provide means of the same variable that appear in Table IV.1 broken

Table IV.1
Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Max
Cars	500,990	104.617	36.750	1	262
Distance (mi)	500,990	733.673	485.146	1.000	3,315.400
Weight	500,990	11,973.400	5,394.147	1	559,100
Revenue ($)	500,990	177,103.700	137,057.300	1	8,441,737
Rate ($/(Tons·Mi))	500,990	0.062	0.988	0.00000	321.902
Weight	500,990	26,548.490	11,721.250	2	1,677,300
RRs Used	500,990	1.156	0.394	1	6
Nearby Mine Count	500,990	0.003	0.004	0.0002	0.025
Nearby Mine Prod	500,990	7.165	7.585	0.243	39.955
Nearby Plant Count	500,990	0.001	0.0005	0.0001	0.007
Nearby Plant Cap	500,990	0.001	0.001	0.0001	0.007
Origin RR Comp	500,990	6.021	1.862	1	8
Term RR Comp	500,990	3.525	1.470	1	8

[9]The rail rate is measured as $Rate = \frac{Revenue}{Tons*Miles}$

Figure IV.7
Annual Trends in Coal Commerce

(a) Tonnage Shipped

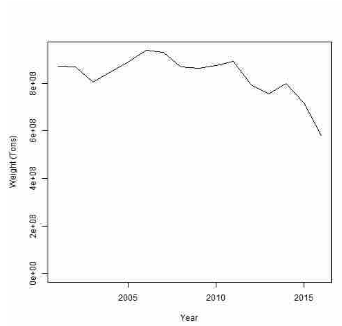

(b) Quantity of Coal Shipments

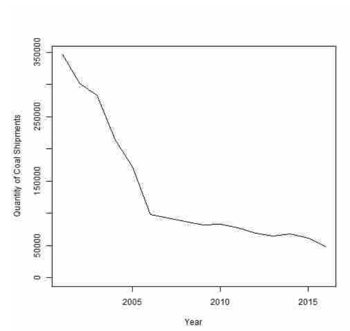

(c) Coal Shipment Revenues

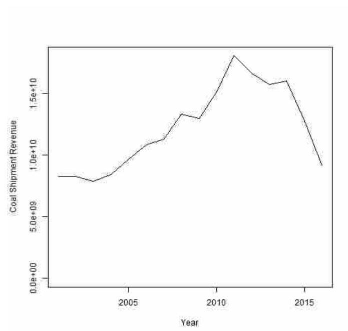

(d) Coal Shipment Rate

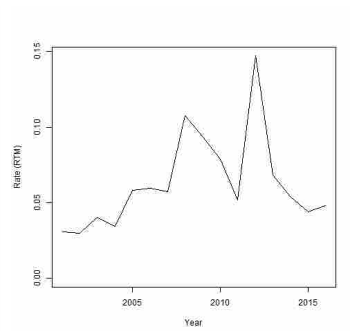

down by basin.

A few things should be apparent. First, the rate charged to transport coal from the Powder River Basin is the lowest of all basins. This lower rate seems to be a result of some form of economies of scale, as shipments from the Powder River Basin travel the farthest of shipments from any other basin and also carry the most cars.

As has been previously pointed out, the Powder River Basin is characterized by very few mines that are very large, while Appalachia and the Illinois Basin are composed of numerous smaller mines. This can be seen in my two measures of coal supply near the origin of shipments. The Powder River Basin has the lowest measured count of mines near the origin, but has the highest measured mine production near the origin. This also indicates that these two separate measures capture different aspects of local coal-mining markets, with the mine count capturing the competitiveness of the local market and the scale capturing the market size. In my basin heterogeneity analysis, I include both coefficients to capture how these shippers react to these differences in market structures between basins.

In Figure IV.8 I highlight more historical differences between basins. In all subfigures, all aggregation is done using the expansion factor that was previously discussed. Panel (a) displays the total annual rail traffic from coal transport coming from each basin. It is apparent that the Powder River Basin has the most coal rail traffic and has had the largest decline of the major US basins, but it should also be noted that the traffic in other basins is also falling. A similar decline can be seen in panel (b), which shows the total revenue derived from coal shipments originating in each basin. While the Powder River Basin still makes up the largest portion of rail revenues of all basins, it should be noted that Appalachia constitutes a comparable amount of revenue. In conjunction with panel (a), this again demonstrates the differences in the rate charged for coal shipments from each basin.

In panels (c) and (d), I present two measures of coal mine supply near the origin of the shipments, with my measure of coal mine count in panel (c) and my measure of

Table IV.2
Summary Statistics By Basin

Basin	Appalachia	Illinois	Powder River	Uinta
Rate ($/(Ton· Mile))	0.065	0.078	0.05	0.077
Weight (Ton)	9114.034	11439.044	14459.593	10392.404
Revenue ($)	148039.355	88911.306	218258.172	216338.509
Distance (Mi)	407.978	209.403	1081.794	943.883
Cars	83.682	100.221	123.074	94.932
Nearby Plant Count	0.001	0.001	0.001	0.001
Nearby Plant Capacity	0.001	0.001	0.001	0
Nearby Mine Count	0.007	0.002	0.001	0.001
Nearby Mine Prod	3.069	1.694	12.053	2.767
Origin RR Comp	4.99	5.37	7.308	4.492
Terminal RR Comp	3.352	3.94	3.652	3.626
RRs Used	1.072	1.212	1.19	1.386
Obs	178593	38946	240181	23161

Figure IV.8
Annual Trends in Coal Commerce by Basin

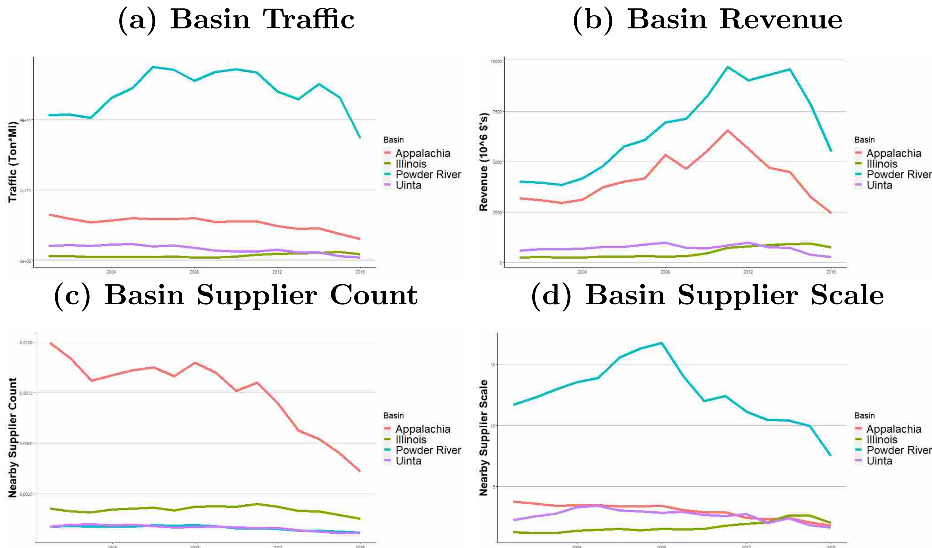

coal mine scale in panel (d). These two figures further highlight the differences in the coal mining markets in Appalachia and the Powder River Basin. One can see how the number of suppliers near a typical origin of a shipment in Appalachia dominates any other basin in panel (c), but the sheer scale of the mining operations in the Powder River Basin dwarf all other areas in the US in panel (d).

As a final way to discuss the differences in shipping strategies between basins, I present the average annual rate charged per shipment for shipments originating from each basin in Figure IV.9. All average rates are computed using my derived expansion factor. While rail rates are noisy from year to year, it can be seen once again that coal shipments from the Powder River Basin typically are typically the least expensive expensive per unit.

IV.7 Empirical Results

Table IV.3 contains parameter estimates of of equation (1) to predict coal rail rates, quantity shipped per shipment, and revenue generated per shipment. Appendix

Figure IV.9
Annual Rail Rates by Basin

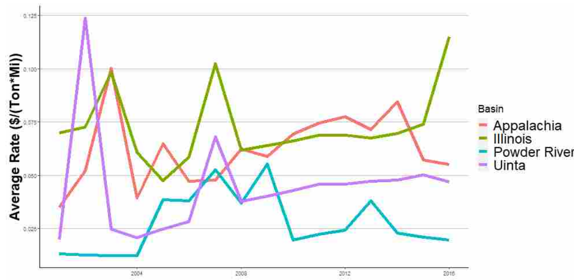

Tables A.IV.3A-A.IV.3E contain the results of estimating my model with the power plant coal demand and mine coal supply variables respecified as radii around the terminus and destination of the shipment. I only include the estimates of my four variables that quantify the mine coal market supply and the power plant coal market demand in the appendix tables.

I first want to bring attention to my controls. The natural log of the distance a shipment travelled and the number of cars that are on a shipment have negative relationships to the rate charged for a shipment, indicating that there are economies of scale in shipping. Revenue earned from shipments rises as the number of cars on a shipment and the distance a shipment increases as one would expect.

The primary coefficients of interest are my measures of demand for coal by power plants and the supply of coal by mines. In the rate column, the coefficient attached to the total number of power plants near a shipment is positive, indicating that as the number of power plants in an area increases, the rate charged to ship coal to that area increases. This indicates that when controlling for scale, coal power plants with more market power may have more bargaining power. The estimated effect of my measure of market size on the transportation rate is negative, indicating that there exist economies of scale in shipping coal to larger markets.

The coefficient estimates on the effects of power plant market characteristics on the the size of the shipment and the revenue derived from shipments tell similar stories.

Table IV.3
Shipment-Level Results

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	0.174^{***}	-0.031^{***}	0.143^{***}
	(0.003)	(0.002)	(0.002)
Plant Cap.	-0.153^{***}	0.036^{***}	-0.117^{***}
	(0.002)	(0.001)	(0.002)
Mine Count	0.193^{***}	-0.023^{***}	0.171^{***}
	(0.003)	(0.002)	(0.002)
Mine Prod.	-0.133^{***}	0.026^{***}	-0.107^{***}
	(0.002)	(0.001)	(0.002)
Origin RR Comp	0.068^{***}	0.011^{***}	0.079^{***}
	(0.004)	(0.002)	(0.003)
Terminal RR Comp	0.020^{***}	-0.004^{**}	0.016^{***}
	(0.003)	(0.002)	(0.003)
Distance	-0.569^{***}	-0.009^{***}	0.422^{***}
	(0.001)	(0.001)	(0.001)
Cars	-0.058^{***}	1.011^{***}	0.953^{***}
	(0.001)	(0.0003)	(0.001)
RRs Used	0.307^{***}	-0.011^{***}	0.296^{***}
	(0.004)	(0.002)	(0.004)
Intermodal	-0.044^{***}	0.0002	-0.044^{***}
	(0.004)	(0.002)	(0.003)
Constant	2.476^{***}	4.835^{***}	7.312^{***}
	(0.077)	(0.041)	(0.066)
Observations	500,990	500,990	500,990
R^2	0.817	0.990	0.978
Residual Std. Error (df = 500814)	0.872	0.468	0.748

Note: *p<0.1; **p<0.05; ***p<0.01

Fixed effects for primary railroad, origin state, destination state, and year are used in all models presented above.

As the generation capacity near the terminus of a coal shipment increases, rail firms tend to send larger quantities of coal on each shipment. Although rail firms charge lower rates on these shipments, the size of the shipments is large enough that the net effect on revenue is positive as shown in column (3). Conversely, as the number of plants near the terminus of a shipment increases, the size of each individual shipment falls which leads to less revenue per shipment.

The relationship between the supply of coal and shipment characteristics is similar to that of coal demand and shipment characteristics. The coefficient attached to the number of coal mines in an area is positive in the rate regression and negative in the regressions on shipment weight and shipment revenue. This could indicate that shippers have less pricing power when there are fewer relevant agents near the origin of a shipment. Ultimately, the revenue per shipment goes down as the number of mines in an area increases all else being held equal.

The coefficients attached the total coal production in an area are all negative, demonstrating that shipments originating in areas with a large supply of coal are charged lower rates, ship less coal per shipment and ultimately bring in less revenue per shipment on the intensive margin. This demonstrates that as the supply of coal contracts, a rail firm chooses to charge a higher rate for each individual shipment.

This general story can be seen under my alternative specifications of my coal supply and coal demand variables and in my annual level regressions. In all of the Appendix Tables A.IV.4A-A.IV.4E, the signs attached to my four main variables of interest are unchanged.

IV.7.1 Robustness Check

As a robustness check, I aggregate my models up to the annual-railroad-origin-destination level and run my models with the same parameters. All aggregation was done using the expansion factor to create weighted averages for all variable means and to create accurate variable sums. The result of the annual model estimation is presented in Table IV.4. As was the case with my base model, the log of all non-indicator variables are used in the aggregated models. The regressions presented in Table IV.4 are no longer weighted by the expansion factor because the expansion factor was used to aggregate variables. Appendix Tables A.IV.4A-A.IV.4E contain the results of estimating my model with the power plant coal demand and mine coal supply variables respecified as radii around the terminus and destination of the shipment. For the sake of brevity, I only present the coefficient estimates of the four variables of interest in Appendix Tables A4.A-A4.E.

Attention should be brought to two things in the annual-level models. First, neither the model fit nor the sign and magnitude of the controls vary greatly from the shipment-level regressions presented above. Second, although the magnitude of my four measures of coal market supply and power plant market demand are diminished, the signs of all the variables are the same in all three regressions. The same trends can be seen in Appendix Tables A.IV.4A-A.IV.4E as well.

IV.7.2 Basin Heterogeneity

As was discussed previously, characteristics of each basin vary dramatically. The Powder River Basin has a small number of exceptionally large mines, while the Illinois Basin and Appalachia have numerous smaller mines. While my measures of mine coverage capture some of this heterogeneity, there is still reason to suspect that a coal shipper treats market characteristics differently between basins. For example, one may expect that a shipper servicing Appalachia would respond more strongly to a mine shut down than one servicing the Powder River Basin. I present the coefficient estimates of my two measures of coal supply in each basin for on all three of my outcome variables in Table IV.5. The top panel contains the results on the rate charged, the middle panel contains the results on the size of the shipment and the bottom panel contains the results for revenue earned per shipment. All other covariates are identical to the models models I presented in Table IV.3.

It should be apparent that across all specifications the sign, the magnitude and the significance of the coefficients vary greatly, which provides evidence that when choosing their optimal shipping prices, rail firms treat coal from each basin as separate commodities. Of particular interest, the rate charged by shippers has a positive relationship to the number of nearby operational mines in Appalachia, but all other basins in the United States have an estimated negative coefficient. This indicates that shippers raise the price to ship coal out of Appalachia when there are many operational mines, likely due to negotiating power. As can be seen in the bottom panel, this coefficient estimate translates into more revenue earned per shipment when there are more operating mines in the area.

The same shipping rate story does not hold in any of the other major basins in the US, particularly in the Powder River Basin. Shipper in the Powder River Basin appear to only respond to nearby coal output rather than the number of competitors when making their pricing decisions. In particular, in the Powder River Basin coal shippers set lower prices when the quantity of coal produced in the area is large. These higher quantities of coal produced in the area are correlated with larger shipments and ultimately more revenue earned per shipment.

IV.8 Conclusion

I investigate the relationship between the fall of coal in the United States and its effects on the rail industry. I set out to discuss how two separate components of the coal industry, namely coal mining and generating electricity using coal, affect rail rates, coal rail traffic and revenue earned from shipping. Coal and railroads are particularly strong complements to each other, with the majority of coal in the

United States being shipped via railroads and the largest portion of rail commerce in the United States being coal. So, the decline of one of these industries will naturally have a profound impact on the other.

Using various definitions of the supply and demand for coal that are relevant to a particular shipment, I run various models connecting coal supply and demand to the rate charged to ship coal, the quantity of coal carried on a shipment and the the total revenue earned on a shipment. At first glance, rail rates appear to change through two channels: the number of agents in a relevant market and the size of the total supply or demand in the relevant market. For both mines and power plants, it appears that rail firms raise their prices when there is an in the count of nearby agents, which can be explained by an increase in the number of parties demanding a shipment. Conversely, the rates charged by a rail firm appear to fall as the relevant markets grow in total size, which indicates that there are economies of scale in shipping to areas with high demand.

While these market characteristics succeed in explaining the entire US market in broad terms, it doesn't fully capture how rail firms may react to changes in different basins. To address this I run separate regressions for each of the four major basins in the United States, and it is clear that shippers view changes in basins differently. In particular, shippers in the Powder River Basin do not appear to take the number of mines into account when setting rates, and shippers in Appalachia appear to raise rates when there are many active mines in basin.

Table IV.4
Annual-Level Results

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	0.129^{***}	-0.026^{***}	-0.140^{***}
	(0.009)	(0.004)	(0.029)
Plant Cap.	-0.114^{***}	0.039^{***}	0.371^{***}
	(0.007)	(0.003)	(0.022)
Mine Count	0.147^{***}	-0.012^{***}	-0.147^{***}
	(0.009)	(0.004)	(0.027)
Mine Prod.	-0.121^{***}	0.016^{***}	0.144^{***}
	(0.006)	(0.003)	(0.020)
Origin RR Comp	-0.001	-0.003	0.077^{**}
	(0.011)	(0.005)	(0.035)
Terminal RR Comp	-0.065^{***}	-0.001	-0.095^{***}
	(0.010)	(0.004)	(0.030)
Distance	-0.559^{***}	0.0002	0.188^{***}
	(0.004)	(0.002)	(0.013)
Cars	-0.104^{***}	1.027^{***}	0.357^{***}
	(0.003)	(0.001)	(0.008)
RRs Used	0.284^{***}	0.008	-0.075^{*}
	(0.013)	(0.006)	(0.040)
Intermodal	0.298^{***}	-0.068^{***}	0.531^{***}
	(0.034)	(0.016)	(0.105)
Constant	1.119^{***}	5.045^{***}	12.637^{***}
	(0.196)	(0.090)	(0.598)
Observations	34,987	34,987	34,987
R^2	0.761	0.970	0.336
Residual Std. Error (df = 34811)	0.406	0.186	1.235

Note: *p<0.1; **p<0.05; ***p<0.01

Fixed effects for primary railroad, origin state, destination state, and year are used in all models presented above.

Table IV.5
Shipment-Level Basin Heterogeneity

	Appalachia	Powder River	Illinois	Uinta
	(1)	(2)	(3)	(4)
		Log(Rate)		
Mine Count	0.207***	−0.007	−0.607***	−0.215***
	(0.003)	(0.007)	(0.031)	(0.027)
Mine Prod	−0.191***	−0.046***	0.151***	0.142***
	(0.002)	(0.004)	(0.014)	(0.020)
R^2	0.772	0.787	0.740	0.709
		Log(Weight)		
Mine Count	−0.016***	0.006	0.023	−0.027
	(0.001)	(0.004)	(0.016)	(0.023)
Mine Prod	0.062***	−0.006**	−0.002	0.035**
	(0.001)	(0.002)	(0.007)	(0.018)
R^2	0.996	0.974	0.993	0.826
		Log(Revenue)		
Mine Count	0.191***	−0.001	−0.584***	−0.242***
	(0.002)	(0.005)	(0.027)	(0.016)
Mine Prod	−0.129***	−0.052***	0.149***	0.177***
	(0.002)	(0.003)	(0.012)	(0.012)
R^2	0.989	0.977	0.987	0.948
Observations	178,593	240,181	38,946	23,161

Note: *p<0.1; **p<0.05; ***p<0.01

CHAPTER V

DISSERTATION CONCLUSION

In this dissertation, I investigate three how three aspects of the coal market have changed in the modern world: coal mines, fossil fuel power plants and the rail industry used to connect most mines and power plants in the United States. In the first chapter, I demonstrate that although environmental policy played a part in coal mine closure, the largest contributor to the fall in coal mine participation is the sharp drop in the natural gas price due to fracking. I also show that fixed costs in coal mining are very high. The results of the research suggest that policies targeted at providing incentives to reopen mines will be largely unsuccessful.

Although natural gas's rise to prominence is the largest reason that coal mines close down, it is not the only reason. Carbon cap-and-trade programs are a relatively new policy that have further stymied coal use at power plants. It has previously been shown that carbon cap and trade programs cause a substitution away from coal and towards natural gas. In the second chapter of my dissertation, I show that this substitution is heterogeneous throughout the year, with the largest substitutions away from coal occurring in the fall and the spring when the demand for electricity is the lowest. This means that a utility that wishes to shut down coal plants will have issues if it unable to find a suitable power source in the higher demand summer and winter months.

The first two chapters of my dissertation demonstrate that both the number of suppliers and the market demand for coal are dropping in the United States. My third chapter discusses how this has a non-trivial spillover into the railroad industry, which has a traditionally strong tie to coal markets. I find that rail rates and revenues respond to shrinkages in the number of operational coal mines and coal power plants, and that shippers view changes in market characteristics differently between basins. The results of this research can be used to inform regulators how to properly set guidelines for rail rate setting.

APPENDIX

Table II.5A
Base Linear Probability Model Results

VARIABLES	(1)	(2)
Real Nat. Gas Price	0.00132**	0.000791
Log(Avg # of Employees)	0.240***	0.250***
Log(Total Extraction)	0.118***	0.100***
Log(Total Extraction)^2	-0.00706***	-0.00528***
Multi-Mine Firm (0/1)	-0.0411***	-0.00510
Total Mines in Company	-0.000280***	9.73e-05
Cumulative Violations Incurred by Mine	9.20e-06***	3.10e-05***
γ_0	0.174***	-0.0987***
γ_1	0.0371***	0.0454***
γ_2	0.0135***	0.0484***
γ_3	-0.00820*	0.0421***
γ_4	-0.0327***	0.0278***
γ_5	-0.0597***	0.00702
γ_6	-0.0580***	0.0141**
γ_7	-0.0733***	0.00711
γ_8	-0.0812***	0.0202***
Clean Air Interstate Rule (0/1)	-0.0127***	-0.00268
NAAQS Revision (0/1)	0.00539	0.00550
Time Trend	0.00123***	-0.00182***
Mean Shift for t>2009q1	0.309***	0.0229
Trend for t>2009q1	-0.00162***	-0.000115
Observations	93,837	89,856
Number of Mines Used	3,981	3,884
Mine Effects	RE	FE

AR(1) Standard errors used
*** p<0.01, ** p<0.05, * p<0.1

Table II.5B
Base Linear Probability Model Results

VARIABLES	(1)	(2)	(3)
Real Nat. Gas Price	0.00439***	0.00349***	0.00332***
Log(Avg # of Employees)	0.167***	0.216***	0.227***
Log(Total Extraction)	0.0866***	0.0919***	0.0897***
Log(Total Extraction)^2	-0.00538***	-0.00616***	-0.00609***
Multi-Mine Firm (0/1)	-0.0518***	-0.0297***	-0.00505
Total Mines in Company	-0.000320***	-0.000165**	0.000135
Cumulative Violations Incurred by Mine	-2.45e-05***	1.50e-05***	2.35e-05***
γ_0	0.312***	0.205***	0.194***
γ_1	-0.304***	-0.199***	-0.176***
γ_2	-0.355***	-0.227***	-0.200***
γ_3	-0.381***	-0.239***	-0.209***
γ_4	-0.393***	-0.243***	-0.211***
γ_5	-0.411***	-0.257***	-0.225***
γ_6	-0.402***	-0.244***	-0.211***
γ_7	-0.400***	-0.239***	-0.207***
γ_8	-0.427***	-0.243***	-0.208***
Clean Air Interstate Rule (0/1)	-0.0223***	-0.0261***	-0.0256***
NAAQS Revision (0/1)	0.0109**	0.00787	0.00574
Time Trend	0.00230***	0.00270***	0.00265***
Mean Shift for t>2009q1	0.350***	0.377***	0.343***
Trend for t>2009q1	-0.00184***	-0.00196***	-0.00178***
Observations	87,956	87,956	87,956
Mine Effects	N/A	RE	FE
Number of Mines Used		3,817	3,817

Clustered standard errors used
*** p<0.01, ** p<0.05, * p<0.1

Table A.III.1
Model 1 Coefficent Estimates: Coal

Dep Var: Log(NetGen)	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
RGGI	-0.406***	-0.500***	-0.719***	-0.906***	-0.843***	-0.782***	-0.471***	-0.736***	-0.957***	-0.940***	-0.891***	-0.677***
	(0.117)	(0.0728)	(0.0733)	(0.116)	(0.228)	(0.220)	(0.155)	(0.151)	(0.217)	(0.210)	(0.122)	(0.143)
Log(Capacity)	0.820***	0.879***	0.891***	0.942***	0.916***	0.889***	0.859***	0.883***	0.928***	0.963***	0.900***	0.890***
	(0.0500)	(0.0466)	(0.0390)	(0.0373)	(0.0400)	(0.0409)	(0.0446)	(0.0431)	(0.0480)	(0.0367)	(0.0438)	(0.0438)
State Renewable %	-0.738***	-0.889***	-0.768***	-0.520***	-0.597***	-0.732***	-0.768***	-0.928***	-0.351	-0.583***	-0.642***	-0.899***
	(0.232)	(0.262)	(0.176)	(0.164)	(0.183)	(0.196)	(0.224)	(0.252)	(0.269)	(0.205)	(0.192)	(0.203)
Trend	-0.00274***	-0.00281***	-0.00347***	-0.00424***	-0.00412***	-0.00348***	-0.00331***	-0.00352***	-0.00401***	-0.00401***	-0.00383***	-0.00341***
	(0.000415)	(0.000410)	(0.000444)	(0.000480)	(0.000492)	(0.000461)	(0.000477)	(0.000468)	(0.000527)	(0.000516)	(0.000532)	(0.000488)
Scrubber	0.125***	0.107**	0.141***	0.138***	0.182***	0.183***	0.180***	0.154***	0.164***	0.116**	0.115*	0.108**
	(0.0448)	(0.0476)	(0.0479)	(0.0530)	(0.0525)	(0.0419)	(0.0464)	(0.0431)	(0.0562)	(0.0542)	(0.0607)	(0.0516)
SCR/SNR	0.0867**	0.144***	0.135***	0.193***	0.162**	0.145***	0.158***	0.167***	0.148**	0.109	0.121*	0.153***
	(0.0436)	(0.0447)	(0.0489)	(0.0629)	(0.0661)	(0.0553)	(0.0551)	(0.0554)	(0.0620)	(0.0684)	(0.0631)	(0.0564)
Elec. Precipitator	0.562***	0.443***	0.390***	0.314***	0.347***	0.421***	0.424***	0.378***	0.331**	0.271**	0.381***	0.362***
	(0.132)	(0.125)	(0.125)	(0.121)	(0.119)	(0.114)	(0.115)	(0.116)	(0.129)	(0.117)	(0.127)	(0.124)
Constant	2.538***	2.120***	1.094	-0.443	-0.0978	1.098	1.600**	1.249*	0.0182	-0.0742	0.468	1.351*
	(0.649)	(0.612)	(0.675)	(0.713)	(0.748)	(0.707)	(0.712)	(0.697)	(0.790)	(0.773)	(0.830)	(0.742)
Observations	9,641	9,576	9,461	9,262	9,336	9,536	9,588	9,592	9,448	9,150	9,201	9,331
Number of Plants	724	724	725	725	725	724	726	724	724	721	724	721

Standard errors are clustered at the state level. Unclustered errors and power-plant level clusters result in smaller standard error estimates.
*** p<0.01, ** p<0.05, * p<0.1

Figure A.III.1
Scrubber Use

(a) Coal **(b) Gas**

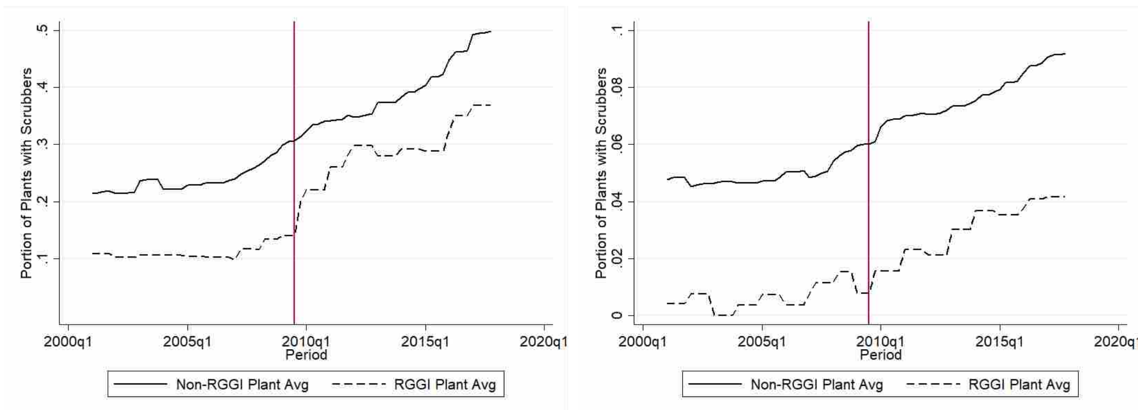

Note: The data for these graphs are compiled from EIA Form 860.

Table A.III.2
Model 1 Coefficent Estimates: Natural Gas

Dep Var: Log	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
RGGI	0.346**	0.276**	0.328***	0.267**	0.431***	0.347***	0.540***	0.350***	0.240**	0.204**	0.209**	0.148*
	(0.141)	(0.129)	(0.112)	(0.110)	(0.110)	(0.0746)	(0.0829)	(0.0905)	(0.104)	(0.0917)	(0.0844)	(0.0879)
Log(Capacity)	0.674***	0.703***	0.669***	0.678***	0.718***	0.693***	0.706***	0.698***	0.658***	0.674***	0.677***	0.679***
	(0.0606)	(0.0556)	(0.0530)	(0.0589)	(0.0542)	(0.0583)	(0.0539)	(0.0598)	(0.0557)	(0.0522)	(0.0569)	(0.0545)
State Renewable %	-0.878	-0.843*	-0.855***	-0.755***	-1.117***	-1.703***	-2.292***	-1.862**	-1.250**	-1.099**	-0.903*	-0.911
	(0.572)	(0.436)	(0.274)	(0.275)	(0.394)	(0.454)	(0.647)	(0.774)	(0.558)	(0.438)	(0.528)	(0.598)
Trend	-0.000986	-0.00120	-0.000951	-0.00114	-0.000938	-0.000847	-0.00125*	-0.00184***	-0.000693	-0.000657	-7.42e-05	-0.000264
	(0.000896)	(0.000911)	(0.000879)	(0.000799)	(0.000990)	(0.000684)	(0.000643)	(0.000534)	(0.000660)	(0.000735)	(0.000936)	(0.000925)
Scrubber	-0.127	-0.286*	-0.136	-0.0488	-0.0675	-0.249*	-0.113	-0.290*	-0.164	-0.106	-0.193	-0.0694
	(0.152)	(0.167)	(0.154)	(0.194)	(0.148)	(0.143)	(0.137)	(0.172)	(0.163)	(0.153)	(0.165)	(0.167)
SCR/SNR	0.650***	0.619***	0.589***	0.618***	0.669***	0.542***	0.513***	0.503***	0.436***	0.461***	0.558***	0.566***
	(0.140)	(0.147)	(0.145)	(0.141)	(0.144)	(0.142)	(0.129)	(0.128)	(0.141)	(0.151)	(0.150)	(0.123)
Elec. Precipitator	-0.611*	-0.620*	-0.558*	-0.582	-0.640*	-0.643*	-0.802**	-0.681*	-0.619*	-0.542	-0.552*	-0.584*
	(0.321)	(0.319)	(0.335)	(0.359)	(0.360)	(0.341)	(0.321)	(0.352)	(0.362)	(0.391)	(0.331)	(0.332)
Constant	3.349**	2.817**	3.399***	3.053***	3.424**	4.068***	3.864***	2.907***	4.230***	3.990***	4.595***	4.382***
	(1.309)	(1.300)	(1.253)	(1.157)	(1.400)	(0.971)	(0.969)	(0.869)	(0.970)	(1.071)	(1.325)	(1.305)
Observations	29,105	28,761	29,088	29,445	30,229	31,075	31,485	31,679	30,949	29,790	29,325	29,487
Number of Plants	2,784	2,783	2,792	2,794	2,805	2,813	2,816	2,819	2,806	2,793	2,782	2,785

Standard errors are clustered at the state level. Unclustered errors and power-plant level clusters result in smaller standard error estimates.
*** p<0.01, ** p<0.05, * p<0.1

Table A.III.3
Model 2 Coal Coefficent Estimates

VARIABLES	(1) Log(MWh)	(2) Log(MWh)	(3) Log(MWh)	(4) Log(MWh)
Jan	-0.366***	-0.333***	-0.335***	-0.322***
	(0.0921)	(0.0883)	(0.0878)	(0.0874)
Feb	-0.429***	-0.410***	-0.410***	-0.393***
	(0.0725)	(0.0714)	(0.0718)	(0.0708)
Mar	-0.678***	-0.639***	-0.630***	-0.620***
	(0.0628)	(0.0703)	(0.0688)	(0.0700)
Apr	-0.972***	-0.913***	-0.908***	-0.907***
	(0.128)	(0.147)	(0.146)	(0.144)
May	-1.059***	-1.022***	-1.007***	-1.001***
	(0.216)	(0.224)	(0.223)	(0.221)
Jun	-0.838***	-0.815***	-0.810***	-0.793***
	(0.229)	(0.230)	(0.231)	(0.231)
Jul	-0.513***	-0.531***	-0.517***	-0.489***
	(0.178)	(0.175)	(0.175)	(0.176)
Aug	-0.802***	-0.792***	-0.793***	-0.768***
	(0.174)	(0.180)	(0.179)	(0.177)
Sep	-0.933***	-0.905***	-0.904***	-0.876***
	(0.244)	(0.254)	(0.254)	(0.252)
Oct	-1.050***	-0.999***	-0.998***	-0.974***
	(0.231)	(0.252)	(0.251)	(0.246)
Nov	-0.936***	-0.874***	-0.876***	-0.857***
	(0.143)	(0.159)	(0.157)	(0.154)
Dec	-0.710***	-0.661***	-0.663***	-0.639***
	(0.139)	(0.148)	(0.148)	(0.145)
Observations	111,711	111,711	111,711	111,711
R-squared	0.135	0.138	0.140	0.152
Number of PlantId	725	725	725	725
All Mos Equal F Stat	91.58	52.99	56.28	69.67
Plant Controls	Y	Y	Y	Y
Time Trend	Y	Y	Y	
State Renewable Use		Y	Y	Y
Weather Controls			Y	Y
Year FEs				Y

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table A.III.4
Model 2 Natural Gas Coefficent Estimates

VARIABLES	(1) Log(MWh)	(2) Log(MWh)	(3) Log(MWh)	(4) Log(MWh)
Jan	0.150	0.211*	0.207*	0.200*
	(0.120)	(0.122)	(0.123)	(0.117)
Feb	0.149	0.189	0.193*	0.174
	(0.112)	(0.114)	(0.114)	(0.109)
Mar	0.225**	0.282***	0.275***	0.270***
	(0.109)	(0.0958)	(0.0977)	(0.0893)
Apr	0.222**	0.251**	0.249**	0.233**
	(0.105)	(0.105)	(0.106)	(0.107)
May	0.323***	0.323***	0.296***	0.287**
	(0.118)	(0.106)	(0.108)	(0.112)
Jun	0.285**	0.293**	0.292**	0.273**
	(0.118)	(0.110)	(0.110)	(0.114)
Jul	0.423***	0.383***	0.365***	0.344**
	(0.129)	(0.132)	(0.134)	(0.138)
Aug	0.224*	0.231*	0.236*	0.213
	(0.122)	(0.125)	(0.126)	(0.132)
Sep	0.281***	0.328***	0.325***	0.344***
	(0.104)	(0.101)	(0.101)	(0.105)
Oct	0.212**	0.301***	0.289***	0.326***
	(0.104)	(0.0984)	(0.0982)	(0.100)
Nov	0.271***	0.368***	0.372***	0.404***
	(0.0847)	(0.0829)	(0.0847)	(0.0791)
Dec	0.124	0.214**	0.210**	0.242***
	(0.0925)	(0.0934)	(0.0931)	(0.0884)
Observations	358,071	358,071	358,071	358,071
R-squared	0.048	0.054	0.056	0.077
Number of PlantId	2,818	2,818	2,818	2,818
All Mos Equal F Stat	24.80	19.61	14.14	15.42
Plant Controls	Y	Y	Y	Y
Time Trend	Y	Y	Y	
State Renewable Use		Y	Y	Y
Weather Controls			Y	Y
Year FEs				Y

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table A.IV.3A
Shipment-Level Results: 200 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	0.004***	0.0001**	0.004***
	(0.0001)	(0.0001)	(0.0001)
Plant Cap.	−0.00000***	−0.00000***	−0.00000***
	(0.00000)	(0.00000)	(0.00000)
Mine Count	0.001***	−0.0001***	0.0005***
	(0.00001)	(0.00000)	(0.00001)
Mine Prod.	−0.000***	0.000***	−0.000***
	(0.000)	(0.000)	(0.000)
Observations	500,990	500,990	500,990
R^2	0.815	0.990	0.978
Residual Std. Error (df = 500814)	0.877	0.468	0.751
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.3B
Shipment-Level Results: 100 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	0.007***	0.0003**	0.008***
	(0.0003)	(0.0002)	(0.0002)
Plant Cap.	−0.00001***	0.00000**	−0.00001***
	(0.00000)	(0.00000)	(0.00000)
Mine Count	0.001***	−0.0001***	0.001***
	(0.00001)	(0.00000)	(0.00001)
Mine Prod.	−0.000***	0.000***	−0.000***
	(0.000)	(0.000)	(0.000)
Observations	500,990	500,990	500,990
R^2	0.816	0.990	0.978
Residual Std. Error (df = 500814)	0.874	0.468	0.748
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.3C
Shipment-Level Results: 50 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	−0.007***	0.001***	−0.006***
	(0.001)	(0.0003)	(0.0004)
Plant Cap.	−0.00001***	0.00000***	−0.00001***
	(0.00000)	(0.00000)	(0.00000)
Mine Count	0.001***	−0.00002***	0.001***
	(0.00001)	(0.00001)	(0.00001)
Mine Prod.	−0.000***	0.000***	−0.000***
	(0.000)	(0.000)	(0.000)
Observations	500,990	500,990	500,990
R^2	0.816	0.990	0.978
Residual Std. Error (df = 500814)	0.873	0.468	0.747
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.3D
Shipment-Level Results: 20 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	0.045***	−0.001**	0.044***
	(0.001)	(0.0004)	(0.001)
Plant Cap.	−0.0001***	0.00001***	−0.0001***
	(0.00000)	(0.00000)	(0.00000)
Mine Count	0.002***	0.0001***	0.002***
	(0.00003)	(0.00002)	(0.00003)
Mine Prod.	−0.000***	0.000***	−0.000***
	(0.000)	(0.000)	(0.000)
Observations	500,990	500,990	500,990
R^2	0.818	0.990	0.978
Residual Std. Error (df = 500814)	0.869	0.468	0.744
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.3E
Shipment-Level Results: 10 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	-0.015^{***}	0.003^{***}	-0.012^{***}
	(0.001)	(0.001)	(0.001)
Plant Cap.	-0.0001^{***}	0.00002^{***}	-0.0001^{***}
	(0.00000)	(0.00000)	(0.00000)
Mine Count	0.003^{***}	-0.0001^{*}	0.002^{***}
	(0.0001)	(0.00004)	(0.0001)
Mine Prod.	-0.000^{***}	0.000^{***}	-0.000^{***}
	(0.000)	(0.000)	(0.000)
Observations	500,990	500,990	500,990
R^2	0.819	0.990	0.978
Residual Std. Error (df = 500814)	0.867	0.468	0.743
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.4A
Annual-Level Results: 200 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	0.001^{***}	0.0002	-0.002^{**}
	(0.0003)	(0.0002)	(0.001)
Plant Cap.	0.00000	−0.00000	-0.00000^{***}
	(0.00000)	(0.00000)	(0.00000)
Mine Count	0.0005^{***}	-0.0001^{***}	−0.00005
	(0.00003)	(0.00001)	(0.0001)
Mine Prod	-0.000^{***}	0.000^{**}	0.000^{*}
	(0.000)	(0.000)	(0.000)
Observations	34,987	34,987	34,987
R^2	0.759	0.970	0.328
Residual Std. Error (df = 34811)	0.408	0.186	1.242
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.4B
Annual-Level Results: 100 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	0.001	−0.0001	−0.004
	(0.001)	(0.0004)	(0.002)
Plant Cap.	−0.00001***	0.00000***	0.00001***
	(0.00000)	(0.00000)	(0.00000)
Mine Count	0.001***	−0.0001***	−0.0002**
	(0.00002)	(0.00001)	(0.0001)
Mine Prod.	−0.000***	0.000***	0.000
	(0.000)	(0.000)	(0.000)
Observations	34,987	34,987	34,987
R^2	0.761	0.970	0.328
Residual Std. Error (df = 34811)	0.406	0.186	1.243
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.4C
Annual-Level Results: 50 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	−0.017***	−0.001	−0.011**
	(0.001)	(0.001)	(0.004)
Plant Cap.	−0.00000***	0.00000***	0.00003***
	(0.00000)	(0.00000)	(0.00000)
Mine Count	0.001***	−0.0001***	−0.0001
	(0.00003)	(0.00001)	(0.0001)
Mine Prod.	−0.000***	0.000***	−0.000***
	(0.000)	(0.000)	(0.000)
Observations	34,987	34,987	34,987
R^2	0.763	0.970	0.328
Residual Std. Error (df = 34811)	0.404	0.186	1.242
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.4D
Annual-Level Results: 20 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	−0.018***	−0.003***	−0.033***
	(0.002)	(0.001)	(0.007)
Plant Cap.	−0.00002***	0.00001***	0.0001***
	(0.00000)	(0.00000)	(0.00001)
Mine Count	0.001***	−0.0001	−0.001***
	(0.0001)	(0.00004)	(0.0003)
Mine Prod.	−0.000***	0.000*	−0.000
	(0.000)	(0.000)	(0.000)
Observations	34,987	34,987	34,987
R^2	0.760	0.970	0.333
Residual Std. Error (df = 34811)	0.407	0.186	1.238
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.IV.4E
Annual-Level Results: 10 mi radius

	Dependent variable:		
	Log(Rate)	Log(Weight)	Log(Revenue)
	(1)	(2)	(3)
Plant Count	−0.005*	−0.002	−0.030***
	(0.003)	(0.001)	(0.009)
Plant Cap.	−0.00004***	0.00002***	0.0002***
	(0.00000)	(0.00000)	(0.00001)
Mine Count	0.002***	0.0001	0.0004
	(0.0002)	(0.0001)	(0.001)
Mine Prod.	−0.000***	0.000	−0.000***
	(0.000)	(0.000)	(0.000)
Observations	34,987	34,987	34,987
R^2	0.760	0.970	0.338
Residual Std. Error (df = 34811)	0.407	0.186	1.233
Note:			*p<0.1; **p<0.05; ***p<0.01

REFERENCES CITED

Ackerman, B. A. and Hassler, W. T. (1980). Beyond the new deal: Coal and the clean air act. *The Yale Law Journal*, 89(8):1466–1571.

Angrist, J. D., Pischke, J.-S., and Pischke, J.-S. (2013). *Mostly harmless econometrics: an empiricists companion*. Cram101 Publishing.

Arrow, K. J. (1962). The economic implications of learning by doing. *The review of economic studies*, 29(3):155–173.

Association of American Railroads (2019). Americas freight railroads under balanced regulation. Technical report, Association of American Railroads.

Bain, D. H. (2000). *Empire express: Building the first transcontinental railroad*. Penguin.

Baldwin, R. (1988). Hysteresis in import prices: the beachhead effect.

Baldwin, R. and Krugman, P. (1989). Persistent trade effects of large exchange rate shocks. *The Quarterly Journal of Economics*, 104(4):635–654.

Bergek, A., Berggren, C., Group, K. R., et al. (2014). The impact of environmental policy instruments on innovation: A review of energy and automotive industry studies. *Ecological Economics*, 106:112–123.

Bernard, A. B. and Jensen, J. B. (2004). Why some firms export. *Review of economics and Statistics*, 86(2):561–569.

Bichescu, B. and Raturi, A. (2015). The antecedents and consequences of plant closing announcements. *International Journal of Production Economics*, 168:197 – 210.

Blonigen, B. A., Liebman, B., and Wilson, W. W. (2013). Antidumping and production-line exit: The case of the us steel industry. *Review of Industrial Organization*, 42(4):395–413.

Borenstein, S., Bushnell, J., Wolak, F. A., and Zaragoza-Watkins, M. (2015). Expecting the unexpected: Emissions uncertainty and environmental market design. Technical report, National Bureau of Economic Research.

Boyer, K. D. (1981). Equalizing discrimination and cartel pricing in transport rate regulation. *Journal of Political Economy*, 89(2):270–286.

Calel, R. and Dechezlepretre, A. (2016). Environmental policy and directed technological change: evidence from the european carbon market. *Review of economics and statistics*, 98(1):173–191.

California Environmental Protection Agency (2015). Overview of arb emissions trading program.

Climate Policy Initiative (2019). California carbon dashboard.

Cole, M. A. (2004). Trade, the pollution haven hypothesis and the environmental kuznets curve: examining the linkages. *Ecological economics*, 48(1):71–81.

Covert, T., Greenstone, M., and Knittel, C. R. (2016). Will we ever stop using fossil fuels? *Journal of Economic Perspectives*, 30(1):117–38.

Dennis, S. M. (1999). Using spatial equilibrium models to analyze transportation rates: An application to steam coal in the united states. *Transportation Research Part E: Logistics and Transportation Review*, 35(3):145–154.

Dilts, J. D. (1996). *The Great Road: The Building of the Baltimore and Ohio, the Nation's First Railroad, 1828-1853*. Stanford University Press.

Dixit, A. (1989). Entry and exit decisions under uncertainty. *Journal of political Economy*, 97(3):620–638.

Dixit, A. (1992). Investment and hysteresis. *Journal of economic perspectives*, 6(1):107–132.

Dunne, T., Klimek, S. D., and Roberts, M. J. (2005). Exit from regional manufacturing markets: The role of entrant experience. *International Journal of Industrial Organization*, 23(5):399 – 421.

Dunne, T., Roberts, M. J., and Samuelson, L. (1988). Patterns of firm entry and exit in u.s. manufacturing industries. *The RAND Journal of Economics*, 19(4):495–515.

EIA (2019). Coal mining and transportation.

Ellerman, A. D. and Buchner, B. K. (2007). The european union emissions trading scheme: origins, allocation, and early results. *Review of environmental economics and policy*, 1(1):66–87.

Energy Information Administration (1991). *Trends in contract coal transportation.*

Energy Information Administration (2019). How much carbon dioxide is produced when different fuels are burned?

Ericson, R. and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies*, 62(1):53–82.

European Commission (2019). Eu emissions trading system.

Eyer, J. and Kahn, M. E. (2017). Prolonging coals sunset: The causes and consequences of local protectionism for a declining polluting industry. Working Paper 23190, National Bureau of Economic Research.

Eyer, J. and Wichman, C. J. (2018). Does water scarcity shift the electricity generation mix toward fossil fuels? empirical evidence from the united states. *Journal of Environmental Economics and Management*, 87:224–241.

Fell, H. and Kaffine, D. T. (2018). The fall of coal: Joint impacts of fuel prices and renewables on generation and emissions. *American Economic Journal: Economic Policy*, 10(2):90–116.

Fell, H. and Maniloff, P. (2018). Leakage in regional environmental policy: The case of the regional greenhouse gas initiative. *Journal of Environmental Economics and Management*, 87:1–23.

Fischer, P. A., Bitzan, J., Tolliver, D., et al. (2001). Analysis of economies of size and density for short line railroads.

Fowlie, M. (2010). Emissions trading, electricity restructuring, and investment in pollution abatement. *American Economic Review*, 100(3):837–69.

Ghemawat, P. and Nalebuff, B. (1985). Exit. *The RAND Journal of Economics*, 16(2):184–194.

Ghemawat, P. and Nalebuff, B. (1990). The devolution of declining industries*. *The Quarterly Journal of Economics*, 105(1):167–186.

Harrison, A. E. and Eskeland, G. (1997). *Moving to greener pastures? Multinationals and the pollution-haven hypothesis*. The World Bank.

Harrison, R. (1994). Trade and industrial policy for a" declining" industry: The case of the us steel industry. In *Empirical studies of strategic trade policy*, pages 131–156. University of Chicago Press.

Heckman, J. J. (1981). Heterogeneity and state dependence. In *Studies in labor markets*, pages 91–140. University of Chicago Press.

Hirth, L., Ueckerdt, F., Edenhofer, O., et al. (2016). Why wind is not coal: on the economics of electricity generation. *The Energy Journal*, 37(3):1–27.

Hitzman, M. (2002). R&d in a declining industry (mining): support for the development of revolutionary technologies? *Technology in Society*, 24(1):63–68.

Holladay, J. S. and LaRiviere, J. (2017). The impact of cheap natural gas on marginal emissions from electricity generation and implications for energy policy. *Journal of Environmental Economics and Management*, 85:205–227.

Jenner, S. and Lamadrid, A. J. (2013). Shale gas vs. coal: Policy implications from environmental impact comparisons of shale gas, conventional gas, and coal on air, water, and land in the united states. *Energy Policy*, 53:442–453.

Jordan, B., Lange, I., Linn, J., et al. (2018). Coal demand, market forces, and us coal mine closures. Technical report, Research for the Future.

Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica: Journal of the Econometric Society*, pages 649–670.

Kim, M.-K. and Kim, T. (2016). Estimating impact of regional greenhouse gas initiative on coal to gas switching using synthetic control methods. *Energy Economics*, 59:328–335.

Knittel, C., Metaxoglou, K., and Trindade, A. (2016). 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*, 32(2):241–259.

Knittel, C. R., Metaxoglou, K., and Trindade, A. (2015). Natural gas prices and coal displacement: Evidence from electricity markets. Working Paper 21627, National Bureau of Economic Research.

Lieberman, M. B. (1990). Exit from declining industries:" shakeout" or" stakeout"? *The RAND Journal of Economics*, pages 538–554.

Linn, J. and McCormack, K. (2019). The roles of energy markets and environmental regulation in reducing coal-fired plant profits and electricity sector emissions. *The RAND Journal of Economics*, 50(4):733–767.

MacDonald, J. M. (1989). *Effects of Railroad Deregulation on Grain Transportation*. US Department of Agriculture, Economic Research Service.

Máñez, J. A., Rochina-Barrachina, M. E., and Sanchis, J. A. (2008). Sunk costs hysteresis in spanish manufacturing exports. *Review of World Economics*, 144(2):272–294.

Meyer, D. W. and Taylor, C. T. (2015). The determinants of plant exit: The evolution of the us refining industry. *FTC Working Paper Series*.

Miller, K. S. and Wilson, W. W. (2017). Governance structure and exit: Evidence from california hospitals. *Review of Industrial Organization*.

Moel, A. and Tufano, P. (2002). When are real options exercised? an empirical study of mine closings. *The Review of Financial Studies*, 15(1):35–64.

Moody, J. (1921). *The railroad builders: A chronicle of the welding of the states*, volume 38. Yale University Press.

Moyer, R. (1965). Requirements contracts and energy transportation developments in coal. *Southern Economic Journal*, pages 331–341.

Murray, B. C. and Maniloff, P. T. (2015). Why have greenhouse emissions in rggi states declined? an econometric attribution to economic, energy market, and policy factors. *Energy Economics*, 51:581–589.

National Archives (2020). Interstate commerce act (1887).

Popp, D. (2003). Pollution control innovations and the clean air act of 1990. *Journal of Policy Analysis and Management*, 22(4):641–660.

Reynolds, S. S. (1988). Plant closings and exit behaviour in declining industries. *Economica*, 55(220):493–503.

RGGI-Inc. (2018). The regional greenhouse gas initiative.

Roberts, M. J. and Tybout, J. R. (1997). The decision to export in colombia: an empirical model of entry with sunk costs. *The American Economic Review*, pages 545–564.

Ruth, M., Gabriel, S. A., Palmer, K. L., Burtraw, D., Paul, A., Chen, Y., Hobbs, B. F., Irani, D., Michael, J., Ross, K. M., et al. (2008). Economic and energy impacts from participation in the regional greenhouse gas initiative: A case study of the state of maryland. *Energy Policy*, 36(6):2279–2289.

Sandbag Climate Initiative (2019). Carbon price viewer.

Schmalensee, R., Joskow, P. L., Ellerman, A. D., Montero, J. P., and Bailey, E. M. (1998). An interim evaluation of sulfur dioxide emissions trading. *Journal of Economic Perspectives*, 12(3):53–68.

Schmalensee, R. L. and Wilson, W. W. (2016). Modernizing us freight rail regulation. *Review of Industrial Organization*, 49(2):133–159.

Slade, M. E. and Thille, H. (2009). Whither hotelling: tests of the theory of exhaustible resources. *Annu. Rev. Resour. Econ.*, 1(1):239–260.

Stover, J. F. (2008). *American railroads*. University of Chicago Press.

Tietenberg, T. (2010). Cap-and-trade: the evolution of an economic idea. *Agricultural and Resource Economics Review*, 39(3):359–367.

Whinston, M. D. (1988). Exit with multiplant firms. *The RAND Journal of Economics*, pages 568–588.

Wilson, W. W. and Wolak, F. A. (2016). Freight rail costing and regulation: The uniform rail costing system. *Review of Industrial Organization*, 49(2):229–261.

Wilson, W. W. and Wolak, F. A. (2018). Benchmark regulation of multiproduct firms: An application to the rail industry. Working Paper 25268, National Bureau of Economic Research.

Winston, C. (2005). *The success of the Staggers rail act of 1980*. AEI-Brookings Joint Center for Regulatory Studies.

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