Essays in Climate Policy and Innovation

Dissertation Advisor:
Professor Joseph E. Aldy

Author:
Megan R. Bailey

Essays in Climate Policy and Innovation

Abstract

This set of studies examines several critical issues related to climate policy and innovation. Innovation is an important facet of environmental policy, generally, as it generates new and improved technologies that bring down the cost of correcting negative externalities related to the environment. For addressing climate change, economists recognize that carbon pricing provides the dual advantage of being the most cost-effective policy option for the near term as well as the future: carbon pricing coordinates emissions reductions among firms so that those with the lowest cost of abatement take action, and it incentivizes firms to innovate to bring down the cost of abatement over the long-run.

In the first chapter of this work, "US Carbon Pricing and Coal Power Plant Productivity", I find that coal power plants in the US Regional Greenhouse Gas Initiative (RGGI) have become significantly less efficient since the beginning of the cap-and-trade program, reflecting a deficit of innovative improvements to plants expected by academic and industry sources. Using econometric decomposition and harnessing the data on the US electricity sector in a new way to craft important operational variables, I find that RGGI coal plant owners accepted the efficiency penalty that mechanically comes with operating at low output rates. This is novel microeconomic evidence that firms may substitute non-innovative for innovative adjustments to an input price shock. I also provide evidence that RGGI coal plant owners have not made investments in efficiency-improving capital as expected; analytical results suggest this may be due to plant owners foreseeing a decline in production. These results are important for climate policy for the US electricity sector. I find a 4.56% carbon emissions rebound effect among RGGI plants in the first four years of the program: though they reduced output and

associated carbon emissions, their decline in efficiency constrained emissions reductions. This rebound effect is small but should be included in cost-benefit analyses of any carbon taxes proposed for the electricity sector. My results suggest that investment behavior will be unique among firms in a declining industry; as coal-fired power is phased out in the US, we can expect to see coal plant owners forego previously profitably efficiency upgrades. Finally, the current US climate policy specifically targets the efficiency of coal-fired power plants. These results show that this is a particularly costly way to address carbon emissions in the US.

In my second chapter, "What do we lose by picking winners? The role of technology-specific clean energy incentives in induced innovation", I examine a common policy combination, technology-specific renewable energy subsidies and carbon pricing. Do such subsidies, when added to a carbon price, incur dynamic social costs in the form of re-directing innovation away from the suite of "optimal", cost-effective technology we would anticipate under just a carbon price? This chapter answers this question with an analytical model and a simulation calibrated to the EU electricity sector. I look at two important cases: the exogenous carbon price (carbon tax) case and the endogenous carbon price (cap-and-trade) case. Results from the analytical model indicate that innovation can follow production changes incentivized by clean energy subsidies. This points to subsidies having the potential to re-direct innovation away from unsubsidized and toward subsidized technologies. When the carbon price is endogenous, as in a cap-and-trade system, innovation directed toward reducing emissions of emissions-intensive inputs (such as coal and natural gas) falls unambiguously, due to the sensitivity of the carbon price to emissions abatement achieved by subsidized renewables. Simulation results point to subsidies reducing the incentive for innovation devoted to non-subsidized technology under high carbon taxes ($30-40$ per ton of CO2). Results are nonlinear in the size of the subsidy and carbon tax, highlighting the roles of both input price and input quantity changes in inducing innovation. When subsidies are used in combination with an emissions cap, the incentive for innovation for all technologies is depressed. Taken together, results point to subsidizing specific technologies as a poor way to guide innovation for clean technology when a carbon pricing scheme is present. Instead of overlapping specific subsidies for renewable

technology on top of carbon pricing schemes, policymakers should use carbon pricing only, adding additional policies if evidence suggests specific market failures. In the third chapter of this work, "Induced innovation, market size, and total value to firms: The case of US electricity", I highlight the importance of an emerging hypothesis in the induced innovation literature, the "market size" effect. This hypothesis holds that firms should innovate to reduce the cost of the input on which they rely most for production. Innovation incentives are often discussed in terms of relative input prices; input price and input quantity (or market size) effects should both play a role in incentivizing innovation but run in opposite directions. I discuss how a firm's elasticity of substitution is a critical parameter that determines which effect dominates. Empirical estimation of induced innovation effects from input price and quantity changes should be guided by the econometrician's knowledge of elasticities of substitution or variables that capture total value changes to firms in the face of input shocks. I propose that the US electricity sector provides a ripe empirical setting for testing the impact of market size on innovation, as the sudden availability of cheap natural gas starting in 2008 accelerated the sector's switch in reliance on natural gas for production. Understanding the role of market size in tailoring innovation has special importance for climate policy. Stringent policy and a subsequent increase in the economy's reliance on less emissions-intensive inputs will impact the direction of innovation, in addition to any policy-led input price changes.

Contents

Acknowledgments

This work would not be possible without the dedicated investment of my dissertation committee: Joe Aldy, Bill Hogan, and Jim Stock. It has been a joy to learn from these scholars and policy professionals. Joe, who served as my committee Chair, has provided above-par mentorship throughout my graduate studies at Harvard.

A very large host of faculty, students, and other scholars have provided me with guidance for this dissertation. In particular, Rob Stavins served as my primary academic advisor for years. Cuicui Chen, Todd Gerarden, A. Patrick Behrer, and Shefali Khanna have generously given ongoing advice for this work. A larger circle of peers and office mates have helped me, day-to-day, shape my thinking and choices. I received valuable feedback from attendees while presenting versions of the chapters in this work in the Harvard Environmental Economics Lunch, the Interdisciplinary PhD Workshop in Sustainable Development at Columbia University, the 2019 Northeast Workshop on Energy Policy and Environmental Economics, the 2019 Harvard Environmental Economics Program Research Workshop, the 2019 Seminar in Environmental Economics and Policy at the Harvard Kennedy School, and the 2020 NBER Productivity Seminar.

For Chapter 1, I am indebted to personnel at the Energy Information Administration (EIA) and the Federal Energy Regulatory Commission (FERC), an anonymous energy company, Sam Korellis at the Electric Power Research Institute, and Una Nowling at Black & Veatch for useful discussions about this work and coal power plant engineering. For Chapter 2, I am grateful Matthew Wittenstein for provision of IEA et al. (2015) data in Excel format.

The National Science Foundation Graduate Research Fellowship made it possible for me to pursue a PhD at Harvard. The Joseph Crump Fellowship provided funding for Chapter 2. The Vicki Norberg-Bohm Fellowship supported my early-stage dissertation research, which involved interviewing stakeholders in the EU Emissions Trading System (EU ETS), including participating firms, regional NGOs and interest groups, and an EU official. I thank all interviewees, who shared their time and thoughts on the EU ETS under the condition of

anonymity.

Completing a dissertation is, of course, a particularly challenging endeavor, and I have been personally supported and intellectually sharpened by many colleagues, friends, and family members during the process. Mark and Annetta Shepard have been close mentors. Bill Clark has helped me discover the rigor and importance of multidisciplinary and interdisciplinary study of environmental issues. My students have taught and inspired me more than they ever suspected. I have been privileged to work and live among stimulating peers in the Harvard Kennedy School, Economics Department, and Business School PhD programs as well as sister programs at MIT. My office mates have been especially encouraging. I have been blessed with a coterie of friends and family members with whom I've been able to share the highest and lowest moments of the PhD journey. You know who you are—and you have been invaluable in me arriving at this moment.

My life companion, Brad, has continually defined the word "support" through the many ways he has offered it over the years.

I have a unique family. At home there is always a science experiment, technical project, or creative enterprise underway. My parents gave me educational and personal opportunities they never had; fostered a culture of discovery, discipline, and concern over all life; and let me go my own way. This dissertation and my completion of a PhD in Public Policy is much to their credit.

I dedicate this work to present and future generations—especially the global poor—who experience the impacts of climate change.

Chapter 1

US Carbon Pricing and Coal Power Plant Productivity

1.1 Introduction

Induced innovation theory suggests that producers will respond to input price changes by directing innovative effort toward reducing the cost of the new, relatively expensive input (Hicks, 1932; Acemoglu, 2003). Therefore, a carbon price should lead firms to direct innovative effort toward reducing the cost of using carbon-intensive inputs.[1] Innovation can take the form of firms inventing and crafting new processes or technology themselves; it can also take the form of firms adopting processes or technology others' have created (Milliman and Prince, 1989). Such innovation is important from a welfare standpoint, as it can bring down the cost of emissions abatement over time. In the United States, there is one carbon pricing scheme, the Regional Greenhouse Gas Initiative (RGGI), that applies to electricity producers that use the most carbon-intensive input: coal. Has RGGI led to innovative changes at coal plants that participate in the program?

For an owner of a coal-fired power plant, improving the efficiency of fuel use is potentially the most cost-effective means of complying with a carbon price and a reflection of innovative

[1] Conditional on input substitution and output adjustments (Acemoglu, 2003).

effort. Previous academic literature suggests that coal plants would respond to a carbon price with efficiency improvements (Linn et al., 2014).[2] Industry publications promote efficiency improvements as a primary mode of compliance with a carbon price. "Heat rate [efficiency] improvement is...the first obvious step to reduce carbon dioxide (CO2) and all other emissions" (Korellis, 2014, p. 1). Efficiency improvements reflect unique action that coal plant operators can take on-the-ground (Bushnell and Wolfram, 2009) as well as technology that plant owners can install (Linn et al., 2014; Nowling, 2015).

Despite expectations that coal plants in RGGI would become more efficient in response to the program, trends in RGGI coal plant heat rates show that RGGI plants have become *less* efficient over time since the beginning of the program (Figure 1.1).[3] Trends in the distributions of RGGI and non-RGGI coal plant heat rates (Figure 1.2) reveal that RGGI plants may have become less efficient compared to other US coal plants during the RGGI program. What could be behind this counterintuitive trend?

I use a difference-in-differences empirical design to estimate the relationship between RGGI participation and coal plant efficiency, conditional on a few common culprits of coal plant efficiency decline. I find that coal-fired power plants in RGGI indeed exhibit a significant decline in efficiency after the start of RGGI, compared to coal plants in the rest of the US. To investigate the mechanism behind this result, I first use a model of coal plant operation and investment to demonstrate that it can be optimal for producers to forego otherwise profitable innovation if they 1) can respond to an input price change along an alternative margin (in this context, through reducing output instead of improving efficiency) or 2) anticipate a future decline in output. I then decompose the change in RGGI plants' efficiency during the

[2]The "carbon price" generated from the RGGI program is a carbon dioxide (CO2) emissions price. As I discuss below in Section 1.2, given the CO2 emissions abatement options for coal-fired power plants in RGGI, the emissions price is effectively a coal tax. Linn et al. (2014) identify the elasticity of US coal-fired unit efficiency with respect to coal price changes and use their estimates to predict efficiency changes that would be the result of CO2 taxation.

[3]The industry standard for quantifying efficiency at power plants is through the heat rate, a measure of the amount of heat energy required to produce a unit of electricity. A heat rate is an inverse measure of efficiency, so an efficiency decline is captured as a heat rate increase. I use the units mmBtu/MWh (million British thermal units per megawatt-hour) for all heat rate values in this paper. A megawatt hour is the amount of electricity produced by an electricity generator with a capacity of one megawatt (MW) running for one hour.

Figure 1.1: *RGGI coal plant net heat rates over time*

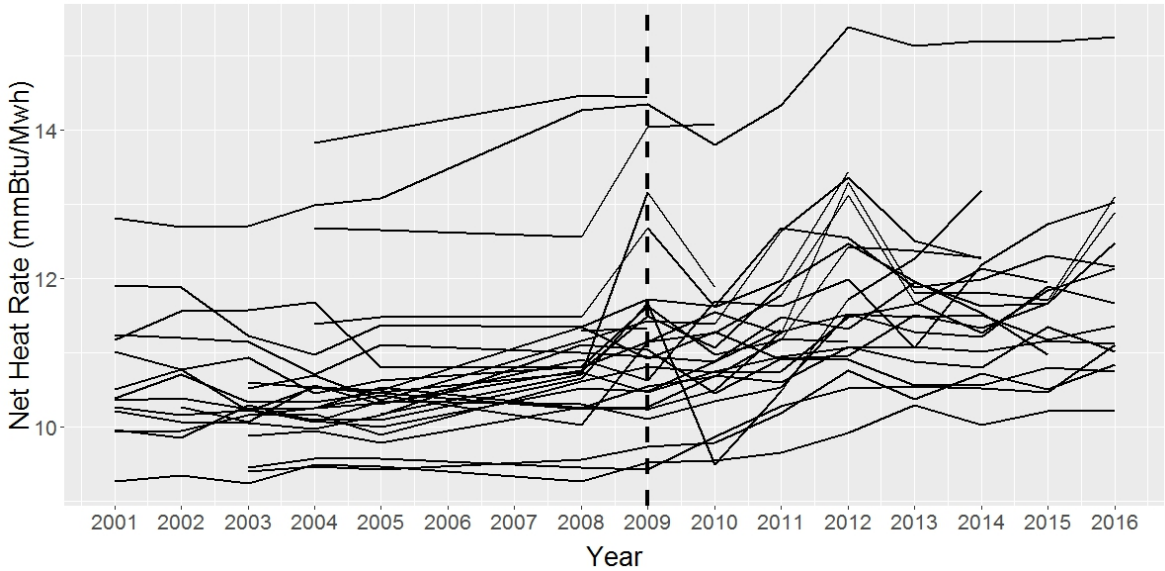

Note: Each line depicts a RGGI coal plants' average annual net heat rate time series, as calculated from the EIA monthly data. The dashed vertical line at 2009 demarcates the beginning of the RGGI program.

Figure 1.2: *The distributions of monthly net heat rates of coal-fired power plants by program participation*

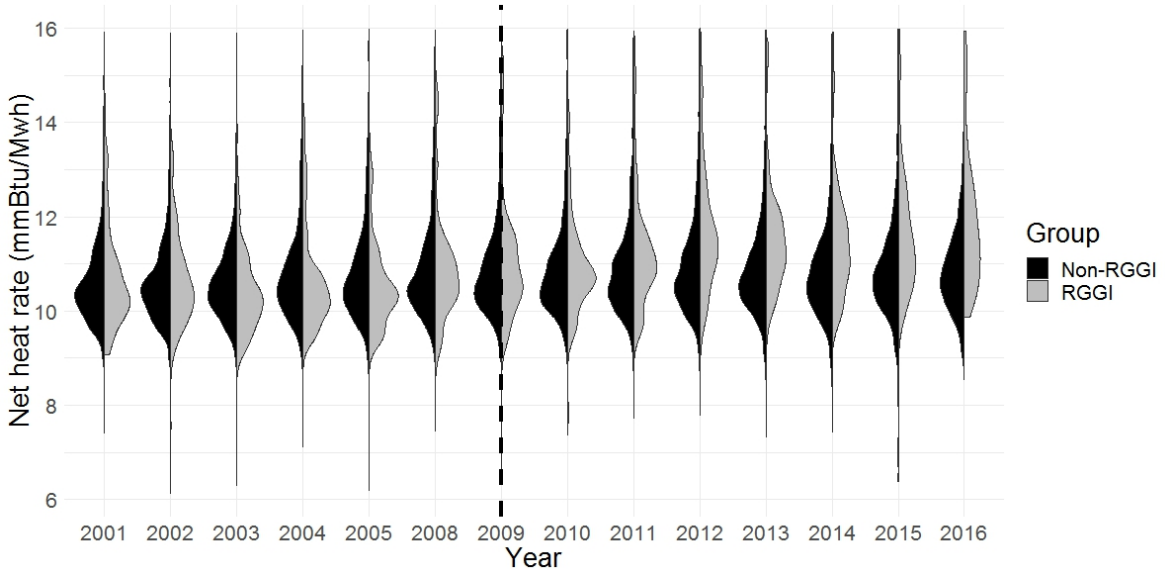

Note: The dashed vertical line at 2009 demarcates the beginning of the RGGI program. Annual distributions are constructed from monthly net heat rate values, from EIA data.

RGGI program into its mechanical sources at a plant level: gross efficiency (the conversion of coal into electricity) and auxiliary power consumption (the usage of electricity on-site). I estimate models of each to determine the roles of output shifts, impending retirement, and other important factors such as compliance with other environmental policies in the RGGI coal plant efficiency decline. I take this estimation approach, informed by the engineering of coal plants, to best identify the sources of inefficiency at coal plants and exploit the detailed data publicly-available on the US electricity sector. To my knowledge, this is the first paper to compute and examine power plant auxiliary power consumption for policy analysis.

I find that the greatest share of RGGI coal-fired power plants' efficiency loss is caused by declining output, or electricity generation, at RGGI plants. During the time period of the RGGI program (after 2009), there has been a dramatic US-wide decline in the share of electricity produced from coal, primarily attributable to the availability of cheap natural gas (a competing fuel source) (Coglianese et al., 2018; Linn and McCormack, 2019). The cost of carbon emissions to producers in the RGGI program further drives a wedge between the costs of producing electricity from coal and natural gas, exacerbating coal's loss of market share in the RGGI region. Fell and Maniloff (2018) find that RGGI coal plants reduced their generation by 10% in the post-RGGI period, due to the program. Anticipated induced innovation—even as estimated in prior academic literature or indicated by industry publications—may not occur among producers exposed to an input price change if they are able to respond along a number of margins.

I also find that, conditional on output changes, RGGI coal plants are not improving efficiency on par with historical precedent for coal plants' responses to input price changes (Linn et al., 2014). This suggests that they are no longer making investments in efficiency-improving capital. Given my second theoretical model result, this may be due to coal plants anticipating a future decline in production. This provides suggestive evidence of the "market share" effect: producers will direct innovative activity toward abundant factors (Acemoglu, 2003).

These results are important for at least two additional reasons. First, they provide an

estimate of a type of "rebound" effect of carbon dioxide (CO2) emissions that we can anticipate with reduced output from coal plants. Based on estimates from Fell and Maniloff (2018), if coal plant efficiency had stayed constant between the pre-program and post-program time periods, RGGI would have been responsible for a corresponding 10% decline in CO2 emissions. However, the implied efficiency decline from this generation shift led to an emissions rebound effect of 4.56%, leading to only a 9.54% reduction in CO2 emissions. Given prior estimates of coal plants' efficiency improvements in response to input price changes (Linn et al., 2014), this rebound effect could be as large as 5.75%. As carbon pricing is expanded and/or coal continues to decline as a share of electricity production, we may continue to see such emission rebound effects. Given the fact that natural gas plants' efficiency is even more sensitive to output shifts than coal plants' (Lew et al., 2012), a phase-out of natural gas through a high carbon price or increase in renewables on the electricity grid may lead to a rebound effect among natural gas plants as well.

Second, the set of plants subject to this first-best climate policy are exhibiting an efficiency trend that is opposite to not only expectations but also the contemporary proposed federal climate policy, the Affordable Clean Energy (ACE) Rule. The ACE Rule allows US states to voluntarily create efficiency improvement policy for their coal-fired power plants. Coal plants in RGGI have not found it cost-effective to improve their efficiencies in the last decade. Theory suggests that the ACE Rule is less cost-effective than a carbon pricing scheme by not allowing abatement costs to be efficiently allocated across plants; results presented in this study suggest that the rule additionally would mandate a form of abatement among all plants that would not counterfactually be their cost-effective form of compliance under a carbon price.

This paper contributes to the empirical literature on induced innovation, particularly the growing literature on induced innovation from pricing environmental externalities. Broadly, this literature examines the role of changing input prices (as a proxy for externalities pricing) or of direct externalities pricing on innovation outcomes.[4] As predicted by theory for macro-level

[4]Some of the foundational work in this area includes Newell et al. (1999), Popp (2002), and Jaffe and Stavins (1995). More recent empirical work includes Noailly and Smeets (2015), Aghion et al. (2016), and Calel and Dechezlepretre (2016).

(Hicks, 1932) and micro-level (Milliman and Prince, 1989) responses, most studies find that these price changes induce significant, positive changes in invention, innovation (development of an idea into a process or product), or diffusion of energy-saving or pollution abatement technology.[5] This paper is most closely related to work that captures changes in process improvements or diffusion of technology that contributes to an efficiency or emissions-rate change. Linn (2008), for example, found that owners of coal-fired power plants subject to the Nitrogen Oxides (NOx) Budget Trading Program, a cap-and-trade program for NOx emissions in the eastern US, significantly reduced their plants' NOx emissions rates during the program. I bridge this literature with the larger empirical literature on innovation and technological change, which provides empirical examples of role of market share. For example, Bustos (2011) finds that Argentinian firms that gained market share in the export market after a regional trade liberalization exhibited significant increases in spending on technology upgrades. Acemoglu and Linn (2004) document that the potential market share for new drugs is significantly associated with entry of several measures of innovative new drugs.

I build on and contribute to the literature that examines how owners of coal and natural gas power plants respond to incentives. Understanding firm behavior in this industry is critical for good environmental policy: the US electricity sector is a major contributor to CO2 emissions and local pollutants.[6] This paper is similar to Linn et al. (2014), in which the authors estimate US coal plants' efficiency response to changes in delivered coal prices and find that a $10 per ton CO2 price should induce a 0.6% improvement in plant-level efficiency, on average. Knittel (2002) documents the efficiency improvements of fossil plants in response to a suite of state-level incentives set up through cost-of-service regulation. Post industry restructuring, US coal plants that were divested from former utility ownership improved their efficiency (Chan et al., 2017) and re-negotiated long-term coal contracts to obtain a lower cost of fuel (Cicala, 2015). Preonas (2019) finds that mark-ups from coal suppliers respond to the relative natural

[5] I refer here to Schumpeter's trichotomy as a classification of the manifestation of "innovation" in an economy (Schumpeter, 1939).

[6] In 2018, the electricity sector was responsible for about 33% of US CO2 emissions. Source: https://www.eia.gov/tools/faqs/faq.php?id=77&t=11

gas to coal price, indicating that a carbon price, which imposes a relatively higher price on coal than natural gas plants, can be offset via upstream pass-through. Given the evidence confirming that coal plants have historically responded to efficiency and other incentives in line with theory, my results are even more surprising—and point to the need for nuanced theory to capture the specific tradeoffs that producers face in this industry, in a time of declining market share for coal plants.

This paper adds empirical evidence to the literature that asks: What would happen to emissions if the US adopted a carbon price applicable to the electricity sector? Recent work focused on emissions changes from intensive-margin generation shifts illustrates that potential emissions reductions achieved from carbon pricing are dependent on the availability of cheap natural gas (Cullen and Mansur, 2017) and built natural gas capacity (Knittel et al., 2015). Similarly, the potential of renewable energy to reduce emissions depends on low natural gas prices (Fell and Kaffine, 2018) and is regionally heterogeneous (Holladay and Lariviere, 2017). I find that emissions reductions, which are primarily driven by the market share of coal declining relative to natural gas, comes with a rebound effect when coal plants are operating at historically unprecedented low output rates.[7]

In the section that follows, I detail the context in which coal-fired power plants in RGGI are operating as well as the determinants of plant-level efficiency. I follow this discussion with a model of coal plant operation in Section 1.3. In Sections 1.4-1.7, I present my empirical strategy, a description of my data, and my results. Section 1.8 includes robustness checks. I conclude in Section 1.9.

[7]Holland et al. (2018) find that the shift in generation from coal to natural gas is one of the two major contributing factors in the decline in damages from electricity generation emissions, 2010-2017. EIA (2017) shows that the decline in emissions intensity of electricity from fossil sources—which captures the shift in production from a reliance on coal to natural gas—has contributed more to decline in carbon dioxide emissions from electricity generation each year since 2005 than the increased reliance on zero-emissions sources. See the Sept. 2018 report on U.S. Energy-Related Carbon Dioxide Emissions, 2017, from the EIA at https://www.eia.gov/environment/emissions/carbon/pdf/2017_co2analysis.pdf

1.2 Operating a coal plant: market and engineering details

1.2.1 Market context

Coal-fired power plants in the United States operate in a complex market that has made them less competitive in the last decade. The share of coal in electricity production declined from 48% in 2008 to 27% in 2018.[8] The price of natural gas, a competing fuel source, delivered to the electricity generation industry plummeted from about $9 per million British thermal units (mmBtu) to $5.50 per mmBtu (2010 USD) between 2008 and 2009 and has remained relatively cheap since, thanks to the large-scale deployment of fracking technology to remove previously unavailable natural gas from shale deposits. (See Figure 1.3.) Over the same time frame, demand for electricity has fallen, in tandem with the Great Recession. Around 2007, the use of renewables for electricity generation began to increase and has been increasing steadily in many of regions. Additionally, coal plants must comply with a suite of federal and state environmental policies. Industry literature raised concern about about the tight margins that owners of many types of power plants faced in response to changing natural gas prices and demand as early as 2010 (Filsinger, 2010). Linn and McCormack (2019) find that the change in the natural gas prices, demand shifts, and increased wind generation reduced both output and profits of coal plants in the eastern US, 2005-2015. Coglianese et al. (2018) find that 92% of the total decline in coal production in the US between 2008 and 2016 was due to the change in the relative price of natural gas to coal impacting coal's share in electricity generation. Both papers find a relatively small role for compliance with environmental regulations in impacting coal plants' generation.

1.2.2 The Regional Greenhouse Gas Initiative

In addition to facing the market forces described above that have made coal plants less competitive in the wholesale electricity market over the last decade, owners of all fossil-fired

[8] Author's calculations from EIA data available here: https://www.eia.gov/energyexplained/electricity/electricity-in-the-us.php

Figure 1.3: *Price series of coal and natural gas delivered to RGGI states*

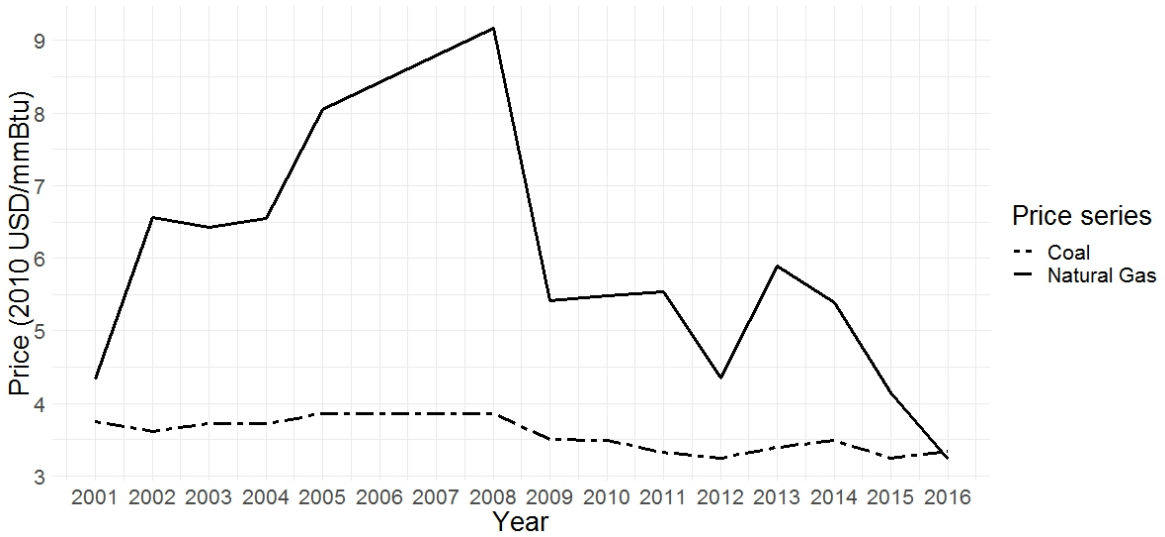

Note: Data is from the EIA SEDS database and converted to 2010 USD per mmBtu.

plants in the US Northeast and Mid-Atlantic states participating in the Regional Greenhouse Gas Initiative (RGGI) have been required, since the beginning of the cap-and-trade program in 2009, to purchase and submit carbon allowances for their CO2 emissions.[9] The product of voluntary state participation and coordination, RGGI imposes a regional cap on CO2 emissions from fossil-fired electricity generation sources. Specific participation requirements are set by states, though typically sources of at least 25 MW in size are included. A regional auction of RGGI allowances is held quarterly; fossil plant owners may purchase them from these auctions or through secondary market transactions.

The program includes a reserve price in allowance auctions that is the minimum price of CO2 emissions, when demand for allowances falls short of supply in the regional auction. This reserve price creates a price floor for RGGI allowances. The initial reserve price was set at $1.86 per RGGI allowance (one ton of CO2 emissions) in 2009 and was indexed to Bureau of

[9] All fossil-fired electricity generating plants in participating states are covered by RGGI. Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New York, Rhode Island, and Vermont currently participate. New Jersey participated through 2011, and New Jersey, Pennsylvania, and Virginia are planning to re-join or join the program.

Labor Statistics Consumer Price Index until 2013.[10] After initial demand for allowances in 2009 led to an allowance price of about $3, demand in 2010-2012 was low enough such that the price of allowances was at or just above the price floor. The RGGI program therefore effectively served as a small carbon tax prior to 2013. (See Figure 1.4 for the RGGI allowance price series.) When, in 2013, information that the annual cap for 2014 would be less than half the annual cap for 2009-2011 and would decline 2.5% each year thereafter, the price of allowances began to climb to just over $8/ton of CO2 in 2015. Therefore, for the time period of my sample (2009-2016), the RGGI price was largely exogenous to producers, either being at the program's price floor or being driven by cap changes.

Figure 1.4: *Price series of RGGI allowances*

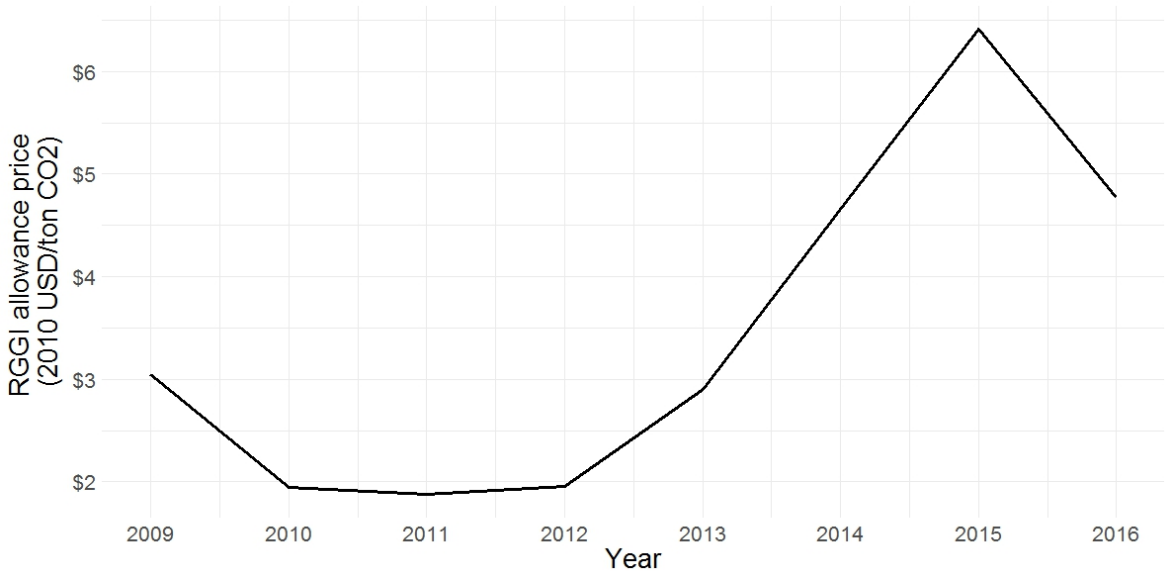

Note: Data is from the RGGI program website and converted to 2010 USD per ton of CO2.

Though the RGGI allowance price was close to the program's price floor during three of the initial years of the program, over the time frame of the RGGI program captured by my data (2009-2016), the RGGI price has been 6-19% of the price of coal delivered to RGGI states, a non-trivial cost for coal plant owners. The price of allowances generated from the program

[10]See the Auction Notice for March 2013 auction, available at: https://www.rggi.org/auctions/auction-materials/auction-materials-archive

penalizes electricity production from coal plants more than that from natural gas plants, due to the difference in carbon intensity between these fuel sources and the efficiency of each type of plant. Over the course of the RGGI program in my data (2009-2016), the average RGGI allowance price translated into an average tax of $3.98 per MWh of electricity produced from coal plants and $1.61 per MWh from natural gas plants.[11]

1.2.3 Coal-fired power plants, emissions, and efficiency

The impact of a carbon emissions price

Coal-fired power plants do not have cost-effective means of reducing CO2 emissions through emissions control equipment (as they can sulfur emissions, for example, with scrubbers). Very few coal plants in RGGI states had natural gas units when the program started. Therefore, to coal plant owners, the RGGI CO2 price is tantamount to a coal tax. To reduce expenditures on emissions allowances, plant owners have the options of becoming more efficient, reducing generation, or retiring.

Conditional on using the same type of coal, the efficiency, or heat rate, of a coal-fired power plant maps one-to-one with its carbon emissions intensity. Heat rate improvements are a primary mode of emissions reduction for coal plants. One industry article aptly notes: "Heat rate improvement is...the first obvious step to reduce carbon dioxide (CO2) and all other emissions" (Korellis, 2014, p. 1). Heat rate improvements are also means by which coal plant owners pursue operational cost savings on an ongoing basis. The cost of heat rate reductions on the order of 0.10% to 3% can typically be recovered in one to five years, with efficiency savings thereafter (Nowling, 2015).

[11]The average RGGI price over the years in my data was $3.43. For this calculation, I use the observed emissions rate and average net heat rate of RGGI coal plants from my data set (from EPA and EIA data sources, respectively), the emissions rate of natural gas reported by the EIA (https://www.eia.gov/electricity/annual/html/epa_a_03.html) and the average heat rate of US natural gas plants, 2009-2016, reported by the EIA (https://www.eia.gov/electricity/annual/html/epa_08_01.html).

Heat rate changes and plant mechanics

Plant-level heat rates can be improved by a number of means. Bushnell and Wolfram (2009) describe a variety of ways that personnel at plants can directly impact the efficiency of plant operations. For example, plant operators can control boiler operations and conditions such as the rate at which pulverized coal is fed into them, the oxygen available for the combustion process, and the timing of soot cleaning. Plant owners can engage in maintenance that contributes to improved efficiency. Repairing steam leaks is one such endeavor. Maintenance of the plant components exposed to extreme heat (for example, the boilers and steam pipes) is particularly important to avoid accidental breakage. There are also investments plant owners can make to improve efficiency. Some of these are forms of equipment to automate the actions just described, such as intelligent sootblowing systems that automatically allocate steam to blow out soot. Others involve upgrading equipment, such as the components of the turbine.[12] In terms of innovation, heat rate improvements include aspects of new ideas and their application that are made on-the-ground by plant personnel as well as adoption of new practices and updated plant equipment.

One important distinction for my analysis is the relationship between a plant's net heat rate, gross heat rate, and auxiliary power consumption. A plant's net heat rate is its total efficiency, the amount of heat input it requires to produce each MWh of electricity it outputs to the electricity grid (net generation). Gross efficiency measures the plant's efficiency in converting fuel into electricity (gross generation), before it uses some of that electricity on site for plant operations (auxiliary power consumption). Many aspects of plant operation, such as its capacity factor and the choices personnel can make at the plant described by Bushnell and Wolfram (2009) impact gross efficiency. Maintenance to plant boilers, turbines, or generators will impact gross efficiency. A plant uses auxiliary power to run fans that assist in plant operation, environmental control equipment, and any other components of plant operation

[12] The Regulatory Impact Analysis of the Affordable Clean Energy Rule describes some of these capital investment options. See https://www.epa.gov/stationary-sources-air-pollution/regulatory-impact-analysis-repeal-clean-power-plan-and-emission.

that require electricity.

The impact of changes in output

Aside from intentional changes that a coal plant owner or operator would make, a coal plant's efficiency has a mechanical relationship with its output. Coal plants are designed to run their units at about 95% of their total output rate, or capacity factor.[13,14] When they do so, they are able to operate at optimal efficiency. As plants run their units at lower capacity factors, their heat rates rise. The relationship is nonlinear and exponential such that running a unit at very low capacity factor leads to a very high heat rate. Therefore, coal plants running less due to the exogenous market forces described above as well as the RGGI program result in direct efficiency loss.

If coal plants fluctuate their output in response to factors such as changing electricity demand and electricity generated from renewables (called "cycling" or "load-following"), they can experience reduced efficiency (EPRI, 2018). Because I do not observe a difference in coal unit cycling over time or between RGGI and non-RGGI coal units, I do not include it in this study.

The impact of environmental control equipment

For my study, US Clean Air Act regulations are important to consider, as compliance choices also impact efficiency. I provide an overview of the major policies with which coal plants have had to comply during my sample period (2001-2016) here and note how compliance impact plant efficiency.

A series of cap-and-trade programs have required fossil-fired electricity units to reduce summertime NOx emissions. The Ozone Transport Commission (OTC) NOx Budget Program was initiated in 1999 for the Northeastern US, and the NOx Budget Trading Program expanded

[13] A typical coal-fired power plant has several units, which consist of a boiler-generator combination.

[14] A plant's capacity factor is its generation divided by its total potential generation in a particular time period. Throughout this paper, I express capacity factor on a 1-100 scale, the percent of the plants' total output that is realized in a given time period.

the program to the Southeast and Midwest US in 2004 (Linn, 2008). Coal-fired power plants can reduce NOx emissions in a number of ways, including installation of Selective Catalytic Reduction (SCR) units, which capture NOx emissions post-combustion, or reducing NOx formation in boilers (Linn, 2008). The former method can incur a small (1% or less) net heat rate penalty on plants by requiring more on-site electricity.

The Clean Air Interstate Rule (CAIR) (2009-2014) and Cross-State Air Pollution Rule (CSAPR) (2015-present) essentially replaced the NOx Budget Trading Program with a more stringent summertime NOx cap-and-trade program. They also have included annual NOx and sulfur dioxide (SO2) caps (Coglianese et al., 2018). Coal-fired power plants can reduce sulfur emissions by installing flue gas desulfurization units (FGD units or "scrubbers"), which remove SO2 from emissions post-combustion or adopt a number of options that have a higher variable cost but much lower fixed cost of capital. Scrubbers can incur a net heat rate penalty of about 1% on plants via, like SCR units, requiring plants to use more electricity on site to run them. Figure 1.5 shows that the share of coal plants in RGGI states that had installed scrubbers increased markedly around the time of the start of the program (2009) and surpassed that of the rest of the US in 2010. I therefore look at the role of scrubbers in RGGI coal plants' heat rates. Another option plants have for reducing sulfur emissions is to switch to using Powder River Basin (PRB) coal from specific mines in Montana and Wyoming, which has a lower sulfur content. Depending on which coal type plants are switching from, this can lead to a net efficiency loss of 2-2.5%. However, I do not see evidence of RGGI coal plants having made such a switch in fuel purchase data.

Finally, the Mercury and Air Toxics Standards (MATS), which began in 2016, requires coal power plants to meet emissions standards for mercury and other toxic air pollutants. While some forms of mercury emissions controls may incur a small heat rate penalty on plants (EPA, 2018), only three RGGI plants had mercury controls during the RGGI program in the time frame covered by my panel. I exclude mercury controls from my analysis for this reason.

Figure 1.5: *The fraction of plants with scrubbers by group over time*

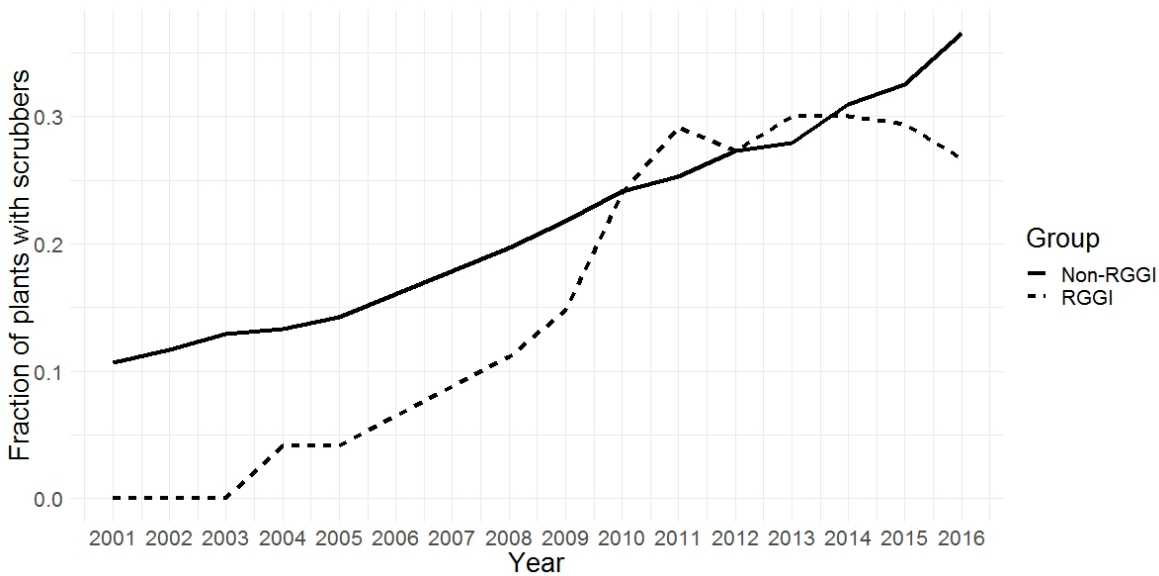

Note: Data on whether plants have scrubbers is from the EPA CEMS data. The universe of coal plants is defined by the criteria described in Section 3.4.

1.3 A model of coal plant operation and investment

In this section, I provide a model of how coal plant owners make production and efficiency choices. I examine both static and dynamic investments in efficiency improvements in a two-period framework.[15]

Let the efficiency of a plant, e_t, be a function of the state of its efficiency, ω_t, and within-period expenditures on efficiency, γ_t:[16]

$$e_t = g(\omega_t, \gamma_t) \tag{1.1}$$

[15]Plant owners can make short-run (within-period) heat rate adjustments by, for example, instructing plant personnel to optimize boiler operations or conducting maintenance. Long-run, between-period changes include major upgrades or the installation of new equipment. See Section 1.2.3 for more detail. The period length depends on the amount of time required to make major changes at the plant and start to experience efficiency improvements after those changes. Many capital-intensive forms of improving efficiency will impact efficiency in one to five years (Nowling, 2015). Linn et al. (2014) aggregate their data to 5-year periods to capture long-run heat rate changes at US coal plants.

[16]In this model, e_t is the inverse of the heat rate of a coal plant. Example units for e_t include MWh/mmBtu.

I consider ω_t to be a state variable that is altered by the plant owner's choice to invest in long-run efficiency improvements. Let the dynamic movement of ω_t be given by:

$$\omega_{t+1} = h(\omega_t, i_t) \tag{1.2}$$

in which i_t is investment in long-run efficiency improvements.

In each period t, the firm chooses a quantity of electricity to produce, q_t, that is a function of the plant's efficiency and f_t, the amount of fuel (coal) it uses: $q_t = e_t f_t$.[17,18,19,20] It also chooses within-period efficiency changes, γ_t, and investments that impact the next period's efficiency, i_t, to maximize profit. Its objective function is:[21,22]

$$\max_{q_t, i_t, \gamma_t} \pi_t(\omega_t) = p_t^w q_t - c(q_t, p_t^c, \gamma_t, i_t) + \delta E[\pi_{t+1}(\omega_{t+1})] \tag{1.3}$$

in which p_t^w is the wholesale electricity price, p_t^c is the carbon price, $c(\cdot)$ is a cost function, and $\delta E[\pi_{t+1}]$ is the discounted expected profit in period $t+1$. Equation (1.3) is a dynamic, recursive optimization problem that could be used to evaluate decisions over an infinite horizon

[17] For simplicity, I assume a one-firm, one-plant relationship.

[18] The fuel choice is in terms of heat input. Example units include mmBtu.

[19] Coal plant efficiency is also, mechanically, a function of output. I show in Appendix section A.2 that modelling it as such does not appreciably change the results from this model.

[20] Here, $q_t = e_t f_t$ is a simplified production function. Fuel and operations and maintenance (O&M) costs are the most important variable costs in coal plant operation. I choose not to include O&M costs that do not contribute to plant efficiency. The majority of capital is installed during plant construction. Labor changes are not considerable in this industry. Linn et al. (2014) similarly model coal plant efficiency choices in a two-period model that abstracts from capital and labor choices. Knittel (2002) uses a production function for electricity generation that does include capital and labor.

[21] Profit-maximization applies most closely to the incentive structure faced by owners of plants in deregulated markets. Over 80% of plants in RGGI are deregulated. Owners of plants that operate in traditional vertically-integrated, regulated settings will have a different objective function and, depending on regulatory incentives, different incentives for efficiency improvements than are given by this model. See Section 1.8 for a robustness check I conduct to ensure that my main results are not due to regional differences in plant regulatory status.

[22] I assume wholesale electricity markets are competitive. I also assume that the carbon price is exogenous, which is reasonable for the RGGI case. (See Section 1.2.2.) I do not include commitments in the capacity market. (The capacity market is a market in which plant owners commit to be available to provide electricity during a certain future time. It serves as a form of non-generation revenue for many plant owners.) I also do not model firm exit. Though there is considerable retirement of coal plants during the time of my sample, this retirement is due to exogenous forces (namely the availability of cheap natural gas). It is most interesting to examine efficiency choices among remaining coal plant owners, as future entry is not expected.

or fewer periods.

I use Equation (1.3) to look at decisions made in two periods only. The time frame of my data only includes seven years after the RGGI program start date, so it is plausible that RGGI coal plant owners only made one long-term efficiency investment decision in response to the carbon price. Additionally, the program has coincided with the era of low natural gas prices in the US; it is likely that coal plant owners have employed a "wait-and-see" approach to investment. (More precisely, owners may have regarded the information set about investments in future periods to be limited.) I transform Equation (1.3) into a two-period problem in Equation (1.4) below. To fully separate the production and efficiency choices, it is helpful to consider f_t the production choice variable. Finally, I explicitly define the terms of the cost function and assume a linear approximation of production costs.[23] The two-period objective function is:

$$\begin{aligned} \max_{f_t,\gamma_t,i_t} \pi_t(\omega_t) = p_t^w e_t f_t - p^f f_t - m p_t^c f_t - \gamma_t - i_t + \\ \delta E[(p_{t+1}^w e_{t+1} f_{t+1} - p^f f_{t+1} - m p_{t+1}^c f_{t+1} - \gamma_{t+1})] \end{aligned} \tag{1.4}$$

subject to $f_t \in [L, U]$.

The terms of the cost function include: the cost of production, $p^f f_t$ (the product of the cost of fuel, p^f, and f_t); the cost of complying with the carbon price, $m p_t^c f_t$, in which m is the emissions rate of fuel; and expenditures on changing efficiency for the current and future periods, γ_t and i_t, respectively.[24] I assume that production has lower and upper bounds: $f_t \in [L, U]$. The lower bound, L, is either zero or a minimum value necessary for the plant to maintain function, depending on the length of the period under consideration. The upper bound, U, depends on the plant capacity (as well as period length).

By taking first order conditions (FOCs) of (1.4) with respect to each of f_t, γ_t, and i_t, I

[23]Note that a linear approximation is reasonable, as my data provide one-month or longer time periods for operational variables and therefore reflect average prices and production. If one were to use this model to evaluate decisions made within shorter time periods (such as production decisions for real-time dispatching of electricity), one may want to use a nonlinear cost function.

[24]I assume that the price and carbon intensity of the fuel supply for the owner's coal plants have been determined through long-term contracts.

gain some intuition about how firms make these choices, given the carbon price.

These FOCs are:[25,26]

$$\frac{\partial \pi}{\partial f_t} = p_t^w e_t - p^f - mp_t^c = 0 \tag{1.5}$$

$$\frac{\partial \pi}{\partial \gamma_t} = p_t^w f_t \frac{\partial e(\cdot)}{\partial \gamma_t} - 1 = 0 \tag{1.6}$$

$$\frac{\partial \pi}{\partial i_t} = -1 + \delta E[(p_{t+1}^w f_{t+1} \frac{\partial \omega_{t+1}(\cdot)}{\partial i_t})] = 0 \tag{1.7}$$

Rearranging (1.5) results in:

$$e_t = \frac{p^f + mp_t^c}{p_t^w} \tag{1.8}$$

Equation (1.8) shows that an optimizing firm will maintain a level of efficiency that is directly proportional to its total marginal cost of operation from production and compliance with the carbon price ($p^f + mp_t^c$). Conditional on the firm having optimized over output, a higher carbon price induces greater efficiency. This provides a theoretical microeconomic foundation for examining efficiency improvements as a function of carbon pricing in this setting. The efficiency choice is also related to the cost of fuel, p^f, directly; a higher coal price should lead the firm to improve efficiency. As described above, Linn et al. (2014) found that US coal plant owners have historically improved plant efficiency in response to coal price changes, indicating that a carbon price would have the same effect.

[25] Assuming an interior solution for production. Otherwise, Equation (1.5) is an inequality, with the direction of the sign determined by whether production is at the lower or upper bound.

[26] Assuming $\frac{\partial e(\cdot)}{\partial \omega_{t+1}}$, the change in efficiency from the inherited state of efficiency, is one.

Combining (1.8) and (1.6) results in:

$$\frac{e_t}{f_t} = (p^f + mp_t^c)\frac{\partial e(\cdot)}{\partial \gamma_t} \tag{1.9}$$

Equation (1.9) shows how the firm trades off its efficiency and production choices in the current period, in the face of a carbon price. As the carbon price rises, the firm will either become more efficient or decrease production, all other terms held constant. The greater the change in plant efficiency with respect to within-period efficiency expenditures, $\frac{\partial e(\cdot)}{\partial \gamma_t}$, the greater the incentive to improve efficiency and/or decrease production. This term captures some of the heterogeneity among firms in the propensity to make adjustments in response to a carbon price.

Equation (1.9) motivates my empirical strategy of estimating the relationship between the RGGI allowance price and coal-fired power plants' efficiency. After I find that RGGI coal plants' efficiency *decreased* during the RGGI program period, Equation (1.9) motivates my second strategy of estimating the relationship between RGGI coal plants' efficiency and output. In practice, the efficiency of coal plants declines as plants reduce output, due to the mechanical relationship between efficiency and output, as described in Section 1.2.3. Equation (1.9) shows that, depending on the magnitude of that response, plant owners should improve or "let go" of efficiency to operate optimally from an economic standpoint.[27]

Rearraging (1.7) results in:

$$\delta E[(\frac{\partial \omega_{t+1}(\cdot)}{\partial i_t})] = \frac{1}{\delta E[p_{t+1}^w f_{t+1}]} \tag{1.10}$$

which tells us that investments in next-period efficiency improvements, i_t, are chosen such that their expected impact on efficiency ($\delta E[(\frac{\partial \omega_{t+1}(\cdot)}{\partial i_t})]$) is inversely proportional to expected marginal revenue ($\delta E[p_{t+1}^w f_{t+1}]$). Lower expected marginal revenue raises the standard for efficiency investments in terms of their efficiency-improving potential. Higher expected marginal

[27]Fell and Maniloff (2018) found that RGGI coal plants reduced output, due to the program, in its first four years. I examine the role of output in efficiency changes to see whether this explains the efficiency decline I observe among RGGI plants. This is described further in Section 1.4.2.

revenue means that less impactful efficiency investments are profitable. If firms are at all considering or suspecting retirement to be on the horizon in coming years (the special case of $f_{t+1} = 0$), capital-intensive efficiency improvements are not profitable. Notably, $(\delta E[p^w_{t+1} f_{t+1}])$ for a given firm is highly sensitive to the electricity market's reliance on coal-fired power. This is a microeconomic example of what Acemoglu (2003) calls the "market size" effect in innovation incentives: if the electricity sector reduces production from coal relative to other sources, then the incentive for firms to adopt new, improved equipment falls. If RGGI coal plant owners had unique expectations about future revenue, this could have led them to invest in efficiency-improving capital differently than owners of plants operating in other regions.

Equation (1.10) motivates my empirical strategy of looking at whether RGGI coal plant owners improved the heat rates of their plants over the long-run (with my data aggregated to two-year time periods), conditional on output. (See Section 1.4.4.) This allows me to assess whether RGGI coal plant owners made investments in efficiency-improving capital as would normally be incentivized by a carbon price.[28]

Equation (1.10) motivates some additional empirical choices I make. First, I estimate the average difference between RGGI and non-RGGI coal plant net heat rates for all coal plants as well as for a subset that had not announced a planned retirement date to the Energy Information Administration (EIA) (Section 1.4.1). This is a simple way to examine whether anticipated retirement is behind the unique heat rate trend among RGGI plants. Second, based on observed heat rate trends in my data, I take the perspective that equation (1.10) could be relevant for relatively short time periods and control for coal units being within six months of retirement, when I estimate a model of coal unit gross heat rates (Section 1.4.2). Finally, I examine time trends in my models of coal unit gross heat rates and coal plant auxiliary power consumption to see if there is any indication of long-run efficiency improvement or deterioration, conditional on output changes.

[28]Note that, because the RGGI program allows participants to bank allowances, the RGGI price captures expectations about the stringency of the program. I do not test the impact of expectations about future revenue changes from other market forces.

1.4 Empirical strategy

1.4.1 The regional difference in coal plant heat rates

To estimate the statistical difference between RGGI and non-RGGI coal plant net heat rates during the time period of the RGGI program, I estimate the coefficients of the following two models:

$$log(HR_n)_{it} = \beta_0 + \beta_{RP}\mathbf{RGGI}_i * \mathbf{POST}_t + \beta_R\mathbf{RGGI}_i + \beta_P\mathbf{POST}_t + \beta_{\mathbf{X}}\mathbf{X}_{it} + \theta_i + \gamma_t + \epsilon_{it} \quad (1.11)$$

$$log(HR_n)_{it} = \beta_0 + \beta_{RP}p_t^R + \beta_{\mathbf{X}}\mathbf{X}_{it} + \theta_i + \gamma_t + \epsilon_{it} \quad (1.12)$$

With equation (1.11), I use a difference-in-differences design to estimate the relationship between coal plant heat rates and participation in the RGGI program. With (1.12), I estimate the relationship between coal plant heat rates and the RGGI allowance price. For each, my outcome is the log net heat rate of coal plant i in month t, $log(HR_n)_{it}$.[29]

In equation (1.11), the coefficient of interest, β_{RP}, is on the interaction term $\mathbf{RGGI}_i * \mathbf{POST}_t$, which is an indicator for coal plants that are in RGGI during the RGGI program period (post-2009). As in any standard difference-in-differences design, I include independent indicator terms for the group of interest, RGGI coal plants, $\mathbf{RGGI}_i$, and the time period of interest, the RGGI program period, $\mathbf{POST}_t$. Identification of β_{RP} requires parallel trends between non-RGGI and RGGI coal plant log net heat rates in the pre-program period. I examine this assumption by estimating (1.11) with the RGGI program start year set to placebo years and discuss results in Section 2.6. Interpretation of β_{RP} as a causal parameter additionally requires no other common shocks to RGGI coal plants' heat rates aside from the RGGI program, post-2009. The results from these models are intended to show in a

[29] I use log heat rates as outcome in all models of heat rates in this paper for two reasons: 1) the distribution of heat rate values in the monthly data I use has a long right tail and 2) this is common practice in the literature. See Appendix section B for details on data handling for heat rate variables.

high-level, stylized way the statistical significance and magnitude of the the regional difference in heat rates between RGGI and non-RGGI plants during the RGGI program, so I do not include other factors that may be driving this difference or interpret β_{RP} as a causal program effect. Instead, I discuss the role of factors other than the RGGI program below, in my set of estimation strategies designed to determine the mechanism underlying this regional difference I describe by estimating (1.11) and (1.12). I include plant fixed effects (θ_i), so that I am capturing the average plant-specific change in net heat rates among RGGI plants during the program. I include month fixed effects (γ_t) to capture aggregate changes in plant heat rates due to seasonality.[30]

I also include a few time-varying plant- or state-specific observables (in $\mathbf{X}_{it}$) that are known contributors to plant efficiency: the presence of scrubbers and the price of coal. Scrubbers, which plants install to comply with sulfur emission standards, typically impose a heat rate penalty on plants of about 1%. Linn et al. (2014) find that US coal plant owners change their heat rates in response to delivered coal prices. The inclusion of these observables is to demonstrate the regional difference in heat rates between RGGI and non-RGGI coal plants in the RGGI program period, conditional on a few variables that could be common reasons for the difference.

To estimate (1.12), I regress the log net heat rate of coal plant i in month t on the quarterly RGGI allowance price as well as the same variables in $\mathbf{X}_{it}$, plant fixed effects, and month fixed effects. By estimating this equation, I obtain β_{RP}, the average percent change in plant-level net heat rates associated with a \$1 change in the RGGI price. As with equation (1.11), I am using this estimation to capture the significance and magnitude of the change in RGGI plant heat rates that is associated with RGGI participation, an association I for which I pursue explanation in analyses below. (Here RGGI participation is represented by the stringency of the RGGI program in each quarter through the RGGI price.) I set the allowance price equal to zero for plants which are never included in RGGI and for RGGI plants in the pre-RGGI

[30]Coal plant heat rates fluctuate throughout the year due to changes in electricity demand and temperature. On average, heat rates are worse in the spring and late summer months due to these impacts.

period.

For both models above as well as all models below, I exclude coal plants from Pennsylvania and Ohio. Fell and Maniloff (2018) show that, during the first three years of RGGI (2009-2012), electricity production in "Leaker" states Pennsylvania and Ohio significantly increased their generation, and exports of electricity from those states into RGGI states increased. Because coal plants in these states may have been quasi-treated by the RGGI program, excluding them is important for identifying the unique trend among RGGI plants compared to plants in the rest of the US.

I estimate both models above for two samples: the set of all coal plants in my panel and a set that is restricted to plants that have neither retired nor announced a planned retirement date to the EIA. Though plant owners may certainly be considering retirement in absence of having made that announcement, subsetting on this group allows me to see whether results are driven by a decision to retire. This choice is motivated by equation (1.10) in my model above.

1.4.2 The mechanism of the regional difference in coal plant heat rates

Decomposition of the source of efficiency change

To uncover the mechanism of the difference in heat rates between RGGI and non-RGGI plants during the RGGI program years, I first decompose the change in plant-level net heat rates into changes in their gross heat rates and auxiliary power consumption. This allows me to identify, from an engineering perspective, the source of efficiency change at these plants. Additionally, it allows me to exploit the detail in the data available on coal plant gross heat rates, which is available at unit-level. Finally, coal plants' gross heat rates and auxiliary power consumption are impacted by different factors. Broadly speaking, gross heat rates are more likely to be influenced by generation shifts, and auxiliary power consumption is more likely to be influenced by the use of environmental control equipment. Both of these have changed for RGGI plants in particular. As seen in Figure 1.6, RGGI coal plants have reduced their generation as a fraction of their capacity in the post-RGGI period more so than other coal plants. Identifying the specific source of efficiency changes at RGGI coal plants will help me

identify the likely determinants of the efficiency change seen during the RGGI program time.

Figure 1.6: *The distributions of monthly adjusted capacity factors of coal-fired power plants by program participation*

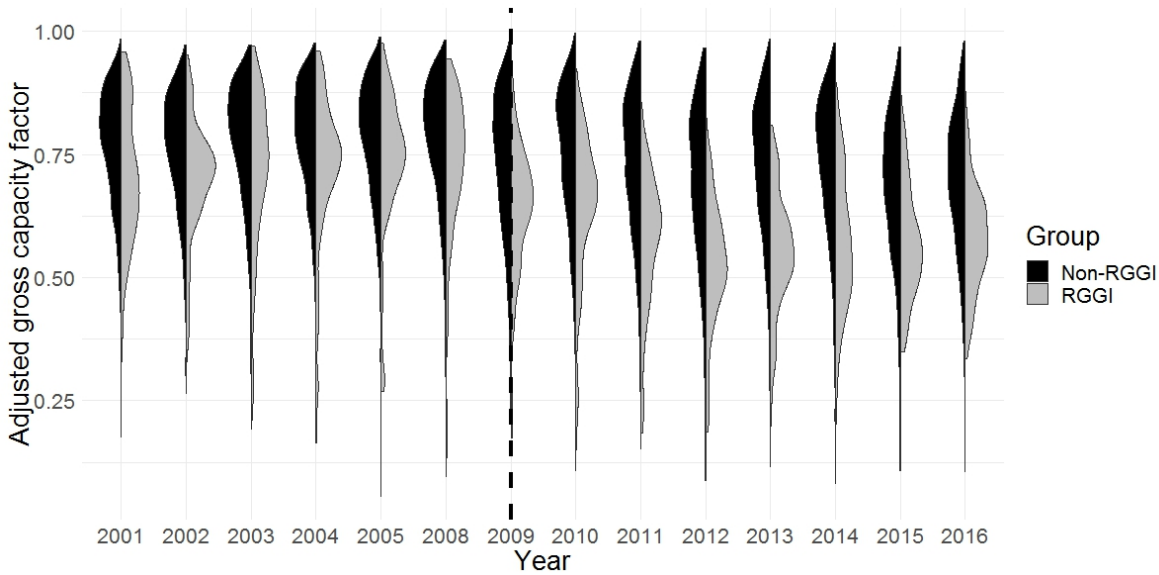

Note: The dashed vertical line at 2009 demarcates the beginning of the RGGI program. Calculation of the adjusted gross capacity factor variable is in Section 1.4.2.

A coal-fired power plant's net heat rate (HR_n), mechanically, is given by following relationship between its gross heat rate (HR_g) and ratio of its gross generation to net generation ($\frac{gen_g}{gen_n}$), a reflection of its auxiliary power consumption:

$$HR_n = HR_g * \frac{gen_g}{gen_n} \tag{1.13}$$

Or, more explicitly, using HI to denote plant-level heat input (fuel used in mmBtu):

$$\frac{HI}{gen_n} = \frac{HI}{gen_g} * \frac{gen_g}{gen_n} \tag{1.14}$$

The terms of this relationship are separable if one takes logs of both sides of equation (1.13):

$$log(HR_n) = log(HR_g) + log(\frac{gen_g}{gen_n}) \tag{1.15}$$

I use (1.15) to conduct the following decomposition of the time series of changes in annual RGGI plant net heat rates. Note that this does not involve estimating any coefficients but simply transforming these variables and taking differences over time. I do this by plant and then average the change in net heat rate between years due to each variable on the right-hand side over all plants.[31]

$$\Delta_{by} log(HR_n) = \Delta_{by} log(HR_g) + \Delta_{by} log(\frac{gen_g}{gen_n}) \tag{1.16}$$

in which Δ_{by} denotes the change in the variable it precedes between base year b and year y. Holland et al. (2018) use a similar, identity-based approach in decomposing the determinants of the change in emissions damages from the electricity sector.

Analysis of gross heat rates

To quantify the role of generation and other factors in influencing coal plants' gross heat rates, I estimate the following model:

$$log(HR_g)_{it} = \beta_0 + \beta_{CF} ACF_{it} + \beta_{\mathbf{X}} \mathbf{X}_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{1.17}$$

In equation (1.17), $log(HR_g)_{it}$ is the log monthly gross heat rate for coal unit i. ACF_{it} is an adjusted capacity factor, calculated as the unit's gross generation (in MWh) in month t divided by the product of the unit's capacity (in MW) and the number of hours in the month that the unit was running.[32] I use the adjusted capacity factor as a measure of each unit's

[31] I create annual heat rate and generation measures for each plant by summing heat input and generation variables over all months. Creating annual values this way (instead of taking means) allows me to include heat input that is recorded for plant-month observations when there is zero generation. Correspondence with the EPA has indicated that the CEMS data can detect heat input for coal units that are producing or available to produce steam but not generating electricity. This decomposition therefore captures changes in long-run heat rates at plants that include all fuel used at the plant.

[32] Because unit-specific capacity is not available from the EPA CEMS hourly operations data, I proxy a unit's maximum historical generation over the years 2001-2018 for its capacity. The EIA form 860 data reveals that unit-level capacity changes among coal units in the time frame of my panel are very rare occurrences. When I merge the EPA and EIA boilers and visually inspect matches, I see that a unit's max historical generation from the EPA data is reasonably close to its listed nameplate capacity in the EIA data. Kotchen and Mansur (2014) use a similar approach in identifying EIA units likely subject to EPA's proposed Carbon Pollution Standards.

monthly generation rate, conditional on running. It importantly separates output changes that are due to low generation rates from output changes that are due to time that units are not running. A unit's gross heat rate is sensitive to the former through an engineering relationship.[33] I observe the relationship between unit monthly log gross heat rates and adjusted capacity factors to be slightly nonlinear, so include a quadratic term in equation (1.17).[34]

The capacity factor coefficients capture the average relationship between coal units' gross heat rate changes and capacity factor changes. This is a causal relationship that results from the mechanical link between efficiency and output of coal units. My coefficients capture the causal relationship conditional on there being no time-varying, unit-specific omitted variables jointly associated with gross heat rates and intensive-margin output changes.[35]

$\mathbf{X}_{it}$ includes other variables of interest and some control variables. As variables of interest, I include indicators for whether a unit will be "retired" in six months or less and whether the plant has a scrubber installed. I define "retirement" as a unit ceasing to operate for any

[33]RGGI plants, in particular, have shut off their coal units for days, weeks, or even months at a time, increasingly over time. This trend seems unique to the region, based on my inspection of many unit-level time series plots of gross generation for RGGI and non-RGGI plants from the EPA CEMS hourly data. A naive construction of a plant, monthly capacity factor would include the unit's capacity multiplied by the number of hours in a month as the denominator. The resulting variable captures both plant operation at low generation rates and the amount of time units are not running in one monthly capacity factor. Note also that these long shut-off periods results in unit-month observations for which the unit's capacity factor is zero, and these observations are necessarily dropped from estimating (1.17). Finally, aggregate measures of heat rates, which I use here, may also be impacted by units being turned off and on, as extra fuel is needed to heat boilers from being in a "cold" state. To test for this, I included a variable for "number of starts" in my estimation and found it to be small and statistically insignificant.

[34]For every estimation in this paper that includes unit or plant capacity factor on the right-hand side, I determined the appropriate functional form relationship for the model with my own plots of the the data and model residuals.

[35]One may be concerned that output rate changes are simultaneously determined with efficiency changes, given that plant efficiency impacts the marginal cost of operation. See Section 1.8.2 for evidence that there is not an issue with simultaneity in coal plant heat rates and capacity factors. One may also suspect that these coefficients capture, in addition to the mechanical change in efficiency from output changes, short-run adjustments that coal plant owners make to plants, as seen in Equation (1.9). In Section 1.6.5, I do not find evidence that coal plant owners are adjusting short-term heat rates in a way that is consistent with the incentives from a carbon price, conditional on fluctuations in heat rates that occur from output. Therefore, coal plant owners are likely not trying to compensate for output-related efficiency loss in the short term, and the heat rate-capacity factor relationship I estimate here can be interpreted as the causal, engineering relationship.

reason.[36] Equation (1.10) shows that anticipated future reduction in output can lead coal plant owners to defer investments in efficiency improvements. If units are slated to retire, even basic unit maintenance may not be cost-effective. While unit-level plots of gross heat rates revealed that some RGGI coal units had increasing heat rates for about two years before retirement, I examine the six-month window of time before retirement as this is the window in which there would be the largest average impact of units being pre-retirement on their gross heat rates, if any. Scrubbers typically impact the amount of auxiliary generation that plants use, but they can also impact the gross heat rate of units by decreasing boiler efficiency.

Linn et al. (2014) find that, conditional on a coal unit's capacity factor at time t, coal plant owners will make efficiency improvements in response to an input price shift. This motivates inclusion of the coal price as a control variable in (1.17). Additionally, I include year fixed effects or a time trend in various versions of (1.17) to look at potential efficiency changes, conditional on generation shifts. Specifically, given an average RGGI price of \$3.43 from 2009-2016, RGGI coal unit net heat rates should drop 0.2% during the RGGI program period (Linn et al., 2014). This does not give a precise prediction about the magnitude of potential gross heat improvements but indicates that they may exist. I estimate (1.17) on the sample of all US coal units and independently for the sample of RGGI units to explore trends by group.

The model includes unit and month fixed effects. The coefficients I obtain by estimating (1.17) give the average association between coal units' monthly gross heat rate and changes in each right-hand side variable, averaged across all units, conditional on common seasonal shocks.

I use coefficients from estimating (1.17) in a decomposition, described below, that quantifies the role of each right-hand side variable in RGGI coal units' gross heat rate changes during the RGGI program period.

[36]The EPA unit-level operational data from the CEMS often shows long time periods of no operation for units, even though they are listed as operational. Because the EPA and EIA data I use cannot be matched at a unit-level, I cannot verify official retirement status for units with the EIA data. I therefore assume units are retired if they cease to operate. Additionally, I observe units being converted to using other fuels, such as natural gas. I consider a coal unit "retired" if it is converted to another fuel.

Analysis of auxiliary generation

To understand the determinants of changes in RGGI plants' auxiliary power consumption, I estimate the following two models:

$$Aux_{it} = \beta_0 + \beta_{CF} CF_{it} + \beta_{\mathbf{X}} \mathbf{X}_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{1.18}$$

$$Aux_{it} = \beta_0 + \beta_{CO} CapOn_{it} + \beta_{\mathbf{X}} \mathbf{X}_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{1.19}$$

In (1.18) and (1.19), Aux_{it} is the percent of electricity that coal plant i uses on site as auxiliary power during month t (specifically, $(\frac{gen_n - gen_g}{gen_g}) * 100$).[37] CF_{it} is the plant's unadjusted capacity factor. (The unadjusted capacity factor is $\frac{gen_g}{MWhMax}$, in which $MWhMax$ is the maximum amount of electricity a plant could produce in a month, given its nameplate capacity, computed by multiplying its capacity (in MW) by the number of hours in a month). I look at the role of a plant's unadjusted capacity factor in its auxiliary power consumption to estimate the extent to which power consumption is associated with total generation fluctuations. (In contrast to the gross heat rate model above, I anticipate that auxiliary power consumption at a plant responds to total generation shifts, those due to low unit utilization and especially those due to turning units off. This unadjusted capacity factor captures both margins of output changes. Using this variable also allows me to estimate the total impact of generation shifts on auxiliary power consumption for my decomposition, described below.)

To look at the role of turning units off, a phenomenon that has occurred at RGGI plants in particular with increasing duration over the post-RGGI period, I instead include $CapOn_{it}$ in equation (1.19), the percent of a plant's capacity that is used for generation each month. (I construct this variable by multiplying plant capacity (in MW) by the number of hours the

[37] This variable is limited to values between 0 and 100. OLS can lead to out-of-sample predictions of limited outcome variables, and one option for addressing this issue here would be use to the logistic transform of $(\frac{gen_n - gen_g}{gen_g})$. However, the logistic transform is not well-suited to the decomposition analysis for which I use estimated coefficients, described below, as it is very difficult to obtain conditional expectations of the outcome variable after it has been transformed (Papke and Wooldridge, 1996). I find that OLS does not generate out-of-sample predictions for this model. Additionally, if I estimate (1.19) with the logistic transform and run my decomposition with the logitistic transform as the outcome, results are not appreciably different.

plant was running, and dividing that product by the total potential MWh that the plant could produce in a month.) For regressions that include plant capacity factor, I include a quadratic capacity factor term to capture the upward curvature in the relationship between plant auxiliary power consumption and capacity factor that I observe. β_{CF} and β_{CO} will be non-biased if I have not excluded any plant-level, time-varying variables that are jointly correlated with plants' capacity factors and auxiliary power consumption or capacity online during the month and auxiliary power consumption, respectively. As discussed above in Section 1.4.2, these coefficients do not reflect causal relationships, as a plant's generation is based on its marginal cost, which changes with the share of electricity the plant consumes on-site. In Section 1.8 I test for simultaneity bias.

$\mathbf{X}_{it}$ includes the delivered coal price and indicators for whether a plant has an installed scrubber and whether it has a selective catalytic reduction unit. The coal price acts as a direct incentive for coal plant owners to use less electricity on site, and those environmental controls (particularly scrubbers) may cause a plant to use more auxiliary power.[38] Both models include month and plant fixed effects. I include a time trend in one specification to see whether US coal plants have improved their use of power on-site over time, as they have their gross heat rates.

I use estimated coefficients from the model along with the data for RGGI plants only for the decomposition described below, to quantify the role of each right-hand side variable in changes in RGGI plants' auxiliary power consumption.

1.4.3 Decompositions of RGGI gross heat rates and auxiliary power consumption

I use the coefficients I obtain from estimating (1.17) and (1.18) to conduct model-based decompositions to quantify the contribution of each independent variable to RGGI plants' gross heat rate change or auxiliary power consumption, respectively, similar to Coglianese

[38] Only five RGGI plants installed SCR units 2008 or later, so it is unlikely that SCR units have played a large role in RGGI plant auxiliary power consumption or net heat rate changes.

et al. (2018). Specifically, for the model given by (1.17), consider that the change in actual mean log gross heat rate values (across all plants, all months) can be represented as:

$$\Delta_{2009,y} HR_g = [E(HR_g|CF_y, \mathbf{X}_y) + v_y] - [E(HR_g|CF_{2009}, \mathbf{X}_{2009}) + v_{2009}] \tag{1.20}$$

in which $E(HR_g|CF_y, \mathbf{X}_y)$ is the conditional expectation of log gross heat rate values generated by plugging in year y observed independent variable values (including fixed effects) into the linear predictor obtained from estimating (1.17). v_y is the unexplained error, the difference between observed and predicted values, for year y. One can expand (1.20) into a sum of conditional expectations, each generated by changing one variable's values from 2009 to year y values:

$$\begin{aligned} \Delta_{2009,y} HR_g = [E(HR_g|CF_y, \mathbf{X}_y)] - [E(HR_g|CF_{2009}, \mathbf{X}_y)] \\ +[E(HR_g|CF_{2009}, \mathbf{X}_y)] - [E(HR_g|CF_{2009}, \mathbf{X}_{2009})] \\ +[v_y - v_{2009}] \end{aligned} \tag{1.21}$$

Each term yields the change in the outcome, the mean annual log gross heat rate, due to the variable that is changed from the 2009 to y year value. For example, the change in 2010 mean gross heat rate due to a change in capacity factor values between years is $[E(HR_g|CF_y, \mathbf{X}_y)] - [E(HR_g|CF_{2009}, \mathbf{X}_y)]$. (The variables contained in $\mathbf{X}_y$ and the fixed effects can be changed one at time to generate more terms in (1.21), but this is withheld for brevity.) The identification and interpretation of the change in the outcome between years due to each variable follows directly from OLS coefficient identification and interpretation.

I perform a decomposition that is generalized by equation (1.21) with both models (1.17) and (1.18) above.

1.4.4 Did RGGI impact coal plant efficiency change, conditional on output changes?

The results from the analyses above indicate that unit- and plant-level output shifts explain the majority of the change in RGGI coal plants' heat rates. Equations (1.8) and (1.9) show that, conditional on a plant's generation choice, a carbon price provides a direct incentive for efficiency improvements. Linn et al. (2014) find that a $10 per ton CO2 price should induce coal plants to improve their net heat rates by 0.6%, using coefficients from their model of coal unit net heat rates with plant fixed effects. This implies that the RGGI allowance price, which has been an average of $3.43 from 2009-2016, should have led RGGI coal plants to improve their net heat rates by about 0.2%. To look at this for RGGI coal plants, I estimate the following, which is very similar to equation (1.12):

$$log(HR_n)_{it} = \beta_0 + \beta_{RP} p_t^R + \beta_{CF} log(CF_{it}) + \beta_{\mathbf{X}} \mathbf{X}_{it} + \theta_i + \gamma_t + \epsilon_{it} \tag{1.22}$$

To estimate (1.22), I regress coal plants' log net heat rates on the RGGI price (for RGGI plants), their log [unadjusted] capacity factors, $log(CF_{it})$, and control variables in $\mathbf{X}_{it}$, which includes indicators for whether a plant has a scrubber or SCR unit and the coal price. I include the log of capacity factor as it provides the best fit for the model (based on examination of residual plots) and allows me to compare the coefficient on the capacity factor term to the elasticity I estimate below for calculating rebound effects. I include plant (θ_i) and month (γ_t) fixed effects. I include a time trend in one specification, as I found trends in both my gross heat rate and auxiliary power consumption models above. A negative (positive) time trend in net heat rates may indicate that coal plant owners are systematically improving (neglecting) efficiency, conditional on the controls included here. Alternatively, it may indicate the presence of other omitted variables driving a long-run trend in US coal plant heat rates.

Estimating equation (1.22) with monthly data identifies changes in RGGI coal plant heat rates associated with simultaneous, short-run (quarterly) changes in the RGGI allowance price with the coefficient β_{RP}. These changes are due to changes in how plants are managed

(such as how plant operators control boiler conditions), maintenance, and, occasionally, large maintenance projects or installation of new technology. Bushnell and Wolfram (2009) show that individual plant personell can impact the gross heat rate of coal units by 0.01 - 0.03%.

To estimate whether the RGGI price induced coal power plant owners to improve the efficiency of their plants through major plant maintenance projects and/or the adoption of new equipment, I estimate equation (1.22) with my data aggregated into two-year values. This identifies β_{RP} as the average percent change in coal plant net heat rates associated with a $1 change in the RGGI allowance price over two-year periods, capturing how coal plant owners are altering the efficiency of their plants in expectation of the allowance price over the medium-term. I can compare β_{RP} to estimates provided by Linn et al. (2014), cited above.[39] I take two-year averages of all numeric variables and create two-year indicator variables for whether coal plants have a scrubber or SCR unit. For this estimation, I do not have month fixed effects, given that the data is aggregated. I include plant fixed effects. In one specification, I also include a time trend.

Estimating 1.22 with my data aggregated to two-year values also gives me a long-run elasticity of coal plant net heat rate changes with respect to capacity factor changes, β_{CF}, which I use for my rebound estimates. (See section 1.7).

1.5 Data

For this paper, I have assembled two panels of monthly data on coal plant operational variables and characteristics, from combination of public sources on the US electricity sector, for the years 2001-2016. One panel provides data on coal plants; the other provides data at the sub-plant, unit level. I describe these panels in general terms here and discuss more detailed data handling choices in Appendix section B.

For the coal plant panel, I rely on the Energy Information Administration (EIA) Form-767

[39] Ideally I would use five-year periods, as Linn et al. (2014) do, but that would yield only two periods during the RGGI program.

("767") and Form-923 ("923") for data on fuel used and net electricity generated by plants.[40] All US plants with a steam-powered unit 10 megawatts (MW) or larger in capacity were required to report monthly operational variable values on an annual basis on the 767 until 2005. The EIA consolidated the 767 into the 923 starting in 2006 and resumed its monthly survey in 2008. The 923 monthly survey requires plants with at least 10 MW of capacity to report these data monthly if they are part of the EIA's monthly respondent sample. Therefore, my plant-level panel does not have years 2006 and 2007, and the composition of plants changes over time due to the monthly sampling. The missing years and changing plant composition of the panel do not pose a problem, as I do not include dynamic variables in my estimation. Additionally, the compositional change is small: about 98% of plants are kept in the sample pool, year-to-year.

I merge the EIA operational data to coal plant characteristics data available from the EIA Form-860. I identify coal plants as plants that have at least one generation unit that burns a type of coal as its primary fuel in each year, as listed on the 860.(Chan et al. (2017) use the same rule.) I exclude observations associated with plants that have combined-cycle units (coal or natural-gas fired), as these units operate in a fundamentally different way than simple-cycle units that impact their efficiency. I also eliminate observations for which the unit technology cannot be identified. This eliminates 28 plants from my data (5.9%) (15 of which have combined-cycle units) and 2.3% of my observations. The remaining plants define the universe of plants both of my panels. I also obtain the planned retirement date and retirement date each plant owner lists for its units from the 860.

To my EIA panel I merge monthly aggregations of coal plant operational data, including heat input and gross generation, from the EPA Continuous Emissions Monitoring Systems (CEMS) data.[41] These data are generated from the EPA's monitoring of fossil-fired units for compliance with CAA Amendment emission regulations (Cicala, 2019). These data also allow

[40] All EIA form data is available for public download from the agency's website. Currently each form has a page with historical data; the 923 form data, for example, can be found at https://www.eia.gov/electricity/data/eia923/.

[41] See https://ampd.epa.gov/ampd/.

me to observe coal unit CO2 emissions rates.

For my panel of coal unit operations, I create monthly aggregates of unit data from the CEMS data, for the coal plants I identify from the EIA data as described above.

I acquire plant RGGI participation status, the type of control method that plants use for SO2 emissions and NOx emissions, and a number of other plant and unit characteristics from the EPA's Air Markets Program Data database.[42] I match these data to both panels.

I use the EIA's State Energy Data System (SEDS) for annual state average prices of coal and natural gas delivered to the electric generation sector.[43]

For quarterly RGGI auction prices, I extracted the time series of auction results from the RGGI, Inc. website.[44]

Summary statistics, with the heat rate restrictions described in Appendix section B, are shown in Tables 1.1 and 1.2, for RGGI and non-RGGI plants, respectively. While the mean net heat rate of RGGI plants (11.01) is higher than that of non-RGGI plants (10.8) over all years, the mean gross heat rate is not, suggesting that gross efficiency at RGGI coal plants may not be the driver of their change in net efficiency. I test this formally below. The mean auxiliary power consumption is not notably different between groups. Depending on plant age and environmental control equipment, plants may have auxiliary power consumption rates of 5-15% of gross electricity generated (ABB, 2009). (See Appendix section B for how I handle extreme values of this variable.) RGGI coal plants' mean adjusted capacity factor (64.55) is over 10% lower than non-RGGI plants' mean adjusted capacity factor (74.39), indicating that RGGI plants have been running their coal units at lower generation rates than other plants. As discussed above, this can impact plant efficiency, but primarily through plants' gross heat rate.

[42] https://ampd.epa.gov/ampd/

[43] https://www.eia.gov/state/seds/seds-data-complete.php?sid=US#CompleteDataFile

[44] https://www.rggi.org/auctions/auction-results

Table 1.1: *Summary statistics of the 28 RGGI plants in my panel*

Statistic	Mean	St. Dev.	Min	Median	Max
Net heat rate (mmBTU/MWh)	11.01	1.27	6.19	10.77	15.99
Gross heat rate (mmBTU/MWh)	10.20	1.20	6.73	10.05	15.88
Auxiliary power consumption (pct total)	0.08	0.07	−0.51	0.08	0.91
Adjusted capacity factor (100 scale)	64.55	15.59	9.56	66.07	97.08
Plant capacity (MW)	825.61	586.14	70.50	727.80	2,647.00
Coal price (2010 USD/mmBTU)	3.42	0.49	2.37	3.35	5.02
NG price (2010 USD/mmBTU)	5.98	1.64	2.21	5.62	10.15

Table 1.2: *Summary statistics of the 349 non-RGGI plants in my panel*

Statistic	Mean	St. Dev.	Min	Median	Max
Net heat rate (mmBTU/MWh)	10.80	1.00	6.12	10.61	15.99
Gross heat rate (mmBTU/MWh)	10.21	1.06	6.12	10.10	15.99
Auxiliary power consumption (pct total)	0.07	0.29	−42.41	0.07	0.96
Adjusted capacity factor (100 scale)	74.39	13.70	8.79	76.81	99.58
Plant capacity (MW)	1,021.19	769.41	31.90	780.90	4,008.40
Coal price (2010 USD/mmBTU)	2.28	0.66	1.07	2.21	4.07
NG price (2010 USD/mmBTU)	5.53	1.50	1.96	5.36	9.68

1.6 Results

1.6.1 The regional difference in coal plant heat rates

Table 1.3 presents the result of estimating equation (1.11). Columns (1)-(3) present results for estimation on the set of all coal plants in my panel; columns (4)-(6) present results for estimation on the set of plants that have neither retired nor announced a planned retirement date. All models contain plant and month fixed effects, and the only difference among columns (1)-(3) (and also among columns (4)-(6)) is the addition of several variables: an indicator for a plant having a scrubber and the annual average price of coal delivered to a plant's state. The first row shows that RGGI coal plants were about 0.055% less efficient than the rest of the US in the period 2009-2016, after the start of the RGGI program.[45] This plant-level trend is depicted in Figure 1.1. As described in 1.2.3, coal plant owners consider heat rate reductions on the order of 0.10% to 3% with payback periods ranging from one to five years. Therefore, in the seven year period of the RGGI program period in my panel, one might expect to see coal plants improve their efficiency with at least the order of magnitude of this efficiency decline, if not more.

Table 1.4 displays results from estimating (1.12). The results of preferred specifications (3) and (6), which control for plants having scrubbers and the coal price, indicate that, for every $1 increase in the RGGI allowance price, RGGI coal plant net heat rates have risen at a plant level, on average, 0.018%. Given an average RGGI price of $3.43 2009-2016, this coefficient implies that RGGI coal plant net heat rates have been 0.062% higher in these years compared to non-RGGI plants, corroborating the results identified from program status, above.

Figure 1.7 shows results of estimating (1.11) with alternative years for the RGGI program start date (specifically, I change $\mathbf{POST}_t$ to an indicator for years post-2002, post-2003, etc.). The coefficient values that result from using years 2002 and 2003 as placebo program start dates are not significantly different from zero, indicating parallel trends in the early part of my panel. However, the coefficient is rising with each year of estimation, and it is significantly

[45] Note that I multiply the log heat rate by 100 to increase the coefficient values presented in my tables.

Table 1.3: *Results: The regional difference in coal plant heat rates*

(RGGI program identification)

	Dependent variable: Log Net Heat Rate x 100					
	Entire sample			No planned retirement sample		
	(1)	(2)	(3)	(4)	(5)	(6)
RGGI plants, post-RGGI	5.489***	5.396***	5.400***	5.281***	5.175***	5.183***
	(0.886)	(0.916)	(0.917)	(0.870)	(0.901)	(0.903)
Scrubber		x	x		x	x
Coal price			x			x
Observations	47,177	46,628	46,628	46,327	45,792	45,792
Adjusted R^2	0.708	0.705	0.705	0.711	0.709	0.709

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating the coefficients of equation (1.11). The "No planned retirement sample" excludes coal plants that have retired all units or have announced planned retirement for all units. All models include plant and month fixed effects. Standard errors are clustered by plant.

Table 1.4: *Results: The regional difference in coal plant heat rates*

(RGGI price identification)

	Dependent variable: Log Net Heat Rate x 100					
	Entire sample			No planned retirement sample		
	(1)	(2)	(3)	(4)	(5)	(6)
RGGI price (USD per ton CO2)	2.048***	1.942***	1.820***	1.977***	1.868***	1.764***
	(0.235)	(0.252)	(0.255)	(0.229)	(0.245)	(0.249)
Scrubber		x	x		x	x
Coal price			x			x
Observations	47,177	46,628	46,628	46,327	45,792	45,792
Adjusted R^2	0.700	0.699	0.700	0.704	0.704	0.705

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating the coefficients of equation (1.12). The "No planned retirement sample" excludes coal plants that have retired all units or have announced planned retirement for all units. All models include plant fixed effects and month fixed effects. Standard errors are clustered by plant.

different from zero in 2004, 2005, and 2008, the three years in my panel before the start of the RGGI program. This indicates that RGGI coal plant heat rates started to deviate from those of coal plants in the rest of the US before the RGGI program. Murray and Maniloff (2015) note that state announcements to join RGGI began in 2005. Therefore, if RGGI coal plant heat rate changes are program-induced, this pre-trend could reflect anticipatory behavior. The RGGI price should have lead to efficiency improvements, if coal plants had planned to continue operation with the same output once the program starts. However, equation (1.10) suggests that, if RGGI plants planned to reduce generation in response to the program, they may have held off on efficiency-improving investments. Heat rates could decline in RGGI plants before the program due to this investment behavior. Another explanation for the pre-trend is that there is at least one other factor responsible for the different trends in RGGI and non-RGGI coal plant heat rates aside from the RGGI program. I explore additional factors below.

Figure 1.7: *Parallel trends assumption*

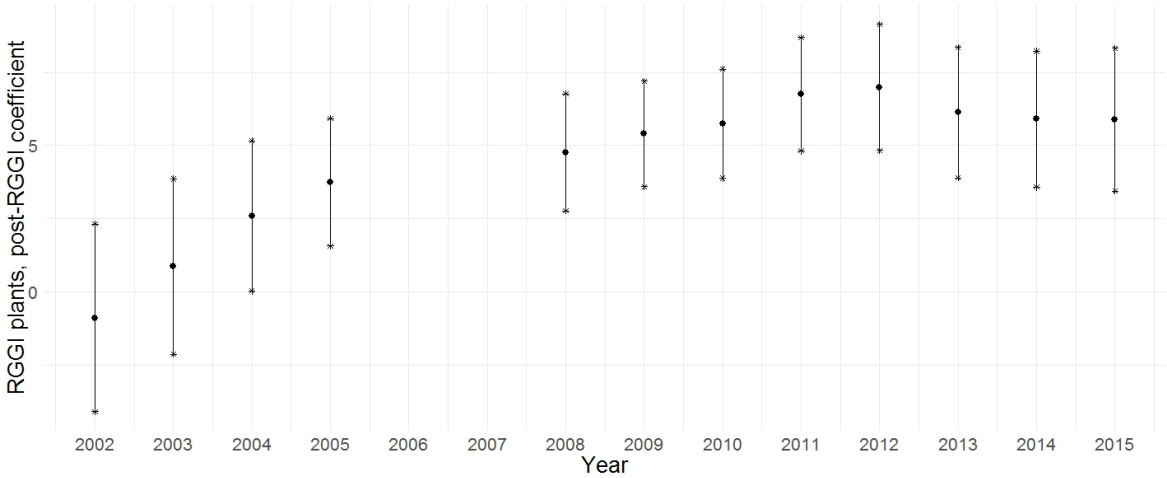

Note: Each dot represents the coefficient value on the term $\mathbf{RGGI}_i * \mathbf{POST}_t$ from estimating equation (1.11) once for each year on the x-axis, with $\mathbf{POST}_t$ set as an indicator for year t being beyond each respective year on the x-axis. Vertical lines with endpoints at stars for each year indicate 95% confidence intervals.

1.6.2 Decomposition of the source of efficiency change

Figure 1.8 displays the results of the decomposition of the time series of RGGI plants' net heat rate changes, 2009-2016. Specifically, it shows the share of each variable in the net heat rate change between 2009 and each subsequent year.[46] Between 2009 and 2010, most of the change in RGGI coal plant net heat rates was due to changes in electricity used on-site. However, in that initial year, plant-level net heat rates did not increase as much as later years (Figure 1.1). From 2010-2014, about 70% of the changes in RGGI plants' net heat rates from 2009 are attributable to changes in plant-level gross heat rates. From 2014-2016, the role of gross heat rate changes in RGGI plants' net heat rate trends since 2009 declined. About half of the total change in plant-level net heat rates from 2009 to 2016 is due to changes in gross efficiency and half is due to changes in auxiliary power consumption, as shown by the 2016 points on the plot. Overall, these results suggest that both changes in RGGI coal plants' efficiency in converting coal energy into electricity, their gross heat rates, and their use of electricity on site, their auxiliary power consumption, are implicated in their overall decline in efficiency, their net heat rates. This motivates exploring determinants of both gross heat rate and auxiliary power consumption changes at coal plants to understand why RGGI coal plants' heat rates changed so dramatically during (and even before) the time of the program. During the years 2010-2014, however, changes in plants' gross heat rates drove changes in net heat rates since the beginning of the program. This suggests that generation shifts may have played a large role in reducing RGGI coal plants' efficiency during that time, as generation shifts are the most likely driver of gross heat rate changes.

1.6.3 Analysis and decomposition of RGGI coal unit gross heat rate changes

Table 1.5 presents the results of estimating equation (1.17) on all US coal units. The top row of Table 1.5 indicates that units' generation rates, conditional on running, are negatively

[46] After calculating each term of equation (1.16) for each plant in each year after 2009, I average each term across all plants. I then express the average, annual change in the log gross heat rate (or the log gross to net generation ratio) between 2009 and each year as a fraction of the average, annual change in log net heat rates. The result is the share of each variable in the net heat rate change between 2009 and each subsequent year.

Figure 1.8: *RGGI net heat rate decomposition results*

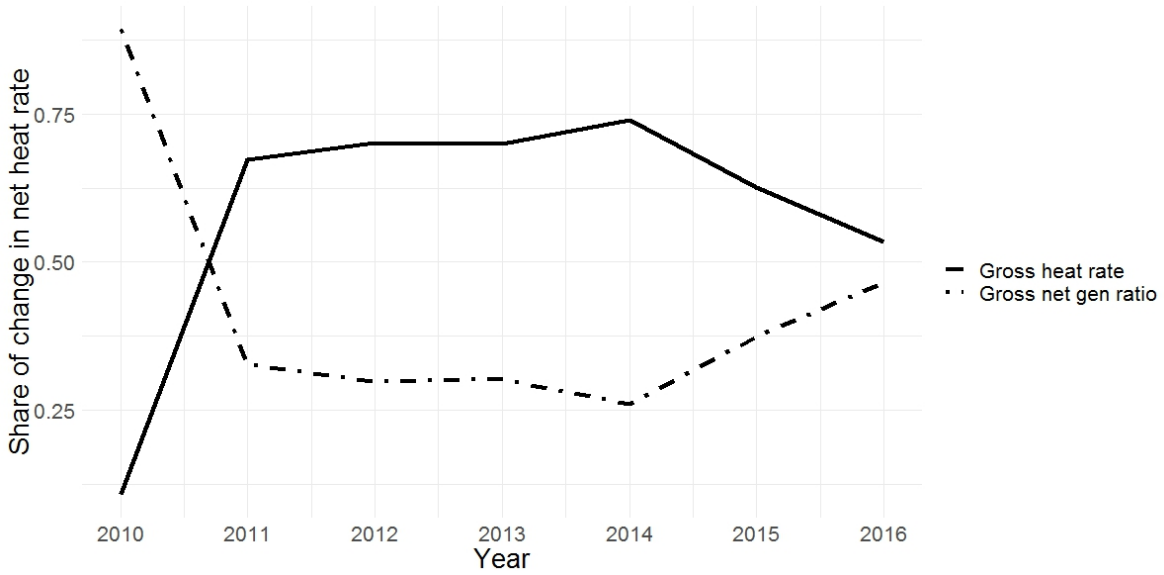

Note: Each line represents the average percent change in RGGI plants' annual net heat rates due to the variables labelled on the right, between each year on the x-axis and 2009. "Gross net gen ratio" is the ratio of RGGI coal plants' gross to net generation.

and significantly associated with their gross efficiency, as suggested by engineering literature. (Because I include a quadratic capacity factor term, the magnitude of the linear term is not directly interpretable.) Though not shown, the quadratic term is positive and statistically significant, which is consistent with the upward curvature of this relationship, particularly at low generation rates. The time trend is small but statistically significant, indicating that, conditional on generation changes, coal plant owners have improved the gross efficiency of their units about 0.0025 - 0.0029% per year, 2001-2016. This indicates that, despite the decline of the market share of coal-fired power plants in electricity generation in the US in the second part of this time period, coal plant owners are likely still maintaining the gross efficiency of their units. Though not shown, the individual year fixed effect coefficients corroborate this; each year fixed effect coefficient after after 2002 has a negative sign. Another explanation for these time trends is that the most efficient units remain in the market as others exit. In any case, there is not evidence of systematic deterioration of units.

Column (4) shows the results for additional variables of interest, an indicator a unit being

Table 1.5: *Results: Gross heat rate model (All US coal units)*

	Dependent variable: Log Gross Heat Rate x 100				
	(1)	(2)	(3)	(4)	(5)
Adjusted capacity factor (100 scale)	-0.513^{***}	-0.527^{***}	-0.514^{***}	-0.505^{***}	-0.504^{***}
	(0.115)	(0.114)	(0.114)	(0.115)	(0.114)
Retiring within 6 months				1.656^{***}	1.656^{***}
				(0.566)	(0.566)
Scrubber				2.427^{***}	2.426^{***}
				(0.802)	(0.802)
Time trend			-0.256^{***}	-0.290^{***}	-0.295^{***}
			(0.033)	(0.034)	(0.036)
Year fixed effects		x			
Coal price					x
Observations	136,897	136,897	136,897	136,897	136,897
Adjusted R^2	0.648	0.659	0.656	0.658	0.658

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating the coefficients of equation (1.17) for all US coal units. All models have a quadratic capacity factor term and unit and month fixed effects. Standard errors are clustered by plant.

within six months of retirement and whether a unit's plant has a scrubber installed. Each incur small gross heat rate penalties: units within six months of retirement are, on average, 0.017% less efficient than before that time. This implies that coal plant owners might "let their units go" before retiring them (or converting them to another fuel). Scrubbers reduce unit efficiency, on average, by 0.024%. Inclusion of the coal price does not impact any of these estimates (column 6).

In Table 1.6, I present the results of estimating (1.17) on RGGI coal units only, to look at any differences in time trends. Columns (4) and (5) of Table 1.6 show that RGGI coal units have also had declining gross heat rates over the entire time period, conditional on generation. This implies that maintenance neglect is likely not behind the net heat rate trend seen in RGGI coal plants, at least the share due to gross heat rate changes. However, this coefficient is half as large as that obtained from estimating the model on all US coal units; RGGI coal plant owners could have been less attentive to maintenance than those running plants in the rest of the US. (Compare the time trend term in columns (4) and (5) of Table 1.6 to those in the same columns of Table 1.5). Interestingly, the time trend does not seem to be identified for RGGI coal units unless the retirement indicator is included. (Compare the time trend term in column (3) to those in columns (4) and (5) of Table 1.6.) Though the retirement term is not statistically significant, this does show some gross heat rate impact of RGGI units being pre-retirement.

Figure 1.9 presents the results of decomposing the changes in RGGI coal units' gross heat rates between 2009 and subsequent years into the role of each right-hand side variable in (1.17), as described by (1.21). I use the coefficients obtained by estimating (1.17) for all coal units and the data for RGGI coal units for this decomposition. I use the model presented in column (5) of Table 1.5. Therefore, this decomposition assumes that the relationships between each right-hand side variable and coal units' gross heat rates are common to all units, RGGI and non-RGGI units. Each line of the figure displays either the observed average annual gross heat rate of RGGI coal units (the solid line) or the predicted annual average gross heat rate of RGGI units by keeping all variables at 2009 values and changing, one at a time, each variable

Table 1.6: *Results: Gross heat rate model (RGGI coal units)*

	Dependent variable: Log Gross Heat Rate x 100				
	(1)	(2)	(3)	(4)	(5)
Adjusted capacity factor (100 scale)	-0.764^{***}	-0.749^{***}	-0.768^{***}	-0.760^{***}	-0.766^{***}
	(0.195)	(0.193)	(0.196)	(0.179)	(0.179)
Retiring within 6 months				1.490	1.509
				(1.745)	(1.750)
Scrubber				3.264	3.367
				(2.387)	(2.384)
Time trend			-0.097	-0.178^{*}	-0.162^{*}
			(0.121)	(0.100)	(0.098)
Year fixed effects		x			
Coal price					x
Observations	8,065	8,065	8,065	8,065	8,065
Adjusted R^2	0.655	0.661	0.656	0.659	0.659

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating the coefficients of equation (1.17) for RGGI coal units only. All models have a quadratic capacity factor term and unit and month fixed effects. Standard errors are clustered by plant.

labelled in the upper left of the figure to new year values. Therefore, the distance between lines represents the additional contribution of the variable labelled for each line to the new year predicted value of the annual average RGGI gross heat rate of coal units. Unexplained or residual variation is represented by the distance between the solid line and the "Predicted" (lower-most, dashed-dot) line. Unexplained variation is 1% to 3% of the true average annual gross heat rate among all units, 2010-2016. i include the time trend in each predicted value line. The figure shows that the contribution of RGGI units' adjusted capacity factors to gross heat rate changes over the RGGI program time period have been relatively large compared to other variables. Scrubbers incur a relatively small and stable gross heat rate penalty. Units being six months from retirement and the coal price play a negligible role in changes in RGGI coal unit gross heat rates from 2009, post-RGGI.

Figure 1.9: *Gross heat rate decomposition results, All US model, RGGI coal units*

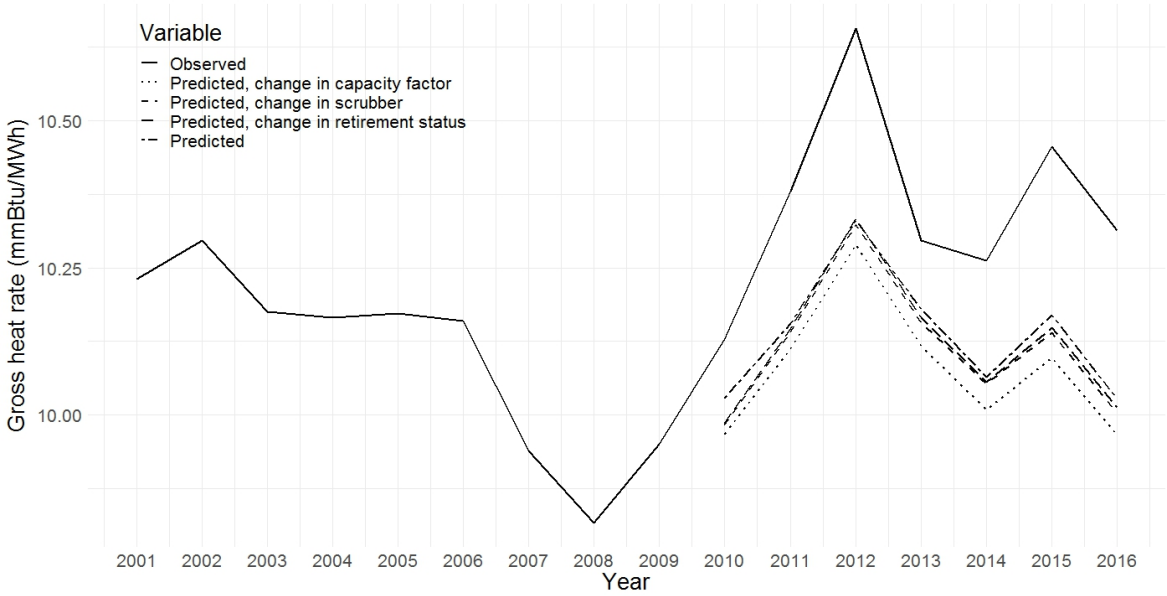

Note: The solid line is the average annual gross heat rate among RGGI coal plants. Each non-solid line represents the average annual predicted value of RGGI plants' gross heat rate while keeping all variables at 2009 values and changing each variable labelled in the top left, sequentially, to the new year values (with each new year labelled on the x-axis). Therefore, the space between lines shows the contribution of each variable to the new year predicted average annual gross heat rate value. See equation (1.21) for decomposition methodology. Predicted values (conditional expectations for each year) are obtained by using estimated coefficients from equation (1.17), displayed in Table 1.5, column (5) and data from RGGI coal plants.

1.6.4 Analysis and decomposition of RGGI coal plant auxiliary power consumption changes

In Tables 1.7 and 1.8, I present results of estimating equations (1.18) and (1.19), respectively. Each column presents coefficient results from adding one variable to the estimation, as shown, as well as a quadratic capacity factor term in models that include plant capacity factor. Results for the scrubber indicator variable in Table 1.7 indicate that scrubbers do lead coal plants to use more electricity on-site. Across specifications, coal plants use a little more than 1% of their electricity on-site when they have a scrubber. The impact of plants having an SCR unit is about half as large.

Table 1.7: *Results: Auxiliary power consumption model (with capacity factor)*

	Dependent variable:			
	Auxiliary power consumption (pct of gross generation)			
	(1)	(2)	(3)	(4)
Capacity factor (100 scale)			-0.109^{***}	-0.105^{***}
			(0.008)	(0.008)
Scrubber	1.716^{***}	1.475^{***}	1.132^{***}	0.876^{***}
	(0.220)	(0.228)	(0.182)	(0.178)
SCR unit		0.773^{***}	0.571^{***}	0.238^{*}
		(0.155)	(0.123)	(0.121)
Time trend				0.066^{***}
				(0.007)
Observations	34,952	34,952	34,952	34,952
Adjusted R^2	0.617	0.622	0.732	0.741

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating the coefficients of equation (1.18). All models have plant fixed effects, month fixed effects, and the coal price. Models that include plant capacity factor also include a quadratic capacity factor term. Standard errors clustered by plant.

Columns (3) and (4) show that, for a 1% decrease in total output, coal plants use about 0.1% more electricity on-site. This is statistically significant and reveals that coal plants' own

Table 1.8: *Results: Auxiliary power consumption model (with capacity online)*

	Dependent variable:			
	Auxiliary power consumption (pct of gross generation)			
	(1)	(2)	(3)	(4)
Percent plant capacity online			-0.036^{***}	-0.032^{***}
			(0.002)	(0.002)
Scrubber	1.716^{***}	1.475^{***}	1.285^{***}	0.891^{***}
	(0.220)	(0.228)	(0.189)	(0.183)
SCR unit		0.773^{***}	0.626^{***}	0.133
		(0.155)	(0.130)	(0.127)
Time trend				0.092^{***}
				(0.008)
Observations	34,952	34,952	34,952	34,952
Adjusted R^2	0.617	0.622	0.695	0.713

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating the coefficients of equation (1.19). Both models have plant fixed effects, month fixed effects, and the coal price. Standard errors clustered by plant.

electricity use is sensitive to output. If output drops by 10%, as it has for RGGI plants during the first four years of the RGGI program in reponse to the program (Fell and Maniloff, 2018), then plants will use 1% more electricity on-site, as if they installed scrubbers. Columns (3) and (4) also show that including plants' capacity factor in the model lowers coefficients on scrubbers and SCR units, indicating that the impact of these forms of emissions control are output-dependent.[47]

The time trend coefficient in column (4) is small but significant. For every year 2001-2016, US coal plants used 0.066% more electricity on-site. Over the entire time period, this is average change amounts to about 1%, as if all US coal plants each installed yet another scrubber. This stands in contrast with the negative time trend I find for coal plant gross heat rates, indicating that, if coal plants are reducing investments in efficiency, they may be the investments that impact electricity-consuming equipment at the plant.

Table 1.8 shows results of estimating equation (1.19), which includes the percent of a plant's total capacity that is online during the month instead of a plant's capacity factor. The results for the environmental control variables are broadly similar. The top row, which displays coefficient estimates for the percent of capacity online variable, shows that coal plants' auxiliary power consumption reduces as units as units have more hours running. Because RGGI plants have been shutting off their units during the RGGI period more so than other plant, this may have been a particular source of inefficiency for RGGI plants.

Results from the decomposition of RGGI coal plants' auxiliary power consumption are shown in Figure 1.10. Only 0.9% to 3.7% of the average annual auxiliary power consumption (across all plants) is unexplained by the model. The figure shows that generation changes have played the largest role in RGGI coal plants' auxiliary power consumption changes in the RGGI program period. Between 2009 and 2016, generation changes are associated with up to 84% of the change in predicted RGGI coal plants' auxiliary power consumption, between 2009 and subsequent years. In contrast, the presence of scrubbers—anticipated to be one of

[47] I also tried a specification in which I interacted the SCR indicator with an indicator for summer months, as NOx caps apply during the summer ozone season. Results did not appreciably differ.

the largest sources of inefficiency at plants—is responsible for only a maximum of 21% of the variation in auxiliary power consumption used at RGGI plants during this period, in terms of annual changes since 2009.

Figure 1.10: *Auxiliary power consumption decomposition results, RGGI coal plants*

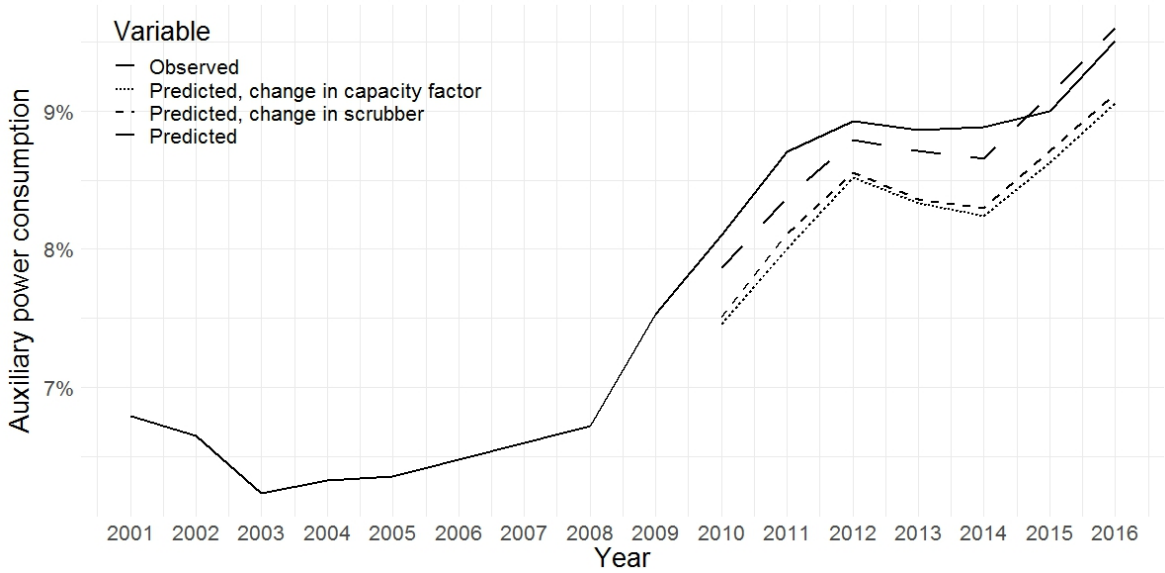

Note: The solid line is the average annual auxiliary power consumption (as a percent of gross generation) among RGGI coal plants. Each non-solid line represents the average annual predicted value of RGGI plants' auxiliary power consumption while keeping all variables at 2009 values and changing each variable labelled in the top left, sequentially, to the new year values (with each new year labelled on the x-axis). Therefore, the space between lines shows the contribution of each variable to the new year predicted average auxiliary power value. See equation (1.21) for decomposition methodology. Predicted values (conditional expectations for each year) are generated from estimating coefficients of equation (1.18) from the set of all US coal plants and using the data from RGGI coal plants. Coefficient values are in Table 1.7, column (4).

1.6.5 Did RGGI impact coal plant efficiency change, conditional on output changes?

Tables 1.9 and 1.10 display the results of estimating equation (1.22) with monthly and two-year data, respectively. The coefficient on the RGGI price in each tells the same story: Conditional on output changes, RGGI coal plant owners have not improved their heat rates in response to the RGGI price. Across specifications in both tables, the coefficient values indicate

that RGGI plants have become significantly less efficient by about 0.009% to 0.01% for every $1 increase in the RGGI price. Inclusion of a time trend in the model, which captures the aggregate change in US coal power plant heat rates from 2001-2016, does not appreciably change these estimates. (Compare columns (1) and (2) of each table.) The fact that the estimations with the monthly data and the two-year data provide similar coefficient values indicates that coal plant owners are neither improving the efficiency of their plants real-time with the carbon price nor making longer-term investments in new capital to reduce the cost of the carbon price to their operations. Given the average RGGI allowance price of $3.43 during these years and using the coefficient value from the estimation using two-year aggregated data and a time trend (Table 1.10, column (2)), RGGI coal plant heat rate rose about 0.034% in total during 2009-2016, over prior years, compared to heat rate changes in other US coal plants. Compared to Linn et al. (2014)'s estimates, this means that RGGI coal plants' heat rates have worsened about 0.23% compared to expectations for a carbon price.

Table 1.9: *Results: Did RGGI impact coal plant efficiency, conditional on output? (Monthly data)*

	Dependent variable: Log Net Heat Rate x 100	
	(1)	(2)
RGGI price (2010 USD per ton CO2)	0.929***	0.905***
	(0.203)	(0.199)
Log capacity factor (100 scale)	−4.370***	−4.194***
	(0.215)	(0.218)
Time trend		x
Observations	46,628	46,628
Adjusted R^2	0.750	0.751

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating the coefficients of equation (1.22). All models have indicators for whether a plant has a scrubber and whether a plant has an SCR unit, the coal price, plant fixed effects, and month fixed effects. Standard errors clustered by plant.

RGGI coal plant owners, forseeing low generation in future periods, could have chosen to forego investments in efficiency-improving measures and capital (Equation (1.10)). Another

Table 1.10: *Results: Did RGGI impact coal plant efficiency, conditional on output? (Two-year data)*

	Dependent variable: Log Net Heat Rate x 100	
	(1)	(2)
RGGI price (2010 USD per ton CO2)	0.979***	1.000***
	(0.335)	(0.336)
Log capacity factor (100 scale)	−5.276***	−4.836***
	(1.065)	(1.309)
Time trend		x
Observations	2,189	2,189
Adjusted R^2	0.873	0.874

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating Equation (1.22) with two-year aggregated data. Both models include indicators for whether plants have a scrubber or have an SCR unit, the coal price, and plant fixed effects. Standard errors are clustered by plant.

explanation may be found in Equation (1.9): the relationship between efficiency (e_t) and the marginal cost of efficiency improvements $\partial c(\cdot)/\partial e_t$ may be important. As coal plants reduce operation in the present term, their efficiency will mechanically fall. At a lower baseline efficiency (say, for example, a net heat rate of 11 instead of 10 mmBtu/MWh), the marginal cost of within-period efficiency improvements may be higher, offsetting the incentive of the carbon price for coal plant owners to improve their efficiency.[48]

The time trend in Table 1.9 also shows that, over 2001-2016, US coal plants have become significantly less efficient, by 0.012%. Given the negative trend I found in coal unit gross heat rates and the positive trend in coal plant auxiliary power consumption, this suggests US

[48]Interestingly, Linn et al. (2014) found that their coefficient estimate, the elasticity of coal unit net heat rates with respect to the coal price, trended upward over time. (See their Figure 4B.) Their estimate is not significantly different from zero in last four years of their panel, 2005-2009. The authors are unable to identify why this is is the case. This suggests that a conservative interpretation of my result is that RGGI coal plants' heat rates have worsened about 0.03% compared to expectations, given coal plants' observed reduced elasticity of net heat rates to input price changes in the years immediately preceding RGGI. Alternatively, my results can be interpreted as suggesting that coal plants are not improving their efficiency as would be expected by Linn et al. (2014)'s long-term average estimates of coal plant owners' responsiveness to input price shifts. Additionally, they corroborate the finding that coal plant owners in the US are not responding to input price changes with efficiency investments as they did prior to 2005.

plants as a whole are becoming less efficient due to using more power on-site, conditional on generation shifts and the environmental control equipment included here.

1.7 Estimation of a "rebound" effect

To calculate an emissions rebound effect from the efficiency loss I observe among coal plants in RGGI, I use my results from Section 1.6.5. The coefficient β_{CF} from estimating Equation (1.22) with two-year aggregated data is the average long-run elasticity of a coal plants' net heat rate with respect to its capacity factor. Table 1.10 shows that a 1% change in a coal plant capacity factor is associated with about a 0.048% increase in its net heat rate, over the long-run, conditional on the aggregate US trend in coal plant net heat rates (column (2), second row). I use this elasticity to estimate the heat rate change associated with RGGI coal plants' reduction in output in response to the program. I then calculate the difference in emissions between RGGI coal plants operating at this elevated heat rate and two counterfactual heat rates: 1) the pre-program average heat rate and 2) the anticipated average heat rate from coal plant owners responding to the RGGI carbon price. Since the RGGI price was approximately at the price floor for the first four years of the program, emissions abatement was not determined by the cap at that time but by the responsiveness of producers to the RGGI price. Therefore, the emissions rebound effect I calculate captures additional damages attributable to heat rate impacts and provides guidance on how a carbon tax might results in a similar rebound effect among coal plants.

In their analysis of the causal impact of RGGI on coal and natural gas generation in RGGI and neighboring states, Fell and Maniloff (2018) found that the program reduced RGGI coal plant capacity factors during the first four years of the program (2009-2012) compared to the five years preceding the program (2004-2008). Though the time period of this estimation does not extend to the end of my panel, RGGI coal plants' heat rates rose the most before 2012 (Figure 1.1). Using my data on net generation from coal-fired power plants in the region and excluding plants that retired, I calculate that RGGI coal plants reduced their total output from 291,795,335 MWhs to 262,615,801 MWhs in total between the four years before

(2005-2008) and after (2009-2012) the program, as a result of this RGGI-induced capacity factor drop.[49] This implies a 10% reduction in CO2 emissions between periods, of 31,904,671 tons, assuming that the pre-period coal plant net heat rate and CO2 intensity of coal used remained constant.[50] If I instead adjust the post-RGGI net heat rate of coal plants upward by 0.507%, as suggested by the elasticity of coal plant net heat rates with respect to capacity factors given in Table 1.10 for a 10% average change in capacity factors, I find the implied CO2 reduction between the RGGI pre-and post-periods is 30,448,828 tons. This is 4.56% less than that achieved without the heat rate change. The RGGI program induced 4.56% less of a CO2 reduction among coal-fired power plants than it would have if the plants' heat rates did not rise. This is a lower bound for the policy-induced emissions rebound effect from the efficiency loss I observe among RGGI coal plants, as it does not include expectations about how coal plant owners would adjust plant efficiency in response to the carbon price. Note also that, because this rebound effect is calculated from RGGI coal plants' net heat rate increase associated with policy-induced output changes, it does not capture additional policy-induced heat rate changes from changes in management practice, maintenance, or investment that were triggered by the program. In welfare terms, using a $31 per ton social cost of carbon (a central estimate for this time period), this emissions rebound effect equates to $45 million dollars in climate damages (Interagency Working Group on Social Cost of Greenhouse Gases, 2016).[51]

I use Linn et al. (2014)'s estimates of coal plants' anticipated net heat rate reductions from facing a carbon price to create an alternative, post-RGGI emissions counterfactual and upper bound for the rebound effect. Linn et al. (2014) estimate that a $10 per ton of CO2

[49]When I exclude retired plants from my sample, I find that the average pre-RGGI and post-RGGI capacity factor for coal plants in my data matches Fell and Maniloff's (2018). The average capacity factor in my RGGI pre-period data is slightly larger than theirs (0.52 vs. 0.50), likely due to the fact that my data is missing years 2006-2007. For this calculation, I calculate the mean annual net generation of RGGI plants over years 2004, 2005, and 2008 and multiply that annual average by four to obtain a ballpark regional production for the four years before the program, 2005-2008. An average capacity factor reduction of 10% to all RGGI plants for the post-RGGI period implies that total generation in the post-period dropped 10% due to RGGI.

[50]The observed CO2 rate for coal used in RGGI states ranges from 0.101 to 0.103 tons of CO2 per mmBtu, 2009-2016. The EPA data from which I observe this seems to have unreliable values before this date. I assume a constant CO2 rate of 0.102 for all rebound calculations.

[51]Note that the CO2 reduction considered here is not a total CO2 reduction acheived from the program, as it does not include changes in emissions from natural gas plant operation.

carbon tax should lead coal plants to reduce their heat rates by about 0.6%, conditional on their capacity factors. This estimate captures heat rate improvements from long-run investments in improved plant capital. The mean RGGI price in the years 2009-2012 was $2.26 per ton of CO2, which means that, if RGGI plants improved their heat rates according to Linn et al. (2014)'s estimates, they would have done so by 0.14% in these years. Such a heat rate reduction, with the 10% reduction in plant capacity factors, implies that RGGI plants would have emitted 32,306,670 fewer tons of CO2 in the years 2009-2012 compared to 2005-2008. This implies an emissions rebound effect is 5.75%. Note that coal plants could have foregone historically profitable, efficiency-improving investments in anticipation from future reduced output from RGGI and/or from the other factors impacting coal plants' decline in market share (particularly cheap natural gas). (See Equation (1.10).) So, this rebound effect captures the total rebound of emissions compared to reasonable ex-ante expectations of emissions reductions among RGGI coal plants. However, it may not be entirely caused by the program itself. Another caveat to this estimate is that Linn et al. (2014)'s coefficient estimates of the elasticity of coal plant net heat rates with respect to the coal price tended toward zero over the years in their panel, as discussed in Section 1.6.5 above. Therefore, this represents an upper bound in the rebound effect and assumes that coal plants would be as responsive to input price changes as historically observed, not as trends immediately prior to the program suggest.

1.8 Robustness checks

1.8.1 Is the unique RGGI net heat rate trend due to regulatory status?

A greater share of coal-fired power plants in RGGI are not regulated, compared to the rest of the US. Regulatory status determines a coal plant owners' incentives for operation, including efficiency changes. Compared to the model presented in Section 1.3, regulated plants may face less of an incentive to improve efficiency in response to changes in input or carbon prices, depending on their cost-of-service compensation set by their public utilities commission

(PUC). Alternatively, regulated plants often face specific efficiency incentives from their PUCs. Cicala (2015) shows that, post-deregulation, coal plants negotiated lower coal prices, indicating that, on average, regulated plants face higher coal prices but do not pursue cost-minimizing practices. Preonas (2019) shows that regulated plants are able to pass on costs associated with relative input price changes to their coal supplier, suggesting potential relative immunity for input price changes. Chan et al. (2017) show that, after deregulation, coal plants improved their efficiency.

Over the course of the RGGI program period, 81-85% of RGGI plants are deregulated, but only 11-16% of non-RGGI plants are. The two groups are heterogenous in terms of efficiency incentives. To test whether this drives my main result, I estimate Equation (1.11) for the subset of plant observations during which plants are deregulated. (The process of electricity restructuring in the US has been an ongoing process since the early 2000's.)

Results are presented in Table 1.11. They are broadly similar to the main results in Table 1.3, though the magnitude of the RGGI program coefficient has been reduced about 11% across all models.

Table 1.11: *Results: The regional difference in coal plant heat rates, deregulated plants only*

	Dependent variable: Log Net Heat Rate x 100					
	Entire sample			No planned retirement sample		
	(1)	(2)	(3)	(4)	(5)	(6)
RGGI plants, post-RGGI	4.693***	4.779***	4.761***	4.554***	4.625***	4.604***
	(1.177)	(1.245)	(1.256)	(1.141)	(1.213)	(1.223)
Scrubber		x	x		x	x
Coal price			x			x
Observations	8,107	7,594	7,594	7,971	7,467	7,467
Adjusted R^2	0.583	0.536	0.536	0.585	0.543	0.543

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating the coefficients of Equation (1.11) on the subset of plants that are deregulated, after deregulation. The "No planned retirement sample" excludes coal plants that have retired all units or have announced planned retirement for all units. All models include plant and month fixed effects. Standard errors are clustered by plant.

1.8.2 Simultaneity of coal plant efficiency and output

The output rate, or capacity factor, of coal plants impacts plant efficiency, due to the way that plants function. When estimating this relationship, one must consider that plant efficiency can also impact output. Coal plant owners in deregulated electricity markets provide bids to the independent system operator (ISO) based on their cost of generation, and the ISO decides their final hourly generation based on plants' production costs and other market-wide factors such as system reliability. Efficiency is one component of the marginal cost of generation. In regions in which the electricity supply is regulated and vertically-integrated, output decisions may be based on a number of additional market-wide factors, such as avoidance of wear-and-tear on the fleet of plants available to run. In both extremes, final output decisions are assigned on behalf of plants but are based in part on plants' heat rates.

To address the issue of simultaneity between coal plants' heat rates and capacity factors in my estimation, I estimate Equation (1.22) as described in Section 1.7 (with two-year aggregates of all variables) with natural gas prices as an instrument for coal plants' capacity factors. The natural gas price, the price of a competing fuel, is exogenous to coal plant efficiency and has dropped since 2008 due to the discovery of fracking technology to remove gas from previously inaccessible deposits in the US.[52] It is highly relevant to coal plant output as the major contributor to coal plants' reduced output, profits, and share in electricity production (Linn and McCormack, 2019; Coglianese et al., 2018) and should therefore be a strong instrument.

Table 1.12 shows the results from this estimation. The large values of the first stage F-statistics and small p-values for the first stage F-tests indicate that natural gas prices are a strong instrument for plants' capacity factors. The coefficient on the log capacity factor term in the first column is slightly larger than the comparable coefficient in the first column of Table 1.10, indicating that the simultaneous determination of coal plants' heat rates and

[52] One may be concerned that coal plant owners' expectations about future operation is one channel through which natural gas prices could impact coal plant efficiency outside of reducing their competitiveness in the wholesale market (and therefore capacity factors). I provide instrumental variables (IV) estimates here with my monthly data as well (which uses the annual state average natural gas price as an instrument), as plant-level changes in efficiency investments due to expectations created by natural gas prices are unlikely to have an impact in a one-year time frame.

capacity factors contributes to a small and unconcerning downward bias in the capacity factor coefficient. The log capacity factor coefficient in column (2) of Table 1.12 is about a third smaller than that in column (2) of Table 1.10. Given the fact that natural gas prices have a strong time trend over this period, this is unsurprising. Because the time trend and the instrumented log capacity factor term are multicollinear, I am unable to refine the coefficient I use for my rebound effect by using these IV results instead. (My preferred specification includes the time trend.) Still, column (1) of Table 1.12 is reassuring that this refinement may not be necessary.

Table 1.12: *IV Results: Simultaneity of coal plant efficiency and output (Two-year data)*

	Dependent variable: Log Net Heat Rate x 100	
	(1)	(2)
Log capacity factor (100 scale)	-5.155^{***}	-3.456^{***}
	(0.697)	(1.045)
Time trend		x
First stage F-stat	397.896	124.667
(p-value)	$(3.33e-80)$	$(4.83e-28)$
Observations	2,189	2,189
Adjusted R^2	0.873	0.872

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating Equation (1.22) with two-year aggregated data, as described in section 1.7, and the price of natural gas as an instrument for log capacity factor. Both models include indicators for whether plants have a scrubber or an SCR unit, the coal price, the RGGI price, and plant fixed effects. Standard errors are clustered by plant.

Table 1.13 shows the results from estimating Equation (1.22) with monthly data and the annual natural gas price as an instrument for plants' capacity factors. This is to check whether the IV results hold on heat rate changes with a shorter time scale than two years, in case the natural gas price is not exogenous to coal plant heat rates over a two-year time period. (That is, coal plant owners' changes in efficiency upgrades from expectations about future operation created by the natural gas price could have an effect on heat rates within two years.) Results are comparable to those in Table 1.9, for which I did not use an instrument. Broadly, the IV

results are similar to the non-IV results. Like the IV results obtained from the data aggregated to two-year time periods, the time trend competes with the variation in plant capacity factors created by the natural gas price, due to the time trend in natural gas prices. The fact that the coefficient on the log capacity factor term in column (1) of Table 1.13 is smaller than that in column (1) of Table 1.9 indicates that there may actually be an upward bias in the capacity factor term from the simultaneity of plant heat rates and capacity factors. This suggests that the log capacity factors term in column (1) of Table 1.12 (which is larger than that in Table 1.10) may be driven in part by changes in expectations created by the natural gas price. These differences in coefficients are not large, and overall the IV results suggest that simultaneity is not an issue in my estimation; coal plants' heat rates are driven by their output.

Table 1.13: *IV Results: Simultaneity of efficiency and output (Monthly data)*

	Dependent variable: Log Net Heat Rate x 100	
	(1)	(2)
Log capacity factor (100 scale)	-4.693^{***}	-3.441^{***}
	(0.593)	(0.789)
Time trend		x
First stage F-stat	2902.98	1019.915
(p-value)	$(0.00e+00)$	$(2.17e-221)$
Observations	46,628	46,628
Adjusted R^2	0.750	0.749

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating Equation (1.22) with monthly data and the price of natural gas as an instrument for log capacity factor. Both models include indicators for whether plants have a scrubber or an SCR unit, the coal price, the RGGI price, and plant fixed effects. Standard errors are clustered by plant. Models are the same as those in Table 1.9, except for the use of the instrument.

1.9 Conclusion

Where there is a change in input prices, there should, according to Hicks (1932), be a change in the direction of innovation. Innovative effort should be directed to reducing the cost of the new, more relatively expensive factor. The change in input prices created by a carbon

emissions price applied to the electricity sector offers a setting for testing this hypothesis and adding to the empirical literature on induced innovation. Additionally, understanding firm innovation responses to market-based environmental policy is important, as they drive the cost of compliance over time.

In this paper, I test whether coal plant owners in the Regional Greenhouse Gas Initiave (RGGI) have responded to the carbon price created by the program by making their plants more efficient, which would reflect several mechanisms of innovative effort. This response is anticipated from prior academic literature (Linn et al., 2014) and industry publications (Korellis, 2014). Through a difference-in-differences empirical design, I find that, during the RGGI program period, coal-fired power plants have become significantly *less* efficient than their peers in the rest of the US.

To uncover the mechanism of this unique trend, I estimate models of the two components of plant efficiency, gross efficiency (the conversion of coal into electricity) and auxiliary power consumption (the usage of electricity on-site). I then use estimated coefficients from these models to decompose the role of their determinants in efficiency changes at RGGI plants. I find that low output is responsible for the majority of changes in both gross efficiency and auxiliary power consumption at RGGI plants over the time period of the RGGI program. Specifically, I use a unique measure of plant output, an adjusted capacity factor, to capture the change in coal plant gross efficiency due to low output and find that RGGI coal plants' gross efficiency is being driven by the mechanical efficiency penalty that comes from operating at very low output rates. This indicates that plant owners are making a tradeoff between reducing generation and improving efficiency in response to input price changes. Low generation also explains the largest share (up to 84%) of changes in RGGI plants' on-site electricity consumption between 2009 and 2016. This implies that RGGI plants' low output has required them to use more electricity on-site as fraction of their total production, despite turning their units off for long periods of time. Notably, the presence of electricity-consuming equipment to control emissions of other pollutants, such as scrubbers to control sulfur emissions, are not largely responsible for RGGI plants using more auxiliary power. This demonstrates that RGGI plants are not

making a substantial efficiency tradeoff in complying with several environmental policies.

Due to the reduced efficiency associated with these output changes, the RGGI program achieved 4.56% less of the CO2 emissions reductions it would have achieved from coal plants if there were no efficiency loss, in its first four years. For policy analyses that estimate emissions reductions from carbon taxes, considering reduced generation from fossil-fired power plants alone does not account for reduced efficiency that occurs at these plants when they operate at low output rates. There is likely an emissions rebound effect for natural gas plants as well, given the fact that their efficiency is even more sensitive to output shifts than coal plants' (Lew et al., 2012). The path toward a clean electricity grid may be "dirtier" along the way.

Conditional on output changes, RGGI coal plant efficiency has not improved as would be predicted by prior work that captures the elasticity of coal plant heat rates with respect to input price changes (Linn et al., 2014). This suggests that, in addition to suffering efficiency loss from generating less, RGGI coal plant owners have not made efficiency investments that one would expect from a carbon price. As I show in my model, anticipated reductions in output in future periods reduces the incentive for improving plant efficiency. (See Equation (1.10)). Given prior literature on the determinants of coal plants' loss of market share (Linn et al., 2014; Linn and McCormack, 2019), plant owners may have expected lower output in the future from increased competition from natural gas plants. The RGGI price may have played a role in this expectation, given the wedge it drives between the cost of electricity generation from coal and natural gas. The program may have also served as a signal of more stringent climate policy in the future. The combined impact of reduced output due to the RGGI program and foregone efficiency improvement investments suggests an emissions rebound effect of 5.75%.

What do these results suggest for policy? As noted above, rebound effects from efficiency loss are important to consider for carbon emissions abatement anticipated from carbon pricing, particularly carbon taxes, applied to fossil-fired power plants. Additionally, these results point to contemporary federal climate policy, the Affordable Clean Energy (ACE) rule, being, at best, a third-best policy. As coal-fired power plants in RGGI have reduced emissions in response to the program, the revealed form of CO2 abatement has not been efficiency, or heat

rate, improvements. As plants have faced reduced generation due to increased competition from natural gas plants and other forces, the revealed course of action for coal plants has not been efficiency improvements. ACE targets coal-fired power plant heat rates. Aside from being a second-best (non market-based) policy in terms of design, it would require coal plant owners to deploy an abatement strategy that is evidently not cost-effective. In the long run, what matters most for reduced CO2 emissions is retirement of fossil-fired electricity generators. (In RGGI states, the combined impact of reduced generation and plant retirement in the years 2009-2016 has led to a 66% drop in CO2 emissions from coal.) One interesting extension to existing cost-benefit analyses of the ACE rule would be additional emissions induced from deferred retirement of coal plants due to forced plant efficiency upgrades.

What do these results imply with respect to the theory of induced innovation? Individual firms, or even sets of firms of a given type (here those that own coal-fired power plants) may not respond to changing factor prices in a way that seems intuitive, is suggested by industry experts, or even estimated in prior academic literature. RGGI coal plant owners found it more cost-effective to respond to the RGGI program by reducing generation, instead of improving efficiency. This paper is the first, to my knowledge, to document that firms able to respond to input price shifts along a number of margins may not exhibit an induced innovation response as anticipated. Additionally, the fact that RGGI coal plant owners did not improve their efficiency despite their output changes suggests support for the "market share" effect (Acemoglu, 2003), that innovation will be directed toward abundant factors in an economy. Note that I do not test this hypothesis rigorously and leave that endeavor to future work. The nature of innovative firm behavior in declining industries is particularly salient for climate policy: a shift away from CO2 intensity in many sectors will inevitably involve industry transition.

Chapter 2

What do we lose by picking winners? The role of technology-specific clean energy incentives in induced innovation

2.1 Introduction

A critical component of cost effective climate policy is incentivizing innovation among emissions-intensive sectors. A carbon price can accomplish this task, absent other externalities, in addition to its primary role in incentivizing emissions reduction (Milliman and Prince, 1989). In practice, however, policymakers often overlap clean energy policies on top of a carbon price. Examples include the state-level Renewable Portfolio Standards in the US and feed-in-tariffs in the EU, which are layered on top of, respectively, various US carbon pricing endeavors such as the Regional Greenhouse Gas Initiative (RGGI) and the EU Emissions Trading System.

If a carbon price is being used to address emissions externalities, these clean energy policies are not justified unless they correct additional market failures. They are often put into place on innovation-inducing grounds, the idea being that new, expensive renewable technologies are

not competitive in a world with initial carbon prices that are, by most social cost of carbon calculations, small. However, our current understanding of market failures specific to various clean energy technologies is in its infancy. At the same time, overlapping these policies can lead to static costs, as expensive forms of emissions abatement (subsidized renewables) replace the sources of emissions abatement that would occur with only a carbon price in place. Moreover, they may undermine the endogenous carbon price set in a cap-and-trade system Fankhauser et al. (2010). In addition to static costs, do these overlapping policies have dynamic costs? For example, do they redirect the path of induced innovation away from the suite of technologies that would be realized under first-best policy? This project aims to answer this question. More specifically, this project aims to characterize the total, additional abatement cost created by common technology-specific ("tech-specific") clean energy subsidies, including any dynamic costs of re-directing innovation.

I first answer this question with a two-period, partial equilibrium model of the electricity sector that incorporates first-best policy. (That is, I include a carbon price and assume knowledge spillovers and other clean energy market failures have already been corrected.) I look at the additional cost incurred in both a static and dynamic sense when the level of support for particular technologies is beyond optimal. The model includes both learning-by-doing (LBD) and private research and development (RD) as sources of innovation. I look at the innovation tradeoffs the electricity sector makes between emissions-intensive ("dirty"), low-emissions ("clean") and unsubsidized, and clean and subsidized technologies when a tech-specific clean output subsidy is put in place.

I find that, in the case of a carbon tax, the electricity sector directs more research effort toward subsidized clean technologies (also "fuels") if the returns to innovation are increasing in the quantity of fuel used. (Innovation here is captured by the second-period reduction in generation cost). Similarly, the sector shifts innovation effort away from unsubsidized fuels if the returns to innovation are increasing in scale. Analytical results for the carbon tax case are therefore ambiguous and reflect recent work showing that innovation incentives are from a combination of input price changes and their resulting impact on input substitution Acemoglu

(2003). In the case of a cap-and-trade system, the carbon price responds endogenously to the emissions intensity of the input mix that producers choose. Therefore, production subsidies for clean fuels cause the carbon price to fall, and we anticipate producers to devote less innovation toward emissions abatement.

I use a simulation of the electricity sector to obtain more concrete results for the tax case and explore results by carbon price and subsidy magnitudes. A simulation also allows me to examine changes in production and research choices according to emissions-intensity of dirty fuels, particularly because switching between coal and natural gas is a common, cost-effective form of emissions abatement in the electricity sector. (Coal is more carbon intensive than natural gas.) I calibrate my simulation to the EU electricity sector, a context in which policy overlap has been a long-standing issue. I add two novel features to the simulation: I incorporate exogenous technical change. I allow dirty technology to experience cost-saving emissions reductions from innovative effort.

I find that, in the case of a carbon tax, small subsidies (on the order of 10% of the cost of production from clean sources) decrease research spending for subsidized clean technology for small carbon taxes ($10 - $20 per ton of CO2). (I abstract from LBD and RD in the simulation and quantify innovation changes through "research spending", which lowers future production or emissions cost.) A subsidy reduces the need to reduce technology cost to allow for fuel substitution. Larger subsidies (on the order of 50% of clean technology production cost) cause research spending for dirty technology to be reduced when carbon taxes are higher ($30-$40 per ton of CO2). When the sector must pay for emissions and would normally invest research to develop cleaner technology, subsidies reduce the need for this research. Very high subsidies (100% of the cost of clean technologies) incentivize higher research spending on subsidized, clean technologies and/or reduce the incentive for research spending on dirty technology (which can be interpreted as developing new clean technology), for carbon taxes in the range of $30-$50 per ton of CO2.

A subsidy for clean technology in the presence of a cap-and-trade system reduces the incentive for research, generally. There is less of an incentive to reduce the cost of clean

technology or the emissions rate of dirty technology. Interestingly, this effect dissipates in subsidy size, implying that, when large subsidies are adopted in a cap case, it is more cost-effective, on the margin, to abate emissions through research spending than through fuel-switching. This may be due to the fact that subsidies have forcibly caused fuel-switching to a large extent.

This paper contributes to the discussion over how climate policies incentivize innovation, which is important for long-run, cost-effective emissions reduction. Goulder and Schneider (1999), Goulder and Mathai (2000), Acemoglu et al. (2012), and Lemoine (2017) characterize optimal climate policy, endogenizing technical change, with general equilibrium models. While this paper is concerned with long-term optimality, it is closer in spirit to recent work that looks at the implications of carbon pricing and research-oriented policy tradeoffs in the electricity sector specifically (Fischer and Newell, 2008; Fischer et al., 2017; Hubler et al., 2015; NAS, 2016). Fischer et al. (2017) suggest using very small tech-specific subsidies to correct for LBD spillovers with clean technology as well as RD subsidies that are 50% of clean technology revenues, along with carbon pricing. In a model similar to theirs, I include exogenous technical change and ask: Given rather high clean energy subsidies we have seen for example, in the EU, how much does innovation deviate from optimal?[1]

Section 2.2 below presents the static cost of overlapping tech-specific clean energy subsidies on top of a carbon price in the framework I plan to use for assessing additional, dynamic costs. Section 2.3 presents my analytical model. I first evaluate the effect of overlapping clean energy subsidies with a carbon tax on innovation outcomes. I then conduct the same analysis with an endogenous carbon price. Section 2.4 describes my simulation, and in Section 2.6 I presents my simulation results. I check the sensitivity of my results to assumptions over the discount rate in Section 2.7. In Section 2.8, I conclude.

[1]By both including exogenous technical change and adjusting other innovation parameters to account for its presence, research spending is less impactful in my model than in Fischer et al. (2017). Fischer et al. (2017) find optimal LBD subsidies of 0.70 cents/kWh for wind and 4.93 cents/kWh for solar, which they compare to observed median EU feed-in-tariffs at the USD equivalent of 6.5 and 33.8 cents/kWh.

2.2 The static cost of tech-specific subsidies

Consider an electricity-producing firm that has a profit-maximizing objective function as follows:[2]

$$\max_{\mathbf{q}} \pi = (\mathbf{p} + \mathbf{s})\mathbf{q} - c(\mathbf{q}) - p_c e^D q^D \tag{2.1}$$

where

$\mathbf{p}, \mathbf{q} > 0$

The firm chooses the production vector $\mathbf{q}$ that contains elements corresponding to the amount of electricity that it will produce from carbon intensive "dirty" fuels (D), unsubsidized "clean" (carbon-free) fuels (UC), and subsidized clean fuels (SC): $\mathbf{q} = \left(q^D q^{UC} q^{SC}\right)$. Electricity generated from all forms receive the same wholesale electricity price p, so the vector $\mathbf{p} = (p\ p\ p)$. However, only the subsidized clean form of energy receives a subsidy, so the subsidy vector $s = (0\ 0\ s^{SC})$, where $s^{SC} > 0$ is a positive subsidy for the input SC. The firm faces a cost of production, $c(\mathbf{q})$. It also faces a carbon price, p_c, so the cost of complying with the carbon price after making input adjustments is $p_c e^D q^D$, the carbon price times the emissions rate of the dirty fuel, e^D, and the quantity of electricity produced from dirty fuel, q^D.

The firm determines production by the following first order conditions (FOCs), obtained by differentiating Equation (2.1) with respect to each element of the output choice vector $\mathbf{q}$:

$$p = c'\left(q^D\right) + p_c e^D \tag{2.2}$$

$$p = c'\left(q^{UC}\right) \tag{2.3}$$

[2]Like Fischer et al. (2017), I take this to be a reasonable assumption, as the sector is trending toward deregulation over time. This model can be used to examine the behavior of an individual firm or the sector as a whole, as Fischer et al. (2017) do. Note, as Fischer et al. (2017), I also assume perfect competition and abstract from market power issues. In the European setting, in which my simulation is based, regional grid interconnection has promoted market competition.

$$p + s^{SC} = c'\left(q^{SC}\right) \tag{2.4}$$

Before the subsidy, the producer choses production from each input to satisfy the following, which is a combination of Equations (2.2)-(2.4), assuming $s^{SC} = 0$:[3]

$$c'\left(q^{D}\right) + p_c e^{D} = c'\left(q^{UC}\right) = c'\left(q^{SC}\right) \tag{2.5}$$

After the subsidy is put in place, Equation (2.5) becomes:

$$c'(q^{D}) + p_c e^{D} = c'(q^{UC}) = c'(q^{SC}) - s^{SC} \tag{2.6}$$

Quite intuitively, the subsidy lowers the marginal cost of producing from SC relative to other fuels. To the extent that these functions are smooth, producers will produce more from SC and less from D and UC in the short term. If the optimal $\mathbf{q}^* = \{q^{*D}\ q^{*UC}\ q^{*SC}\}$ determined by Equation (2.5) has changed to $\mathbf{q}^{**} = \{q^{**D}\ q^{**UC}\ q^{**SC}\}$ via Equation (2.6), after the subsidy, the change in production cost for one firm that takes the subsidy is below in Equation (2.7). If additional abatement is achieved by this production shift, Equation (2.7) captures the cost of this abatement.

$$\int_{q^{SC}=q^{*SC}}^{q^{**SC}} [c'\left(q^{SC}\right)]dq^{SC} + \int_{q^{D}=q^{*D}}^{q^{**D}} \left[c'\left(q^{D}\right) + p_c e^{D}\right] dq^{D} + \int_{q^{UC}=q^{*UC}}^{q^{**UC}} \left[c'\left(q^{UC}\right)\right] dq^{UC} \tag{2.7}$$

Note that the first term is expected to be positive, as the firm is producing more from SC than before. The other two terms may be negative as the firm shifts away from using those fuels. For simplicity, Equation (2.7), the net cost of abatement incurred by the subsidy, can

[3] Note that this model abstracts from switching costs. Instead of determining the moment-by-moment production decisions of an electricity generator, Equation (2.5) determines what a firm chooses over the course of a longer time period, such as a year.

be collapsed into:

$$\int_{\mathbf{q}=\mathbf{q}^{*}}^{\mathbf{q}^{**}} [c'(\mathbf{q})]dq \tag{2.8}$$

If p_c is generated from a cap-and-trade system, then Equation (2.8) captures the additional cost of abatement from the subsidy. If p_c is a carbon tax, then we anticipate additional abatement from the introduction of the subsidy, which yields a social benefit.[4,5] For the carbon tax case, we need to therefore compare Equation (2.8) to how that abatement would be achieved in a cost-effective manner. Let an increase in the carbon tax p_c that leads to the same about of abatement as the subsidy s^{SC} be denoted by p_c'. This new carbon price would lead to a new production vector, $\mathbf{q}'$. The additional cost of abatement from p_c' can be represented by:

$$\int_{\mathbf{q}=\mathbf{q}^{*}}^{\mathbf{q}'} [c'(\mathbf{q})]dq \tag{2.9}$$

The static cost of achieving more abatement using the subsidy s^{SC} is therefore the difference between Equation (2.8) and Equation (2.9), the difference between using the tech-specific subsidy and using the carbon tax:

$$\int_{\mathbf{q}=\mathbf{q}^{*}}^{\mathbf{q}^{**}} [c'(\mathbf{q})]dq - \int_{\mathbf{q}=\mathbf{q}^{*}}^{\mathbf{q}'} [c'(\mathbf{q})]dq \tag{2.10}$$

The value of Equation (2.10) will always be positive, unless a firm, when faced with a

[4] If p_c is endogenous as in a cap-and-trade system, then we would not expect the subsidy to lead to additional abatement, as the increase in production from SC yields carbon abatement that replaces abatement efforts that would otherwise be undertaken. If p_c is a carbon tax, we anticipate that s^{SC} would lead to additional abatement as compared to the no-subsidy case, as producers would be incentivized to produce more from SC—a no-carbon source—than before and continue to face a tax on remaining production from D.

[5] I am not taking a stance on the nearness of the carbon price to the appropriate social cost of carbon (SCC). If $p_c = SCC$, then additional abatement, no matter how conducted, would not be welfare-improving. If $p_c < SCC$, then additional abatement achieved through a technology-specific subsidy may or may not be welfare-improving, depending on the cost of using the technology. The additional (desirable) abatement is, however, more expensive than it needs to be.

more stringent carbon tax, would rely entirely on SC when faced with a subsidy. This is the static cost of tech-specific subsidies. What is the dynamic cost?

2.3 A simple model for the innovation cost of clean subsidies

The previous section characterizes the static change in cost associated with overlapping technology-specific clean energy subsidies with a carbon price. This static cost captures part of the total cost incurred by this overlapping policy. I use the following dynamic model to explore an additional cost: the shift of innovative effort away from the suite of carbon-saving technologies that would be chosen under a sufficiently high carbon price and toward the suite chosen when tech-specific subsidies are overlapped with that carbon price.

I consider a large set of electricity generating firms that each solve the following two-period objective function.

$$\max_{\mathbf{q}_1,\mathbf{q}_2,\mathbf{r}} \pi = [(p_1\mathbf{q}_1 - c_1(\mathbf{q}_1,\mathbf{r})] + \delta[(p_2+\mathbf{s})\mathbf{q}_2 - c_2(\mathbf{q}_2,\mathbf{K},p_c)] \tag{2.11}$$

where

$\mathbf{K} = \mathbf{K}(\mathbf{q}_1,\mathbf{r})$

$p_1, p_2, \mathbf{q}_1, \mathbf{q}_2, \delta > 0$

The firm maximizes the net present value of profit over both periods by choosing a production quantity vector for each period $\{\mathbf{q}_1, \mathbf{q}_2\}$ and R&D investment vector $\mathbf{r}$ in period one, conditional on the exogenous wholesale electricity price for each period, p_1 and p_2; its cost functions, $c_1(\cdot)$ and $c_2(\cdot)$; the carbon price in period two, p_c; the subsidy vector $\mathbf{s}$; and discount rate δ. As above, vectors, denoted in bold, contain elements referring to the three fuel types: carbon intensive "dirty" fuels (D), unsubsidized clean fuels (UC), and subsidized clean fuels (SC). That is, $\mathbf{q}_t = \left(q^D\ q^{UC}\ q^{SC}\right)_t$ and $\mathbf{s} = (0\ 0\ s^{SC})$.

In this set-up, future carbon pricing will induce a firm to abate emissions through using less dirty fuel, substitute toward cleaner fossil fuels, innovate to reduce the emissions rate of

dirty fuel D, and/or innovate to bring down the cost of clean fuels UC and SC. I compare the case of $s = 0$ to the case of $s = (0\ 0\ s^{SC})$ where $s^{SC} > 0$ to characterize how a tech-specific clean subsidy for SC, s^{SC}, impacts the firm's innovation decisions.

The firm's choice of production, $\mathbf{q}_1$, and R&D effort, $\mathbf{r}$, in period one impact second-period cost via their contribution to the knowledge stock $K(\mathbf{q}_1, \mathbf{r})$.[6] Specifically, for D, the knowledge stock from R&D impacts the emissions rate and therefore the cost of complying with the emissions price, p_c. This is how induced innovation for carbon intensive fuel is captured in the model: the firm is able to respond to the carbon price by conducting R&D that impacts its second-period emissions rate. I assume there is no learning-by-doing (LBD) for dirty fuels, given how far out on the experience curve most electricity systems are with fossil fuels. For UC and SC, K generated from both experience—first-period production—and R&D lowers their *production* costs in period two. Clean energy benefits from both learning-by-doing (LBD) and intentional R&D effort. The carbon tax does not apply to production from these fuels.

For D, the second period cost of generation and cost of complying with the carbon price are separable. I assume that the cost of complying with the carbon price takes the form:

$$c_2\left(p_c\right) = p_c e^D\left(K^D\right) q_2^D \tag{2.12}$$

where $e^D(K^D)$ is the emissions rate of D, a function of the second-period knowledge stock related to D, K^D. The quantity of D used in period two is q_2^D.

First-best carbon policy is included by assuming p_c equals the relevant social cost of carbon.[7] First-best innovation policy is included by not modeling spillovers in Equation (2.11); that is, I assume all positive external returns from $\mathbf{r}$ are able to be captured by the innovating

[6] Assuming relevant competitive markets and correction of knowledge externalities, r also captures energy R&D conducted outside the firm and supplied to the firm. In this way, the model characterizes economy-wide shifts in R&D for electricity generation.

[7] I recognize first-best carbon policy would be upstream, applied to fuels entering the economy instead of producers in the markets who generate emissions. I model a carbon price being applied to the electricity generation market given the popularity of this policy choice, although I could certainly modify the model to instead include an input tax on D.

firm.

For the analysis below, it is important to establish assumptions about the derivatives of the functions above. I assume that generation cost functions for all fuels are increasing and convex in production: $c_t'(q_t^f) > 0$ and $c_t''(q_t^f) > 0$ for $f \in \{D, UC, SC\}$. I assume that the cost of R&D effort is also increasing and convex for all technologies: $c_t'(r_t^f) > 0$ and $c_t''(r_t^f) > 0$. I assume that the knowledge stock $K(\mathbf{q}_1, \mathbf{r})$, a function of the first period production and R&D effort vectors, is increasing and concave in both of these arguments: $\frac{\partial K}{\partial \mathbf{q}_1} > 0$ (except for D, as noted above), $\frac{\partial^2 K}{\partial \mathbf{q}_1^2} < 0$, $\frac{\partial K}{\partial r} > 0$, and $\frac{\partial^2 K}{\partial r^2} < 0$. These forms capture both a positive impact of experience and R&D effort on cumulative knowledge, with decreasing marginal returns. An exception is that production does not impact the knowledge stock for dirty fuels: $\frac{\partial K^D(q_1^D, r^D)}{\partial q_1^D} = 0$ where the superscript D refers to all terms relevant to D. Finally, as discussed above, I assume that an increase in the knowledge stock decreases the emissions rate for D and the production cost for UC and SC in period two: $e^{D\prime}(K^D) < 0$ and $\frac{\partial c_2^g(q_2^g, K^g)}{\partial K^g} < 0$ for $g \in \{UC, SC\}$.

Taking derivatives with respect to each element of the firm choice vectors and denoting these elements by superscripts D, UC, and SC, the relevant first order conditions for the firm problem are:

Period one production:

$$p_1 = c_1'(q_1^D) \tag{2.13}$$

$$p_1 = c_1'(q_1^{UC}) + \delta \left(\frac{\partial c_2(q_2^{UC}, K^{UC})}{\partial q_1^{UC}} \right) \tag{2.14}$$

$$p_1 = c_1'\left(q_1^{SC}\right) + \delta \left(\frac{\partial c_2(q_2^{SC}, K^{SC})}{\partial q_1^{SC}} \right) \tag{2.15}$$

R&D choice:

$$c_1'(r^D) = -\delta\left(p_c \frac{\partial e^D(K^D)}{\partial r^D} q_2^D\right) \tag{2.16}$$

$$c_1'(r^{UC}) = -\delta\left(\frac{\partial c_2(q_2^{UC}, K^{UC})}{\partial r^{UC}}\right) \tag{2.17}$$

$$c_1'(r^{SC}) = -\delta\left(\frac{\partial c_2(q_2^{SC}, K^{SC})}{\partial r^{SC}}\right) \tag{2.18}$$

Period 2 production:

$$p_2 = c_2'(q_2^D) + p_c e^D(K^D) \tag{2.19}$$

$$p_2 = \frac{\partial c_2\left(q_2^{UC}, K^{UC}\right)}{\partial q_2^{UC}} \tag{2.20}$$

$$p_2 + s^{SC} = \frac{\partial c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_2^{SC}} \tag{2.21}$$

Equations (2.13)–(2.15) show that the firm chooses period one production by equating marginal revenue (the wholesale electricity price p_1) with the marginal cost of production ($c_1'\left(q_1^{UC}\right)$ for UC, for example) as well as the net present value of the impact of LBD for UC and SC ($\delta(\frac{\partial c_2(q_2^{UC}, K^{UC})}{\partial q_1^{UC}})$ for UC). More specifically, $\frac{\partial c_2(q_2^{UC}, K^{UC})}{\partial q_1^{UC}}$ is the change in the period two production cost that occurs from period one production impacting the knowledge stock for UC.

The conditions for the firm's R&D decisions are described by Equations (2.16)-(2.18). The

firm chooses an amount of investment in R&D effort for each technology such that the marginal cost of R&D effort ($c_1'\left(r^D\right)$ for D) equals the net present value of the returns to innovation in the form of period two cost reduction (compliance cost reduction for D, $p_c \frac{\partial e^D(K^D)}{\partial r^D} q_2^D$, and generation cost reduction for the clean sources, such as $\frac{\partial c_2(q_2^{UC}, K^{UC})}{\partial r^{UC}}$ for UC).

Period two production decisions are described by Equations (2.19)–(2.21); for each, the firm equates marginal cost and revenue. For the decision regarding the amount to produce from dirty fuel, the marginal emissions cost is considered with the marginal production cost (Equation (2.19)). The firm's decision regarding q_2^{SC} is influenced by the subsidy for SC, s^{SC}, which increases revenue for clean fuels. Production of electricity from each fuel is decided jointly with period one production and R&D effort, as the latter affects the marginal cost of emissions control for D and each affect the marginal cost of production for UC and SC, via their influence on K.

2.3.1 Carbon tax case

I now evaluate how tech-specific subsidies impact the innovation effort devoted to each energy input when the carbon price is a tax.

The Impact of a Tech-Specific Subsidy on Induced Innovation for that Technology

How does s^{SC} impact the firm's innovation decisions for SC, compared to the case in which there is no subsidy? Since I assume $\frac{\partial c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_2^{SC}} > 0$, adding a subsidy when there formerly was not a subsidy in place leads the firm to use more SC—the desired result for a production subsidy (Equation (2.21)). This increase in second-period production from the subsidized form of clean energy may effect innovation in two ways: 1) leading the firm to change its first-period production from SC, altering the LBD benefits that accrue to the technology, and/or 2) leading the firm to change its first-period R&D choice for SC. Both changes constitute a re-direction of induced innovation.

We can examine whether the subsidy will impact q_1^{SC} by rewriting Equation (2.15),

totally differentiating the rightmost term:

$$p_1 = c_1'\left((q_1^{SC})\right) + \delta(\frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{(\partial q_1^{SC})^2} dq_1^{SC} + \frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_1^{SC} \partial q_2^{SC}} dq_2^{SC} + \frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_1^{SC} \partial r^{SC}} dr^{SC}) \tag{2.22}$$

Totally differentiating the rightmost term allows us to see how the returns to LBD vary in second-period production (via its second term, $\frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_1^{SC} \partial q_2^{SC}} dq_2^{SC}$) and in R&D (via its last term, $\frac{\partial c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_1^{SC} \partial r^{SC}} dr^{SC}$). It is reasonable to assume that $\frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_1^{SC} \partial q_2^{SC}} \leq 0$, that the returns to LBD are constant or increasing in second-period production.[8,9] If they are increasing in scale, then, because $dq_2^{SC} > 0$ in the presence of the subsidy, we anticipate the firm to increase first-period production from SC in response to the subsidy.

The third term in the totally differentiated representation of returns to LBD, $\frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_1^{SC} \partial r^{SC}} dr^{SC}$, captures whether LBD and R&D are complements or substitutes ($\frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_1^{SC} \partial r^{SC}} < 0$ or $\frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_1^{SC} \partial r^{SC}} > 0$, respectively; keep in mind that a cost reductions constitute innovation here). This becomes relevant if the firm changes its R&D in response to the subsidy ($dr^{SC} \neq 0$).

By looking at Equation (2.18), we see that the increase in q_2^{SC} may influence the firm's decision related to R&D devoted to SC. Totally differentiating this term, which represents the returns to R&D, we can express Equation (2.18) as:

$$c_1'\left(r^{SC}\right) = -\delta\left(\frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial r^{SC} \partial q_1^{SC}} dq_1^{SC} + \frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial r^{SC} \partial q_2^{SC}} dq_2^{SC} \frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{(\partial r^{SC})^2} dr^{SC}\right) \tag{2.23}$$

Equation (2.23) shows that the firm chooses to invest in R&D for SC by equating the marginal cost of R&D effort with the net present value of the second period marginal returns, which include how marginal returns vary with the first period production decision ($\frac{\partial^2 c_2\left(q_2^{SC}, K^{SC}\right)}{\partial r^{SC} q_1^{SC}} dq_1^{SC}$) and how they vary with the second period production decision

[8] Learning by doing will likely impact all of future production; it seems unlikely that cost reductions due to experience would decline after a certain amount of future production.

[9] Note that returns to innovation in this context imply a decrease in second-period cost.

$(\frac{\partial^2 c_2(q_2^{SC},K^{SC})}{\partial r^{SC} q_2^{SC}} dq_2^{SC})$. More intuitively, the first term shows that LBD affects the R&D effort choice by nature of whether it serves as a complement or substitute to R&D effort. The second term reflects how R&D effort is impacted by the new production induced by the subsidy, on the margin.

Looking first at the second term, if we assume that R&D effort brings down the marginal cost of production (i.e. R&D returns are increasing in production), then the subsidy will indeed lead to an increase in R&D in response to a greater-than-otherwise production from SC. The cross derivative $\frac{\partial^2 c_2(q_2^{SC},K^{SC})}{\partial r^{SC} \partial q_2^{SC}}$ will be negative, making the entire right hand side of Equation (2.23) more positive, as dq_2^{SC} is positive in response to the subsidy. Any R&D that has greater returns as SC is scaled up, such as storage research that is more profitable for large amounts of clean fuels feeding into the grid, serves as an example of this case.[10] If, however, the future cost reduction from R&D effort is zero (or negative) in quantity produced—say, for example, R&D on the efficiency of solar panels affects the total but not marginal production cost as the cost reduction from the next panel is the same as the previous—then no more R&D effort may be expended in response to the increase in the production of SC from the output subsidy.

The firm's choice of R&D could also be impacted by a tech-specific subsidy through the first term in the returns to innovation in Equation (2.23), $(\frac{\partial^2 c_2(q_2^{SC},K^{SC})}{\partial r^{SC} \partial q_1^{SC}} dq_1^{SC})$. This term captures whether first-period production changes in response to s^{SC} (through dq_1^{SC}) and whether LBD and R&D are substitutes, complements, or neither (through $\frac{\partial^2 c_2(q_2^{SC},K^{SC})}{\partial r^{SC} \partial q_1^{SC}}$). The analysis above indicates that $dq_1^{SC} > 0$ if returns to LBD are increasing in scale. In that case, if LBD and R&D are substitutes, the increase in q_1^{SC} would offset at least some of any potential increase in r^{SC} in response to the subsidy. If they are complements, then the firm will expand its R&D for SC as it scales up first-period production.

In summary, a future production subsidy for a form of clean energy will induce a firm

[10] This case is supported by the functional forms commonly assumed in the literature for renewable production cost functions and knowledge production. For example, Fischer et al. (2017) assume a quadric production cost function that is inversely related to the knowledge stock. In their model, the marginal cost function is decreasing in the knowledge stock.

in a competitive wholesale electricity market (in which there are already policy measures that correct the environmental and knowledge externalities) to: 1) increase second-period production of that clean energy, 2) increase first-period production of that energy if the gains from LBD are increasing in the scale of second-period production, and 3) increase R&D effort toward that form of energy if the gains to R&D are increasing in second-period production. The magnitude of (2) and (3) each depend on the extent to which LBD and R&D are substitutes or complements.

The Impact of a Tech-Specific Subsidy on Induced Innovation for Other Technologies

In the section above, I identified the conditions under which the firm would increase its innovative efforts toward a subsidized form a clean energy. I now examine whether a firm will be incentivized to reduce its innovative effort for other forms of energy when a technology-specific subsidy is put in place.

I start by looking at how s^{SC} impacts production from D and UC. We can combine Equations (2.19)–(2.21) to summarize the firm's second-period production decision:

$$c_2'\left(q_2^D\right) + p_c e^D\left(K^D\right) = \frac{\partial c_2\left(q_2^{UC}, K^{UC}\right)}{\partial q_2^{UC}} = \frac{\partial c_2\left(q_2^{SC}, K^{SC}\right)}{\partial q_2^{SC}} - s^{SC} \tag{2.24}$$

Here I turn from considering the objective function (Equation (2.11)) as depicting the firm's problem to capturing the entire electricity industry's problem. To the extent that producers are heterogeneous in their production costs and the fuel mix from which they produce, (24) will look very different for each producer and result in different fuel mix production decisions when the subsidy is put in place. If we think of (11) as encompassing the entire industry—a summation of the objective functions for all producers—then (24) allows us to examine aggregate, industry-level input choices. Importantly, it allows us to look at input choice shifts for D and UC, assuming that electricity demand is not affected by s^{SC}.

Looking at the far right hand side, we see that the subsidy reduces the marginal cost

of production of the subsidized fuel relative to other fuels. The industry will produce more from SC. Because firms equate the marginal costs of production from all of their fuels via the second-period wholesale electricity price, they will therefore scale back production from UC as they produce more from SC. They will also scale back production from D or switch inputs for production from D (from more expensive gas to less expensive coal, for example), thereby lowering $c_2'\left(q_2^D\right)$ and altering the emissions rate from dirty fuel, $e^D\left(K^D\right)$, to satisfy Equation (2.24). This is done according to the marginal cost of each fuel and, importantly for D, the carbon price.

Equation (2.16), copied below, shows how R&D for D changes as the industry uses less D or changes the fuel mix for production from D.

$$c_1'(r^D) = -\delta\left(p_c \frac{\partial e^D(K^D)}{\partial r^D} q_2^D\right) \tag{2.16 revisited}$$

Since q_2^D enters Equation (2.16) directly, we anticipate R&D directed toward emissions abatement to fall as q_2^D falls. (Keep in mind that $\frac{\partial e^D(K^D)}{\partial r^D}$ has a negative sign.) If producers switch dirty fuels but maintain production from D (q_2^D does not change), then shifts in innovation depend on how the returns to innovation for emissions abatement, $\frac{\partial e^D(K^D)}{\partial r^D}$, change with the new fuel mix. For example, if producers switch to using more of a dirty fuel for which research is more effective at bringing down the emissions rate of that fuel than the previously ($\frac{\partial e^D(K^D)}{\partial r^D}$ becomes more negative), then we anticipate firms to increase emissions-abatement research. Alternatively, it's plausible that research may be less effective at bringing down the emissions rate of the new fuel, leading firms to reduce innovative effort for emissions abatement.

As second-period energy production from UC is scaled back, then changes in first-period production, which accrue LBD benefits, and R&D are determined by Equations (2.14) and (2.17). By totally differentiating the second-period cost function in each and conducting the same analysis as the section above, we would see that the conditions under which innovative effort for SC increases in response to greater production from SC are also the conditions under

which innovative effort for UC decreases when q_2^{UC} declines. Specifically, we would anticipate q_1^{UC} to fall if LBD returns for UC increase in second-period production. We would anticipate r^{UC} to fall if returns to R&D are increasing in second-period production. Each effect is adjusted by the substitutability or complementarity of LBD and R&D.

In sum, we can anticipate a tech-specific clean energy subsidy to reduce the R&D effort that the electricity sector devotes to dirty fuel if firms produce less electricity form dirty fuel. If firms produce the same amount of electricity from dirty fuel but switch fuels, then R&D effort will change with the returns to innovation—how the emissions intensity of fuel changes according to innovation—for the new fuel mix. The effect of a tech-specific subsidy on innovation directed toward unsubsidized clean fuel is conditional on innovation returns, as described immediately above.

Summarizing broadly for the carbon tax case, a future production subsidy for clean energy changes the fuel mix from which producers generate electricity to include more of the subsidized clean energy, less unsubsidized clean energy, and less dirty fuel or a shift toward different dirty fuels. This has the potential to create shifts in the innovative effort that producers expend toward each fuel. For the subsidized clean energy, innovation captured by LBD and R&D will increase, respectively, if the returns to each increase in second-period production. This trend is affected by the complementarity or substitutability LBD and R&D. As the unsubsidized form of clean energy is scaled back, the same conditions determine whether and how much these forms of innovative effort are decreased when a subsidy is put in place. Innovation directed toward emissions reduction for the dirty fuel will unambiguously fall if production from D is scaled back. If production from D stays the same but the fuel mix is changed, innovation depends on the emissions-reducing returns to innovation from the new fuel mix.

In the case of a carbon tax, all of this occurs due to shifts in production decisions that the firm makes in response to the tech-specific subsidy. In the case of a cap-and-trade system, the subject of the next section, shifts in production not only impact innovation decisions directly but also through their impact on the carbon price, which responds endogenously to abatement activity.

2.3.2 Cap-and-trade (endogenous carbon price) case

The previous section examined how tech-specific subsidies for clean energy have the potential to shift induced innovation toward subsidized technologies and away from market-determined clean energy sources and emissions abatement technology for dirty fuel, assuming an exogenous carbon price and correction of other market failures. This section will use the same model to explore how a tech-specific subsidy will affect induced innovation when the carbon price is endogenous, as in a cap-and-trade system.

For this analysis, I use the same objective function for the electricity firm and add a constraint reflecting the fact that the carbon price is now a function of the firm's production vector. The firm operates according to:

$$\max_{\mathbf{q}_1,\mathbf{q}_2,\mathbf{r}} \pi = [(p_1\mathbf{q}_1 - c_1(\mathbf{q}_1,\mathbf{r})] + \delta[(p_2+s)\mathbf{q}_2 - c_2(\mathbf{q}_2,K,p_c)] \tag{2.25}$$

where

$K = K(\mathbf{q}_1,\mathbf{r})$

$p_1, p_2, \mathbf{q}_1, \mathbf{q}_2, \delta > 0$

$$p_c = p_c(Q_2^D) \tag{2.26}$$

Let Q^D represent the aggregate production of D from all firms. When the carbon price is endogenous, if producers make any change in Q^D away from the baseline production that they would choose without the subsidy s, the carbon price p_c will change. In general, substitution of a clean form of energy for the dirty fuel will depress the carbon price, and visa versa. Formally, $p_c'\left(Q_2^D\right) > 0$. Additionally, if Q^D does not change but producers switch dirty fuel sources such that their marginal cost of abatement changes, the carbon price will respond. The carbon price falls when producers switch to cheaper fossil fuels, even if the emissions intensity of electricity production from D rises. This is because the cost of switching yet again

to abate carbon falls, and visa versa. Substitution of one clean form of energy for another, on aggregate, will not affect the carbon price.[11]

As discussed in the previous section, the implementation of s^{SC} necessarily leads to an increase in output produced from SC, a decrease in output from UC, and a decrease in output from D and/or switching to a cheaper D input in period two, compared to the carbon-price only counterfactual. So, we know that the subsidy will cause the carbon price to fall. How does this impact induced innovation?

As p_c falls, there is less of an incentive for firms to innovate to comply with the carbon price. Looking at Equation (2.16) (copied below), the FOC of the firm maximization problem that contains p_c, we can see that the firm's research effort directed toward D falls as the carbon price falls. This is true if, after the subsidy, the firm produces less or the same amount of electricity from D. Fischer and Preonas (2010) show that q^D falls when firms face a falling carbon price in a cap-and-trade system.[12]

$$c_1'(r^D) = -\delta \left(p_c \frac{\partial e^D(K^D)}{\partial r^D} q_2^D \right) \tag{2.16 revisited}$$

Since the other fuels do not face the carbon price, induced innovative effort for them follows the results of the carbon tax case above, adjusted for any additional production shifts in renewables that occur because q^D falls.

The takeaway from this section is that the electricity market's production and induced innovation responses to climate policy and tech-specific clean energy subsidies differ between a setting in which a carbon tax is used as first-best climate policy and a setting in which a cap-and-trade system is used. In a carbon tax setting, the first-best climate policy is not impacted by the tech-specific clean energy subsidy. The subsidy directly impacts firms'

[11] Note that if SC is substituted for UC (market-wide), no additional abatement is achieved from the subsidy effort, only an increase in the social cost of abatement.

[12] Fischer and Preonas (2010) explicitly include production quantities and emissions rates of two types of dirty fuel; their model includes the effect of fuel-switching within D production. This result follows from the fact that, when the carbon price falls, the dirtiest fuel is used more, and, to compensate for emissions, the less dirty fuel is scaled back quite a bit.

production and innovation choices, as described in the previous section. In a cap-and-trade setting, tech-specific clean subsidies impact the first-best climate policy—and the innovation it induces—through the carbon price. Because the carbon price is depressed from the renewables subsidy, innovation for emissions abatement from dirty fuel unambiguously falls.

2.4 Simulation

The results I obtain from my analytical model above point to the role of marginal returns for innovation in production being a driving force in long-run innovation trade-offs. In the case of a carbon tax, whether there are actually long-run innovation trade-offs from policymakers using clean energy subsidies in combination with the tax depends on these returns, by technology. In the case of a cap-and-trade system, we anticipate innovation trade-offs to occur. The subsidies lead to an increased reliance on clean energy, which accomplishes emissions abatement that would otherwise need to be undertaken by producers from dirty fuels. This reduces the marginal cost of abatement (among producers; the marginal social cost would include the cost of the subsidies), the carbon price, which is provides an incentive for inducing long-run innovation.

I use a simulation of the electricity sector to investigate long-run innovation trade-offs in these cases for at least three reasons: 1) For the tax case, a simulation may resolve the question as to whether long-run trade-offs occur. 2) For the cap case, a simulation can provide an estimate of the magnitude of the trade-off between innovation devoted to subsidized and unsubsidized fuels. 3) A simulation provides a framework for investigating production and innovation changes under each policy scenario that occur between dirty fuels of varying carbon intensities. In particular, emissions abatement in the electricity sector often takes the form of producers switching from coal to natural gas. (Coal is more carbon intensive than natural gas.) Including this important detail in analytical models leads to quite detailed models. (See Fischer and Preonas (2010) for an example.)

My simulation includes two novel features: exogenous technical change and innovation devoted to "dirty" technology. Exogenous technical change, defined as technological im-

provement that occurs over time due to unidentified reasons (such as improved processes or materials anywhere along the supply chain for a particular technology), is often confounded with plausible determinants of technological change (such as increases in output, which capture learning-by-doing within a given industry, or RD investment) (Nordhaus, 2014). Given the recent recognition of the importance of including it in learning models, I include exogenous technical change in line with Nordhaus (2014)'s estimates and adjust other innovation parameters accordingly. By including innovation devoted to "dirty" technology, I capture innovations to reduce the emissions rate of carbon intensive fuels that may be crowded out by clean energy supports. These technologies could be modifications of existing dirty technologies (such as carbon capture and storage-like technologies) or breakthrough technologies for new clean energy that would be just as cost effective as existing dirty technologies and offer zero or lowered emissions.

I calibrate my simulation to EU electricity sector data, but it can be easily adapted to other settings. The EU as a whole–and its member states–have long used both carbon pricing and clean energy subsidies to meet emissions reductions goals, so this setting provides a colorful example of this issue. Note that I consider the time period 2011-2015 as a baseline for my simulation; both carbon pricing and clean energy subsidies already tailored technology costs and the suite of technology available for electricity production by this time. My results point to what future trade-offs the region may experience under various policy scenarios. Through The European Green Deal, EU leaders have proposed a 2050 "climate neutrality" objective, so this issue is still relevant to the region.[13] My intention with this simulation is not to provide a policy analysis of options on the table for the EU; instead, I provide a stylized picture of what innovation trade-offs may look like from these common overlapping policies in a developed setting.

In the remainder of this setting, I first describe my model, the data I use for calibration, and other calibration choices. I then present the scenarios I use for investigating whether, and to what degree, long-run innovation trad-offs occur when we "pick winners". The next

[13]https://ec.europa.eu/info/files/communication-european-green-deal_en

section includes my results. Finally, I offer estimates of welfare changes that result from these long-run innovation trade-offs.

2.4.1 Model

I use a two-period model of the electricity sector to examine innovation outcomes from policy changes. Policy changes are made in the future (the second period) and impact innovation decisions in the fist period. This captures investment decisions that are made to adjust for policy.[14,15]

I assume that production ($\mathbf{q}_t$) and innovation ($\mathbf{r}$) choices in this sector are made to meet electricity demand at lowest cost.[16] My objective function takes the form of:

$$\min_{\mathbf{q}_1,\mathbf{q}_2,\mathbf{r}} C = n_1[f(\mathbf{q}_1) + g(\mathbf{r})] + n_2\delta[K(\mathbf{r})f(\mathbf{q}_2) - \mathbf{s}\mathbf{q}_2 + K(\mathbf{r})e(\mathbf{q}_2)] \tag{2.27}$$

subject to:

$$\sum_{i \in F} q_{i,t} \geq D_t \tag{2.28}$$

in which n_t denotes the number of years in period t, $f(\mathbf{q}_t)$ is the cost of producing electricity from the vector of technologies available to the sector, and g($\mathbf{r}$) is the cost of investing in research for all of those technologies. δ is a discount factor based on an annual discount rate that discounts period two years back to the beginning of period one. The term $K(\mathbf{r})$ is an innovation multiplier that translates research investment from period one into reduced production costs for clean fuels and reduced emissions costs (the cost of complying with a

[14]Therefore, I do not include uncertainty over policy, macroeconomic conditions, or other drivers of innovation investment.

[15]It has been suggested that I consider a multi-period model to capture compounding innovation dynamics. I leave this task to future work, given the computational complexity involved. Additionally, a two-period model suits my decision to emphasize the role of research spending over learning-by-doing dynamics in this simulation.

[16]One could also use profit-maximization as an objective, particularly if the electricity sector being modelled is not regulated. Cost-minimization captures well the aggregate objective function of system operators that determine production based on plants' marginal cost to meet demand constraints. In terms of the innovation decision, cost-minimization captures best the decision-making of deregulated plant owners to adopt new technology. As Fischer and Newell (2008), Fischer et al. (2017), and Hubler et al. (2015), I argue that it is the best forward-looking modelling choice, as the sector is tending toward deregulation.

carbon price) for dirty fuels in period two.[17] Subsidies are contained in the vector $\mathbf{s}$; they linearly reduce production costs for clean fuels. Carbon taxes enter Equation (2.27) through the term $e(\mathbf{q_2})$, which is the cost of emissions.

Equation (2.28) is the demand constraint. The sum of production from all technologies i in set F must be greater than or equal to electricity demand for each period, D_t.[18] Technologies available for production include the "dirty" technologies coal and natural gas; baseload (non-changing), zero-emissions technologies nuclear and hydro; and "clean" technologies solar and wind.[19]

In the absence of policy, the production choice is based on the marginal cost of production from each technology. The cost function is:

$$f(q_i) = (Fuel + O\&M)q_i + \alpha_i q_i^2 \tag{2.29}$$

The total cost of producing electricity from technology i is comprised of a linear term, $(Fuel + O\&M)q_i$, a production of q units for each technology is linear the fuel and operations and maintenance (O&M) for much of the range of production. The quadratic term $\alpha_i q_i^2$ captures additional costs (and cost-savings) that arise from scale within and between plants or installations.[20]

[17]Research spending could, in this setting, take the form of energy firms spending to improve or create new technology. It could also take the form of spending on existing technology. The extent to which it represents development or adoption of new technology depends on the relative magnitude of the carbon price. Additionally, research spending for reducing the emissions rate of dirty technologies could be taken as development of less emissions-intensive or even clean technology. I primarily use the term "research spending" to refer to the choice of

[18]I set demand as a hard constraint in this model to avoid the optimization challenges of using a highly inelastic demand function. Fischer and Newell (2008), Fischer et al. (2017), and Hubler et al. (2015) provide examples of including a demand function in a similar model.

[19]I exclude production from oil, as this fuel is used for less than 2% of the EU's electricity and has been declining over time.

[20]For example, if coal-fired power plants are operating at output rates lower than their optimal, they experience efficiency improvements for increasing output. Operating beyond this optimum will incur efficiency penalties and increases in O&M costs. The quadratic term also captures heterogeneity in producers' costs. For example, the cost of building or maintaining a solar installation in one region may differ from another; increasing production from solar technology in a setting in which only high-cost regions are available results in a production cost that increases beyond the average cost for the technology. This term also captures grid constraints and other system-wide factors that lead to nonlinearities in cost by technology. See Section 2.4.2

The innovation, or research spending choice, depends on the relative cost of the research and the benefit gained—in terms of a reduction in the production cost for clean technologies or emissions cost for dirty policies.[21] Research cost is:

$$g(\mathbf{r}) = r_i^{\gamma} \tag{2.30}$$

Research outcomes from R&D spending decline in the marginal dollar spent; the benefit of adopted technology is assumed to be declining in marginal spending on technology ($\gamma > 1$).

Research benefits are determined by the following function.

$$K_i(r_i) = (\frac{S + r_i}{S})^{\sigma} e^{(n_1 - 1)\rho} \tag{2.31}$$

In the objective function, $K_i(r_i)$ serves as a multiplier that reduces the cost of producing electricity from clean technologies or the emissions rate of dirty technologies ($K_i(r_i) \in (0, 1)$). The stock of knowledge, S, is increased with a given investment in either solar or wind technology, $r_{s,w}$. An important parameter is σ, which is either the elasticity of clean energy production cost with respect to research spending (for clean technology) or the elasticity of the emissions rate with respect to research spending (for dirty technology).[22] This functional form for the impact of research spending on production costs follows the learning curve literature (Rubin et al. (2015); Nordhaus (2014)) and the models used in other studies in the electricity sector that incorporate innovation Fischer and Newell (2008), Fischer et al. (2017), and Hubler et al. (2015).

---

for detail on cost data for each technology.

[21] Unlike my analytical model, my simulation does not include learning-by-doing. More specifically, period one production does not impact the cost of period two production. Explicit research spending only impacts production and/or emissions costs in period two. I make this choice for several reasons. First, it keeps the model simple. Second, additional production in period one in an effort to reduce period two costs could be captured by an independent cost term, which could easily be incorporated into a generic "research spending" term. Third, the technological change literature specific to electricity production points to R&D spending playing a larger role in reducing costs than learning-by-doing (Rubin et al., 2015).

[22] Note that when research spending is devoted to reducing the emissions rate of dirty technology, the cost of production from dirty technology is also reduced through a reduction in the emissions cost. See Section 2.4.2 for my calibration decisions.

The second term in $K_i(r_i)$, $e^{(n_1-1)\rho}$, includes exogenous technical change in the model. The exponent ρ dictates the rate of annual reductions in production or emissions cost. I assume this rate is 1% per year, based on Nordhaus (2014).[23]

For dirty technologies, the emissions cost in carbon tax scenarios is:

$$e(q_i) = p^c K_i(ER)_i q_i \tag{2.32}$$

in which p^c is the price of carbon (set exogenously) and $(ER)_i$ is the emissions rate of technology i.

For cap-and-trade scenarios, the objective function is constrained by an exogenously set Cap limiting total emissions:

$$Cap \geq \sum_{i=c,ng} K_i * (ER)_i * q_i \tag{2.33}$$

2.4.2 Calibration

Time periods and discount factor

I set the length of period one to be five years and period two to be thirty years in length. These are chosen to represent the near-term, in which investment decisions are made based on upcoming policy changes, and the medium- to long-term, in which research spending has an impact on technology.

Following Fischer and Newell (2008), Fischer et al. (2017), and Hubler et al. (2015), I use a discount factor δ that discounts period two values back to the beginning of period one and then to the beginning of period one:

$$\delta = (1+r)^{-n_1}((1-(1+r)^{-n_2})(1+r)/r)/n_2 \tag{2.34}$$

I use an annual discount rate of 7% to match the median discount rate used by IEA et al.

[23] Nordhaus (2014) estimated that the rate of exogenous technical change among US non-farm business 1948-2011 to be about 1%.

(2015) in its levelized cost of energy (LCOE) estimates, on which I rely for my cost parameters.

Cost parameters

For my technology-specific cost functions (Equation (3)), I use estimates of the fuel cost and O&M cost from IEA et al. (2015). These costs are components of levelized cost of energy (LCOE) estimates. The International Energy Agency (IEA) and Nuclear Energy Agency (NEA) calculate LCOEs from cost data reported by OECD and a few additional countries for anticipated 2020 costs. LCOE estimates are intended to calculate the long-run average, per-unit cost of generating electricity by technology. I calibrate the parameter a_i so that period one production values equal observed 2015 values, by technology. Cost parameters are shown in Table 2.1.

Table 2.1: *Cost parameters*

Technology	Fuel cost (2013 USD/MWh)	O&M cost (2013 USD/MWh)	Quadratic parameter, α_i
Coal	27045.81	12864.91	0.047
Natural gas	73659.06	8700.70	0.030
Nuclear	9502.71	15984.90	0.000
Hydro	0.00	11392.53	0.000
Solar	0.00	29339.23	0.405
Wind	0.00	29865.99	0.135

Note: Table displays cost parameters for Equation (2.29).

Research cost

Following Fischer et al. (2017), I set γ in Equation (2.30) equal to 1.2, which captures empirical estimates of research output (such as patents) with respect to R&D.

Production calibration

I use electricity production data, by technology, 2011-2015 from the European Network of Transmission System Operators for Electricity (ENTSO-E) for all production-relation

calibration choices.[24] I calibrate period one production values to 2015 data.[25]

Emissions rates

I use average emissions rates of coal and natural gas provided by the Energy Information Administration (EIA).[26]. I also use EIA average heat rates for coal and natural gas plants.[27] Using production data from the European Network of Transmission System Operators for Electricity (ENTSO-E), I weight the coal emissions rate by the shares of EU coal types used for electricity generation (anthracite and lignite).

Knowledge stock

I set the knowledge stock, S in Equation (2.31) so that baseline period one research choices for clean technologies are no more than 2.5% of the total stock, implying a minimum of 40 years of pre-existing research investment. My knowledge stock takes the value of 30,000,000.

Knowledge elasticities

σ

Recall that σ, in Equation (2.31), is either the elasticity of clean energy production cost with respect to research spending (for clean technology) or the elasticity of the emissions rate with respect to research spending (for dirty technology). I set independent values for clean and dirty technologies.

There are a wealth of studies that estimate the elasticities of wind and solar production cost with respect to cumulative capacity (intended to capture learning-by-doing) and R&D

[24] Available at: https://www.entsoe.eu/data/data-portal/#monthly-values

[25] Because I only include electricity from coal, nuclear, hydro, natural gas, solar, and wind in my model, there is production in the region from other sources (namely oil and other renewables such as geothermal and biomass) that is not captured in this model. This model captures 82-87% of regional production captured 2011–2015.

[26] https://www.eia.gov/environment/emissions/co2_vol_mass.php

[27] https://www.eia.gov/electricity/annual/html/epa_08_01.html

spending or other measures of research effort, such as patenting. Rubin et al. (2015) provide a very useful summary. However, Nordhaus (2014) points out that this "learning curve" literature fails to account for "exogenous technical change", change that occurs over time due to unidentified or unobservable forces. He shows that learning rates estimated *with* a time term to capture exogenous technical change are generally lower than those estimated without such a time term. However, they are also uncorrelated with the no-time-term estimates, which indicates that there is little one can do to accurately deflate the upward-biased estimates in the literature to estimates that account for exogenous technical change.

Given the lack of guidance on how sensitive production and emissions costs should be to innovation, I choose values that, conditional on my other parameters, lead to changes in research spending of no more than 50% of the research stock, for most policy scenarios. Despite this being a very generous rule, it resulted in a small elasticity for clean technology: I set σ equal to -0.01, meaning that a 1% change in research spending on wind or solar technology results in a 0.01% reduction in the production cost of those technologies. In contrast, Fischer et al. (2017) use values of -0.15 and -0.3 for wind and solar technologies, respectively.[28]

I set σ for dirty technology equal to -0.02. This means that a 1% increase in research spending results in a 0.02% drop in the emissions rate. This translates into a 0.01% reduction in total per-unit production cost for coal at a carbon tax of $50. Therefore, research spending on dirty technology has an impact on the production cost for coal similar to the impact that research spending has on clean technology production cost, when there is carbon tax close to the social cost of carbon (SCC). This choice regards most emissions-intensive technology as "mature", until a high carbon price incentivizes discovery of new technology for less emissions-intensive production.

[28]These are their elasticities on R&D investment. They have separate LBD elasticities for how the capacity from each technology impacts production cost, of -0.1 and -0.2 for wind and solar, respectively. I do not include LBD in my simulation, for simplicity and on the assumption that, if any LBD effects remain after capturing exogenous technical change, they are small or captured by investment through my "research spending" term. One might consider my σ, the elasticity of production cost with respect to research spending, a weighted average of the R&D and LBD elasticities estimated or used in other literature, adjusted for exogenous technical change. Fischer et al. (2017)'s values are on the low end of the estimated elasticities reported in Rubin et al. (2015)'s survey.

ρ

Recall that, in Equation (2.31, the exponent ρ dictates the rate of annual reductions in production or emissions cost. I assume this rate is 1% per year, based on Nordhaus (2014), and therefore set rho equal to -0.01. Nordhaus (2014) estimated the rate of exogenous technical change among US non-farm businesses, 1948-2011, to be about 1%.

Emissions cap

I calculate pre-policy emissions from the average observed EU production of electricity from coal and natural gas, 2011-2015. I set emissions caps as fractions of this pre-policy emissions amount.

2.4.3 Solution method

I program the equations above in R and minimize Equation (2.27), conditional on the constraints discussed, to obtain production and research spending parameter values, for my "baseline" (no policy) and policy scenarios (listed in Section 2.5. For minimizing Equation (2.32), I use the function "auglag" from the "alabama" package, which is an augmented Lagrangian optimization function that allows for linear and nonlinear equality and inequality constraints.[29]

2.5 Scenarios

To examine whether subsidies for specific technologies crowd out innovation that would otherwise be devoted to carbon abatement technology under carbon pricing, I compare innovation outcomes between the following sets of scenarios:

1. Tax-only and tax-plus-subsidies scenarios
2. Cap-only and cap-plus-subsidies scenarios

[29]Documentation can be found here: https://cran.r-project.org/web/packages/alabama/index.html

For each set, I use either a range of carbon taxes or emissions caps. I also use a range of subsidy magnitudes, as a share of the cost of production from the clean technologies. The tax-only and cap-only scenarios are equivalent from the producer's perspective, as the producer faces either an explicit or implicit carbon price under each. (That is, I do not model uncertainty over or volatility of the realized carbon price that can occur under an emissions cap, which can impact investment behavior. See Borenstein et al. (2019) for a recent treatment of the uncertainty issue.) Because I do not model permit trading, however, the cap scenarios simulate an emissions cap with free allocation of permits. Therefore, the implicit carbon price, the marginal abatement cost, is lower under the cap than the tax scenarios for the system to reduce emissions by a given amount.

I compare various scenarios to a "baseline" scenario which includes not policy. See Table 2.2 for results of this scenario.

Table 2.2: *Baseline Results*

Technology	Period 1 production (GWh)	Period 2 production (GWh)	Period 1 research spending (2013 USD)
Coal	$7.02e+05$	$6.95e+05$	$1.10e+00$
Natural gas	$3.92e+05$	$3.81e+05$	$4.09e+01$
Nuclear	$8.14e+05$	$8.14e+05$	
Hydro	$4.67e+05$	$4.67e+05$	
Solar	$9.45e+04$	$9.90e+04$	$7.04e+02$
Wind	$2.81e+05$	$2.95e+05$	$5.34e+05$

Note: Table displays baseline results from choosing production in both periods and research spending to minimize Equation (2.27).

2.6 Results

2.6.1 Production results

I present the production results from all scenarios here to show that the simulation generates intuitive results from policy changes. Additionally, it is useful to review innovation results after obtaining a sense of how policy impacts production choices, as innovation and production

choices are jointly determined as Equation (2.27) is minimized.

Figure 2.1 shows how the production of electricity from each technology is allocated in a cost-minimizing market when a carbon tax is put in place, as a percent change from the baseline scenario (in which there is no carbon tax). (Only second-period production changes appreciably from policies, as policies are enacted in the second period.) A $5 carbon tax causes production from coal to drop 3.7% and production from renewables and natural gas to increase by 3.6% and 2.6%, respectively. These production changes increase with higher carbon taxes. Coal, the most carbon intensive fuel source, becomes a less cost-effective fuel source with a higher carbon tax. Renewable sources and natural gas are more cost-effective at higher carbon taxes. The gap in the percent change in the usage of each widens as the carbon tax increases, however, as clean energy sources are favored over natural gas.

Figure 2.1: *Tax-only scenario results: Period two production by fuel*

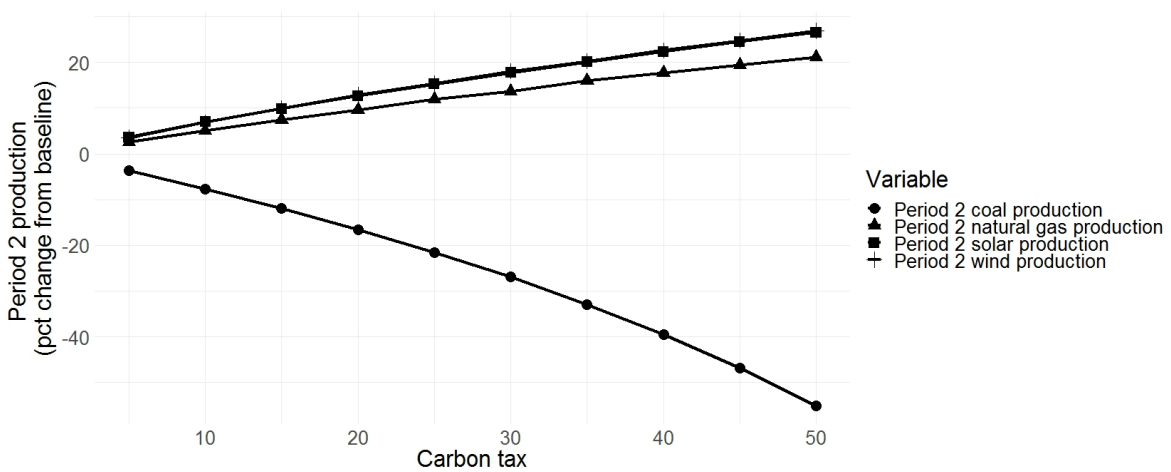

Note: This plot shows the change in period two production that is chosen when a carbon tax is imposed during that period, as a percent from the baseline, no-tax scenario. Results are shown by carbon tax.

When small subsidies are given to production from clean sources (10% of the cost of production from solar and wind, respectively), a $5 carbon tax causes the use of coal to drop 4.7%, natural gas to rise only 0.6%, and the use of renewables to rise 6.5% (Figure 2.2). The trajectory of the percent changes in production from baseline for each technology over the range of carbon taxes is quite similar to the tax-only scenario.

Figure 2.2: *Tax-plus-10% subsidies scenario results: Period two production by fuel*

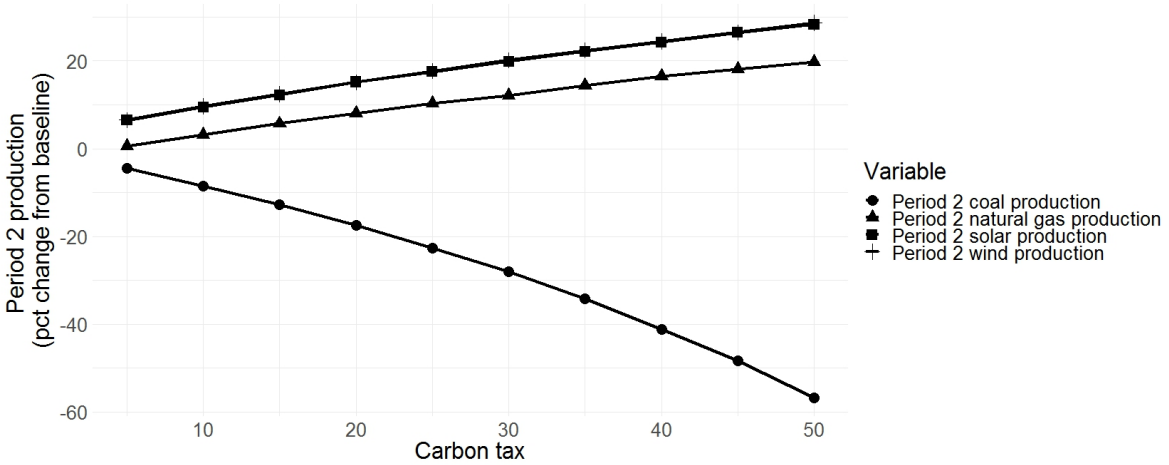

Note: This plot shows the change in period two production that is chosen when a carbon tax is imposed and subsidies equal to 10% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-tax scenario. Results are shown by carbon tax.

Larger subsidies reduce production from both coal and natural gas for small carbon taxes. Figures 2.3 and 2.4 show production results from scenarios in which there is a carbon tax as well as subsidies that are 50% and 100% of the cost of production from clean sources, respectively, in period two. For each set of subsidies, the emissions penalty from higher carbon taxes leads to production from coal being substituted for production from natural gas and clean sources, as in the tax-only and tax-plus-10% subsidy scenarios. The larger sets of subsidies lead to more dramatic substitution between high- and low-carbon fuel sources, however; note that the use of coal drops by over 60% and about 75% in these scenarios for a $50 carbon price, compared to less than 60% for both the tax-only and tax-plus-10% scenarios. For these scenarios with large subsidies, at a low carbon price, production from natural gas falls more as a percent of baseline production than production from coal. This highlights that marginal production decisions depend on the cost of production from each technology, inclusive of all subsidies and emissions costs. With a large subsidy in place and low carbon price, clean technologies are used more. Production must fall from one or both dirty technologies; since there is very little penalty for carbon emissions, relatively cheap coal is favored over natural

gas.

Figure 2.3: *Tax-plus-50% subsidies scenario results: Period two production by fuel*

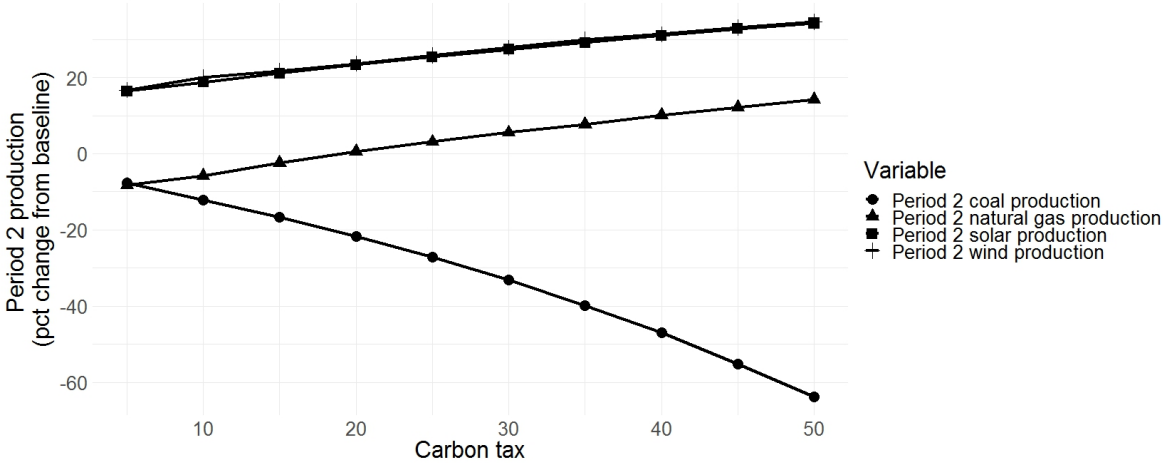

Note: This plot shows the change in period two production that is chosen when a carbon tax is imposed and subsidies equal to 50% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-tax scenario. Results are shown by carbon tax.

Figure 2.4: *Tax-plus-100% subsidies scenario results: Period two production by fuel*

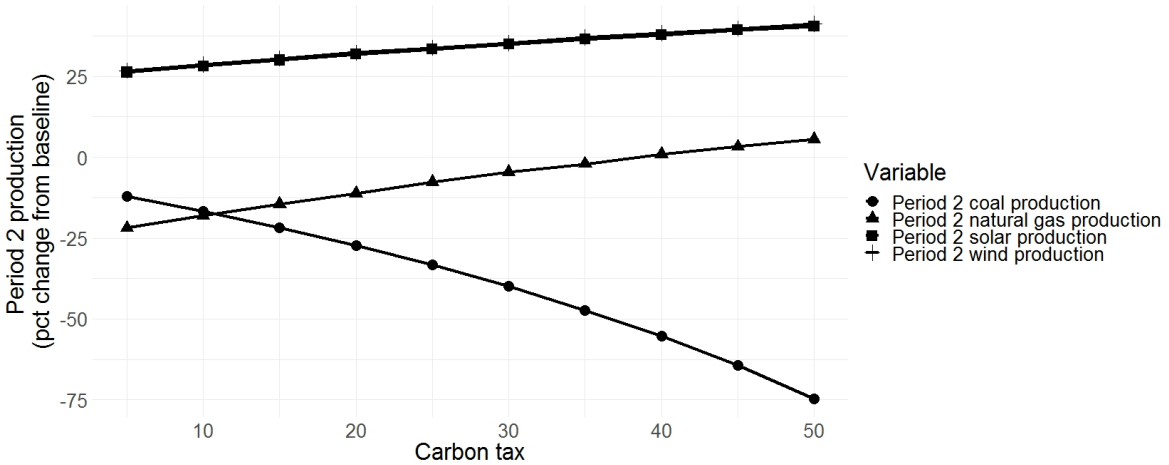

Note: This plot shows the change in period two production that is chosen when a carbon tax is imposed and subsidies equal to 100% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-tax scenario. Results are shown by carbon tax.

Overall, production results from all tax scenarios demonstrate that the simulation generates

expected production results for policy changes. Note that production choices are made jointly with innovation decisions, which work to reduce the cost of production for clean technologies and the cost of emissions for dirty technologies.

The same results for the cap scenarios are show in Figures 2.5-2.8.[30] Recall that these scenarios require the total emissions from both coal and natural gas to be at or below total emissions generated from each tax value on the x-axis. These cap scenarios are analogous to real-world emissions caps that are accompanied by 100% free allocation of permits in accordance with each cap. Therefore, production from carbon intensive technologies is not as expensive *for producers* as it is in the tax scenarios that result in equivalent emissions reductions. Differences in production trends between the tax-only and cap-only scenarios are due to this difference in producer cost as well as difference in innovation choice. Under the cap-only scenario, there is a lot of research spending devoted to reducing the emissions rate of coal for a cap equivalent to a $5 carbon tax, for example (Figure 2.9), whereas under the the tax-only scenario, there is very little research spending devoted to reducing emissions rates (Figure 2.10).[31] The trajectory of production outcomes as the cap decreases (its emissions-equivalent tax increases) is similar to the trajectory of outcomes for the tax-only scenario (Figure 2.1). A small subsidy extends the range of caps for which fuel-switching is not needed to meet the cap constraint (Figure 2.6). Like the high-subsidy and low-tax combination, adding a small subsidy to a small cap incentivizes production from clean technologies and coal. The range of caps over which this occurs is larger for greater subsidy values (Figures 2.7 and 2.8). Large subsidies also require larger changes in the cap stringency for a given magnitude of fuel-switching to occur, as the production from renewables accomplishes a large amount of emissions abatement. When clean technology subsidies are relatively large (80% of fuel cost), even relatively strict caps are easily

[30]For cap scenarios, the simulation meets computational limitations for stringent caps and high subsidies, due to innovation values being pushed to very small values. Therefore, there are some missing values for plots.

[31]My research spending results are much more volatile than my production results, when one compares results between carbon tax or cap stringency, for example. I have intentionally calibrated my model to be less sensitive to research spending than previous research, in order to include exogenous technical change. (See Section 2.4.2.) The effect of this choice is that, when research spending is a cost-effective option for responding to policies, it can take a large amount of spending to have an effect on period two costs.)

met with the heavy reliance on clean technologies and coal (Figure 2.8).

Figure 2.5: *Cap-only scenario results: Period two production by fuel*

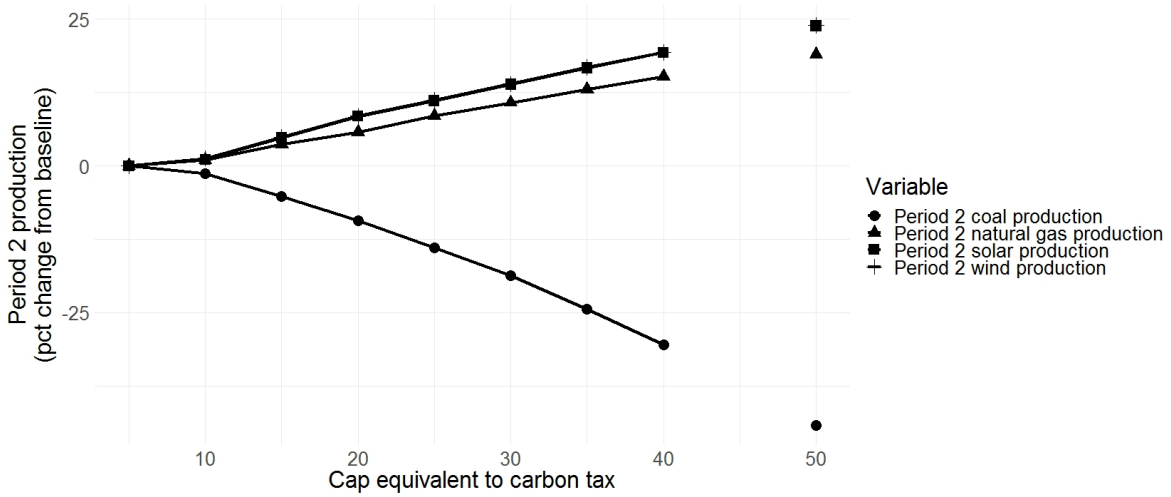

Note: This plot shows the change in period two production that is chosen when an emissions cap is imposed during that period, as a percent from the baseline, no-cap scenario. Results are shown by cap, in terms of each cap's emissions-equivalent tax.

Figure 2.6: *Cap-plus-10% subsidies scenario results: Period two production by fuel*

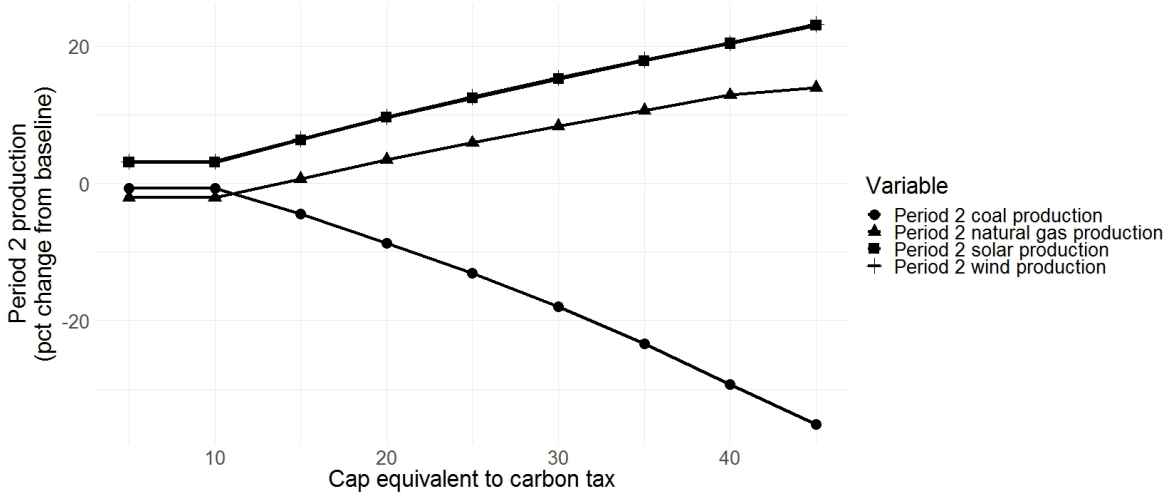

Note: This plot shows the change in period two production that is chosen when an emissions cap is imposed and subsidies equal to 10% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-cap scenario. Results are shown by cap, in terms of each cap's emissions-equivalent tax.

Figure 2.7: *Cap-plus-50% subsidies scenario results: Period two production by fuel*

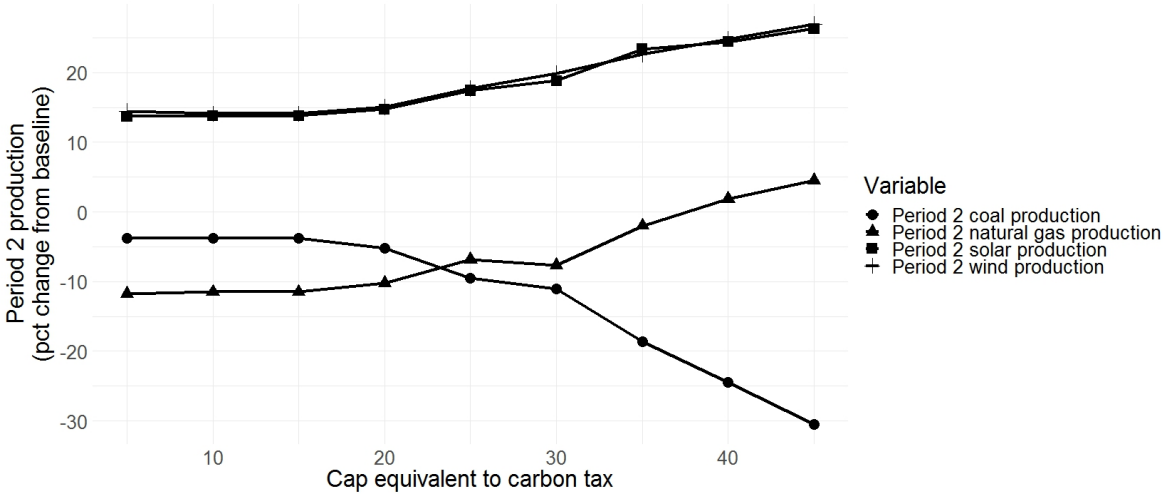

Note: This plot shows the change in period two production that is chosen when an emissions cap is imposed and subsidies equal to 50% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-cap scenario. Results are shown by cap, in terms of each cap's emissions-equivalent tax.

Figure 2.8: *Cap-plus-80% subsidies scenario results: Period two production by fuel*

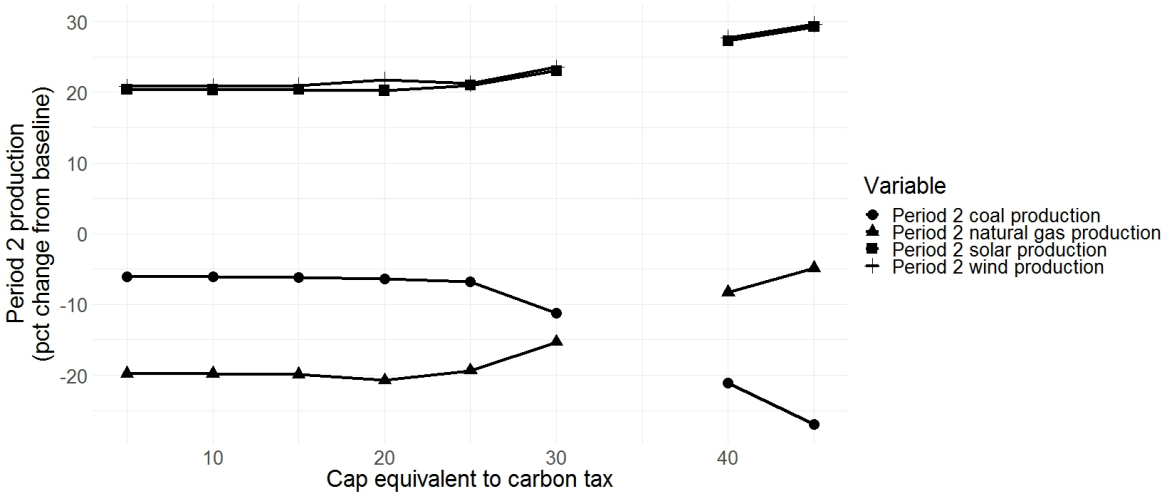

Note: This plot shows the change in period two production that is chosen when an emissions cap is imposed and subsidies equal to 80% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-cap scenario. Results are shown by cap, in terms of each cap's emissions-equivalent tax.

2.6.2 Innovation results

Carbon tax scenarios

Figure 2.11 shows the difference in research spending between the tax-only scenarios and the tax-plus-subsidy scenarios, with subsidies for solar and wind technology that are 10%

Figure 2.9: *Tax-only scenario results: Period one research spending by fuel*

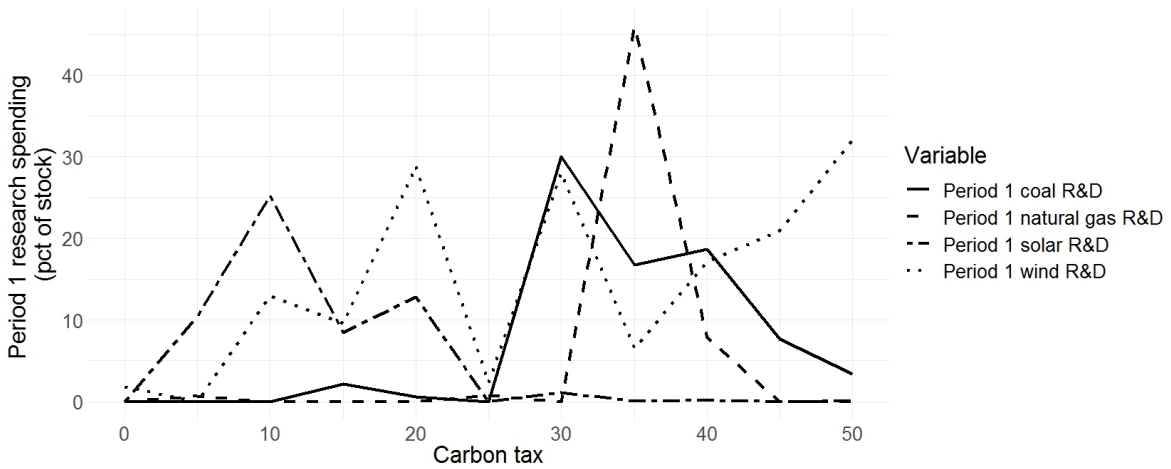

Note: This plot shows the research spending that is chosen when an emissions tax is imposed during during period two, as a percent of the research stock. Results are shown by tax.

Figure 2.10: *Cap-only scenario results: Period one research spending by fuel*

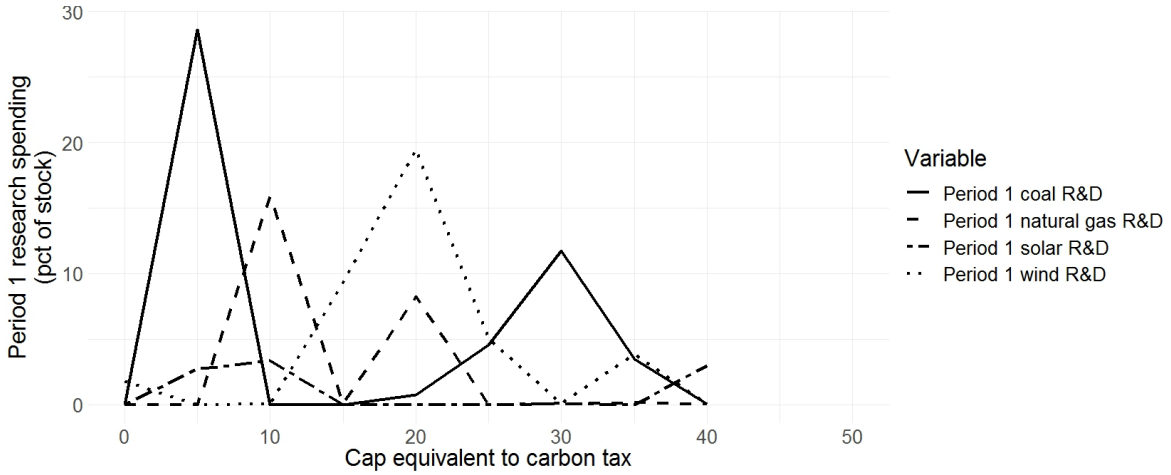

Note: This plot shows the research spending that is chosen when an emissions cap is imposed during during period two, as a percent of the research stock. Results are shown by cap, in terms of each cap's emissions-equivalent tax.

of these technologies' production cost. There is not a clear pattern of changes in research spending from the tax-only scenarios and tax-plus-subsidy scenarios, except perhaps that, at low carbon taxes, there is less research spending devoted toward clean technologies.[32] On the

[32]The subsidy provides a reduction in the marginal cost of production for solar and wind technology, so less research spending is desired for the same purpose. This follows the intuition of Acemoglu (2003) that the

whole, for small subsidies, it does not look like innovation devoted toward dirty fuel is reduced by the presence of a subsidy.

Figure 2.11: *Tax vs. Tax-plus-10% subsidies scenario results: Change in research spending*

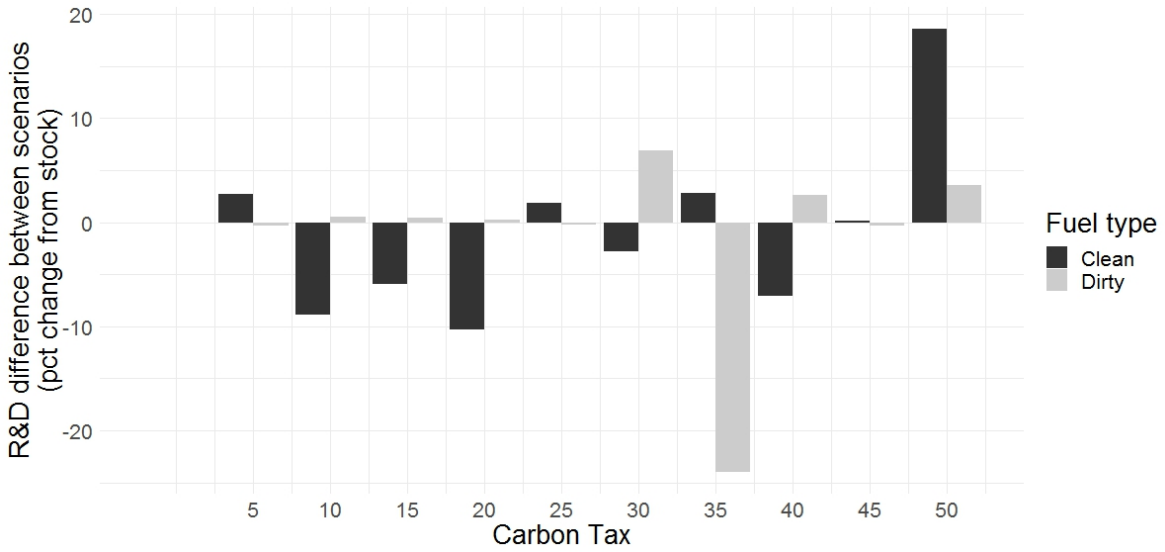

Note: This plot shows the change in research spending between the tax and tax-plus-10% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios.

In contrast, Figure 2.12 shows that the combination of a larger subsidy (50% of the cost of the production from solar and wind) and moderate carbon taxes ($30-$45) leads to a reduction in research spending for dirty technology compared to the tax-only scenarios. For these tax values, research devoted to clean technology is either increased or reduced by a smaller percentage than dirty technology research. The reduction in research for dirty technology is driven by fuel-switching being more cost-effective in these tax-plus-subsidy scenarios than their respective tax-only scenarios; this choice is the result of subsidies being available.

Figure 2.13 shows a similar trend: for carbon taxes in the range of $30-$50, research spending for dirty technology is reduced while research spending for clean technology is either increased or reduced by a smaller percentage. Or, research spending for both types

direction of innovation follows both the marginal cost and the relative abundance of inputs in production. Here, the effect of a marginal cost reduction dominates, as low carbon prices likely do not provide an incentive for more production from clean sources than the subsidy incentivizes.

Figure 2.12: *Tax vs. Tax-plus-50% subsidies scenario results: Change in research spending*

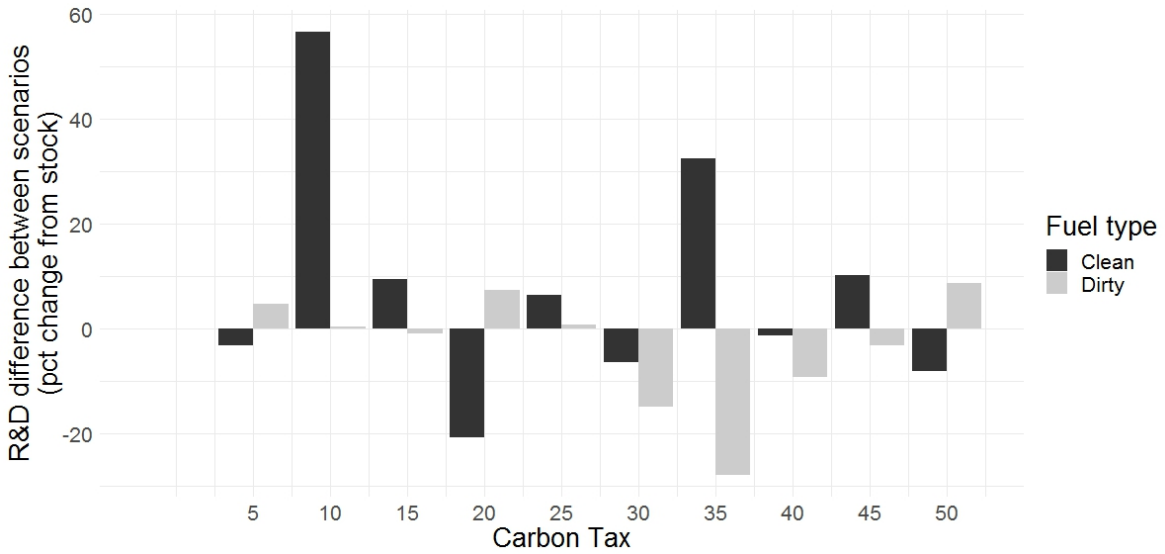

Note: This plot shows the change in research spending between the tax and tax-plus-50% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios.

of technology is increased—but more for clean than for dirty fuels. The take-away from this figure is that large subsidies, when overlapped with high carbon prices, make research spending for reducing emissions from carbon-emitting technologies a less desirable way to comply with the carbon taxes. Instead, it is more cost-effective to use clean technology and spend on research to reduce the cost of production from clean technology as necessary. (Recall that fuel switching increases in subsidy size, for a given carbon tax; compare Figures 2.2 - 2.4.)

Emissions cap scenarios

Figures 2.14 - 2.17 show the change in research investment, by cap stringency, that occurs when subsidies that are 10% - 80% of the cost of production from clean technologies are offered to producers for using clean technologies, in addition to them needing to comply with the carbon cap. Figure 2.14 shows that research for both clean and dirty technologies mostly fall when a small subsidy is combined with an emissions cap. The more stringent the cap, the less this effect occurs. At a cap equivalent to a tax of $10 or less, the subsidy assists producers in meeting the cap. Research spending on emissions rates reductions drops

Figure 2.13: *Tax vs. Tax-plus-100% subsidies scenario results: Change in research spending*

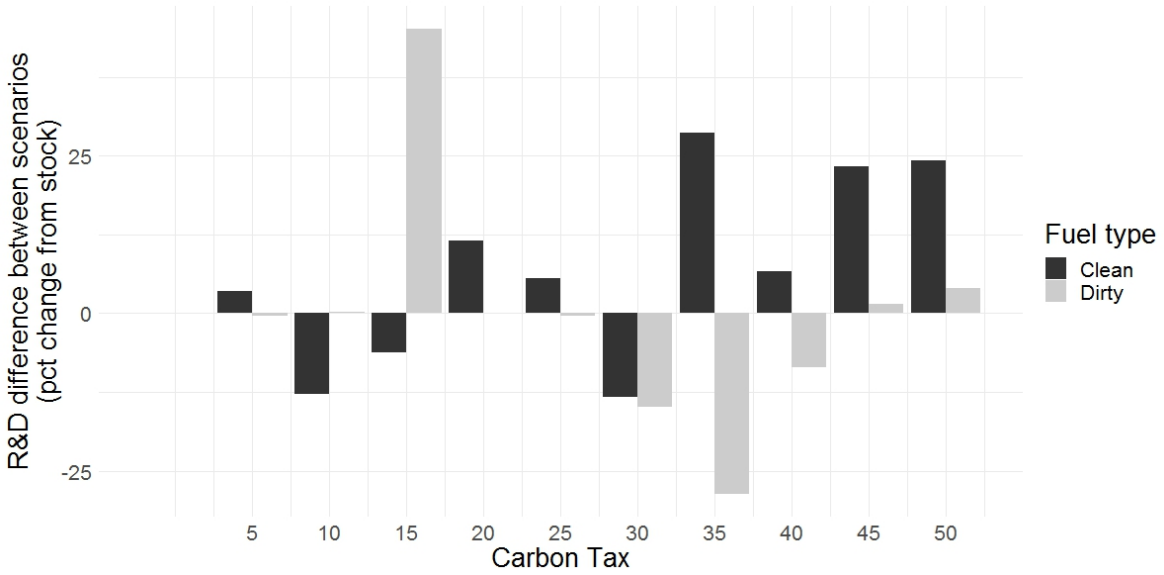

Note: This plot shows the change in research spending between the tax and tax-plus-100% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios.

dramatically. Research spending on clean technologies increases slightly (1-2%). However, when the cap is set at the emission-equivalent of a $15 tax or higher, producers start to orient production toward less emissions-intensive technologies, beyond incentives that the subsidy provides (Figure 2.6). The research spending gap between the cap-only and cap-plus-subsidy case falls as the cap tightens, indicating that research spending toward reducing emissions or using clean technology increases in value with the stringency of the cap (Figure 2.14).

Figure 2.15 shows the same results for subsidies that are 50% the cost of production from clean technologies.[33] Again, for less stringent caps, the cap is easily met with production from clean technology, and, generally, there is a reduced incentive to devote research spending to either technology. For more stringent caps, research spending is more in line with the cap-only scenario.

Figure 2.17 shows the innovation results for subsidies that are 80% the cost of production

[33]This plot excludes results for the caps with the emissions-equivalant tax of $30, which has very large changes in research spending, particularly for clean technology. See Figure 2.16.

Figure 2.14: *Cap vs. Cap-plus-10% subsidies scenario results: Change in research spending*

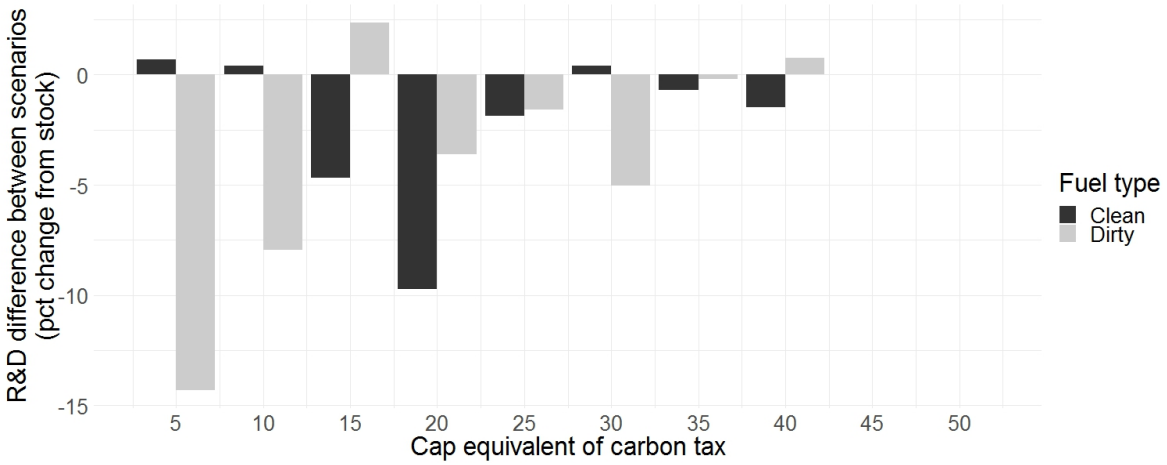

Note: This plot shows the change in research spending between the cap and cap-plus-10% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios.

Figure 2.15: *Cap vs. Cap-plus-50% subsidies scenario results: Change in research spending*

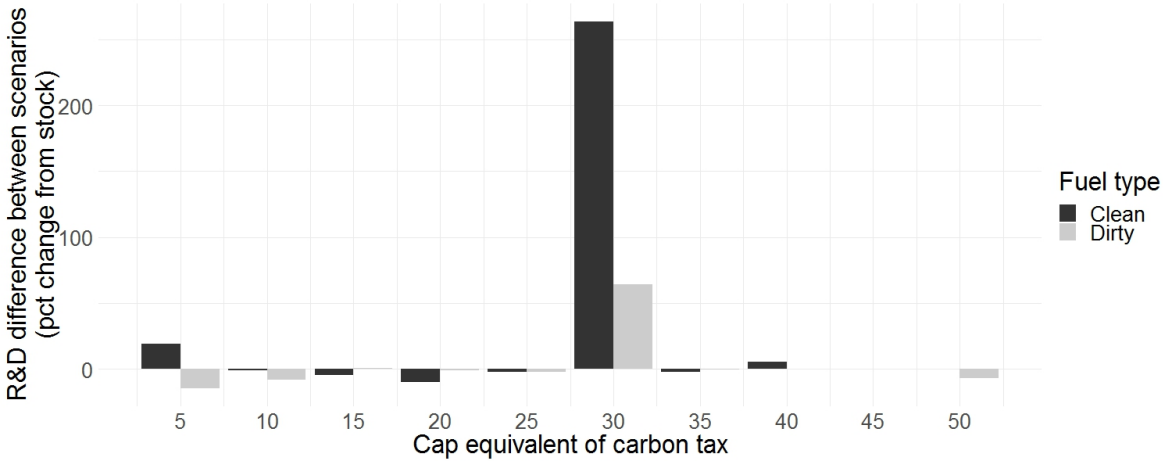

Note: This plot shows the change in research spending between the cap and cap-plus-50% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios.

from clean technologies. There is not a clear trend in changes in research spending when subsidies are added, by cap stringency. In general, though, except for a few caps, there seems to be very little change in research spending between the cap-only and cap-with-subsidy case.

Figure 2.16: *Cap vs. Cap-plus-50% subsidies scenario results: Change in research spending*

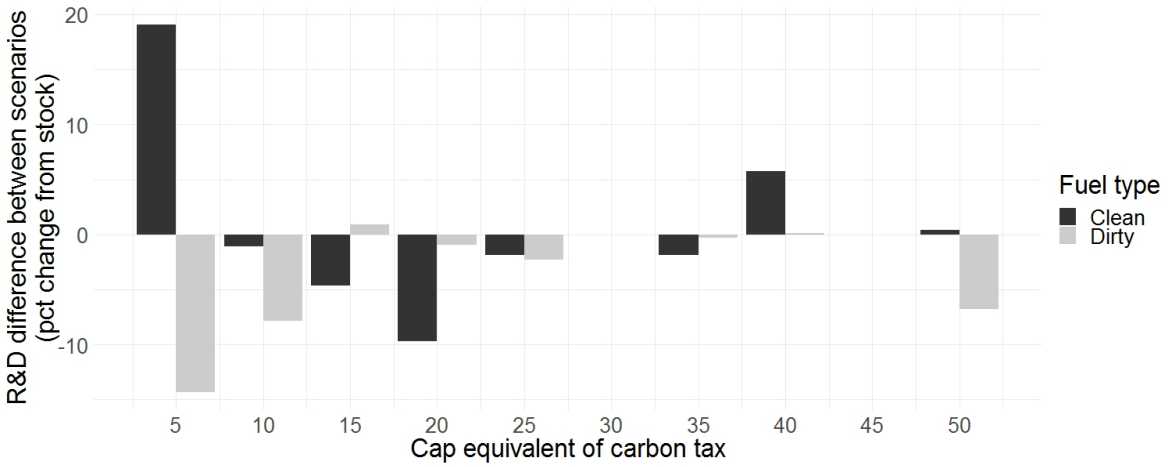

Note: This plot shows the change in research spending between the cap and cap-plus-50% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios. This plot excludes results for the cap that is emissions-equivalent to a $30 carbon tax for graphical clarity.

Figure 2.17: *Cap vs. Cap-plus-80% subsidies scenario results: Change in research spending*

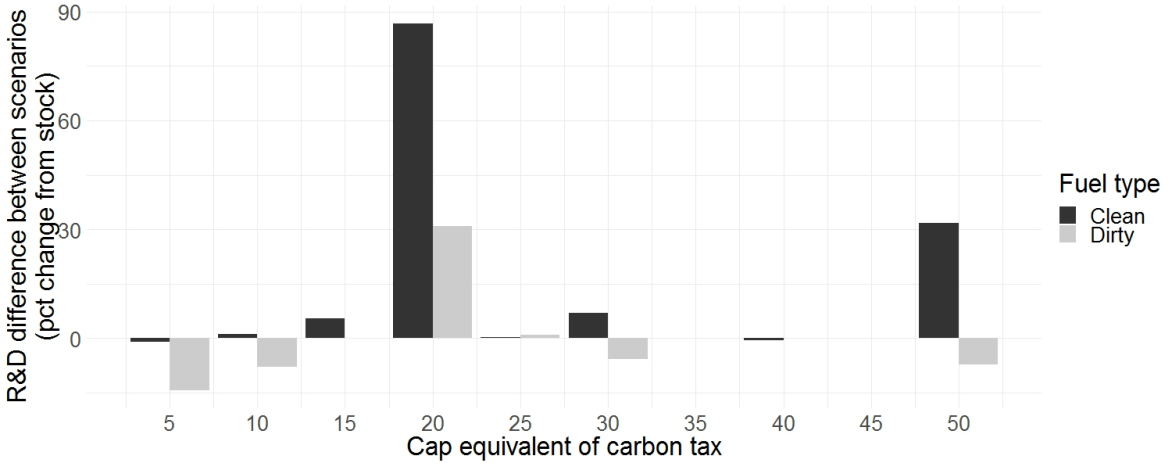

Note: This plot shows the change in research spending between the cap and cap-plus-80% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios.

2.7 Robustness

In this section, I examine whether my results are robust to the discount rate I use. From the perspective of a firm making investments in new technology, assumptions over discount

rates are critical for financing capital and managing risk. From the perspective a social planner, the social discount rate can also encompass social preferences over the timing of costs and benefits. In my modelling above, I assume a 7% annual discount rate to align my model with the assumptions of the LCOE estimates I use for my cost parameters. How would my results change under a smaller discount rate, that, for example, a social planner might use to more heavily weight the welfare of future periods in contemporary policy decisions?

Here I present the results of my tax scenarios under a 3% annual discount rate. (Note that I do not use alternative cost estimates from IEA et al. (2015); I only change r in the model above. Therefore, these results are closer to changing just social preferences than changing social preferences and capital costs to producers.) This has the effect of making all period-two terms more heavily weighted in minimization of Equation (2.27). Notably, the cost of research spending for all technology does not change (since it is made in the first period). However, the benefits are more valuable. Period-two production choices under a 3% discount rate should be the same as the 7% discount rate, however, as they result from incentives in just that period.

Figures 2.18-2.21 show the period two production results for my tax scenarios under a 3% discount rate. As predicted, they are nearly identical to the production results I obtain with a 7% discount rate (Figures 2.1 - 2.4). Overall, the raw research spending results reveal that more research spending is chosen under a 3% than 7% discount rate, over a variety of carbon tax values and subsidy amounts (not shown for brevity). When I compare research under each tax-plus-subsidy scenario to the tax-only scenario (Figures 2.22 - 2.24), I see that, broadly, there is more research spending chosen when subsidies are offered, for carbon taxes under \$40. This suggests that, even though the clean subsidies have made production from clean technology cheaper, the high value of research incentivizes research spending. This stands in contrast to the same results obtained with a 7% discount rate, (Figures 2.11 - 2.13), in which research spending is either reduced in the presence of subsidies (for clean technology) or does not have a clear pattern for carbon taxes under \$40.

Interestingly, carbon taxes of \$40-\$45 lead to reduced research spending under all subsidy amounts, compared to the tax-only scenario. This suggests that the stringency of this emissions

Figure 2.18: *Tax-only scenario results, 3% discount rate: Period two production by fuel*

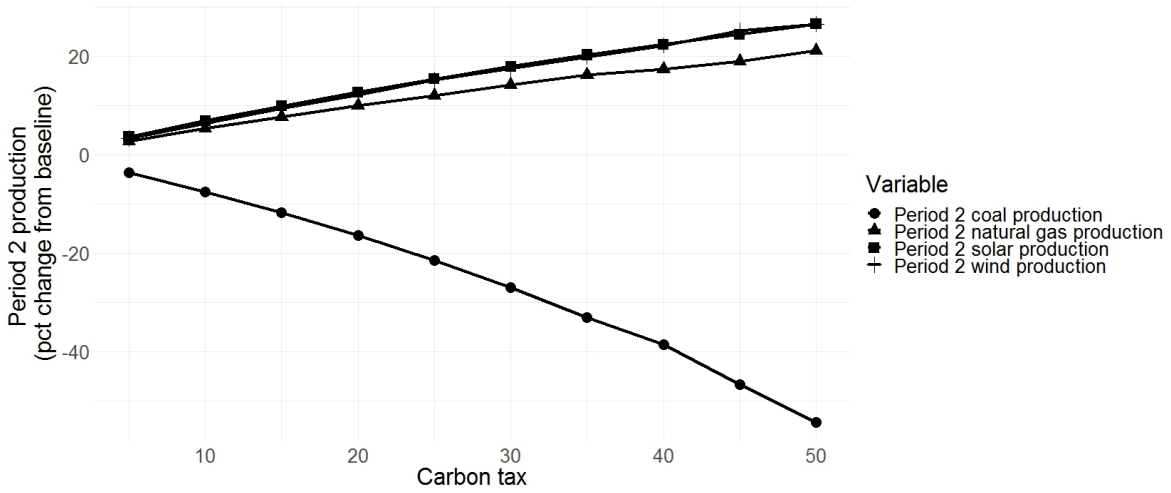

Note: This plot shows the change in period two production that is chosen when a carbon tax is imposed during that period, as a percent from the baseline, no-tax scenario. Results are shown by carbon tax. These results are obtained with an annual discount rate of 3%.

Figure 2.19: *Tax-plus-10% subsidies scenario results, 3% discount rate: Period two production by fuel*

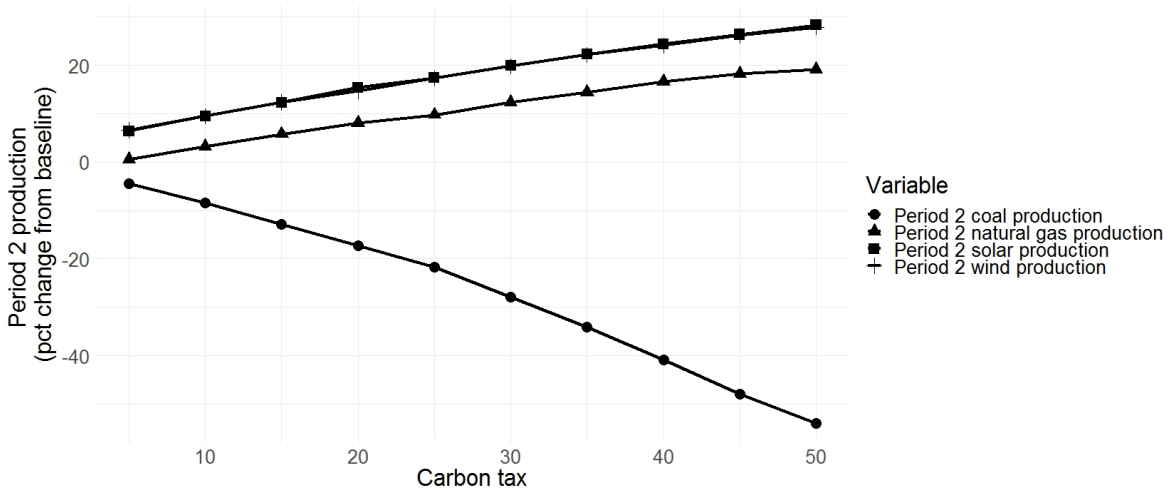

Note: This plot shows the change in period two production that is chosen when a carbon tax is imposed and subsidies equal to 10% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-tax scenario. Results are shown by carbon tax. These results are obtained with an annual discount rate of 3%.

cost makes research spending less cost-effective for reducing that cost than fuel-switching, no matter what subsidies are available.

Figure 2.20: *Tax-plus-50% subsidies scenario results, 3% discount rate: Period two production by fuel*

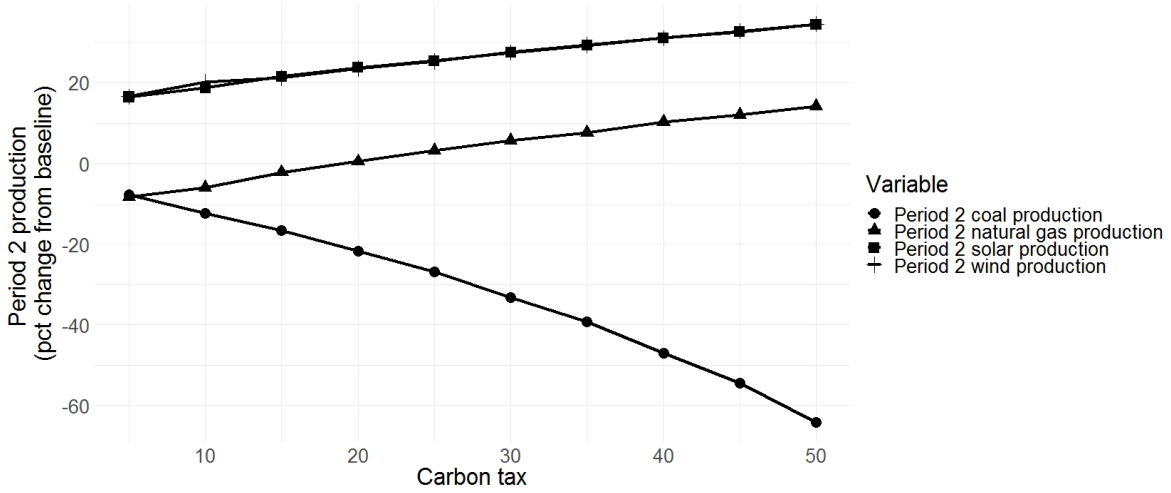

Note: This plot shows the change in period two production that is chosen when a carbon tax is imposed and subsidies equal to 50% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-tax scenario. Results are shown by carbon tax. These results are obtained with an annual discount rate of 3%.

Figure 2.21: *Tax-plus-100% subsidies scenario results, 3% discount rate: Period two production by fuel*

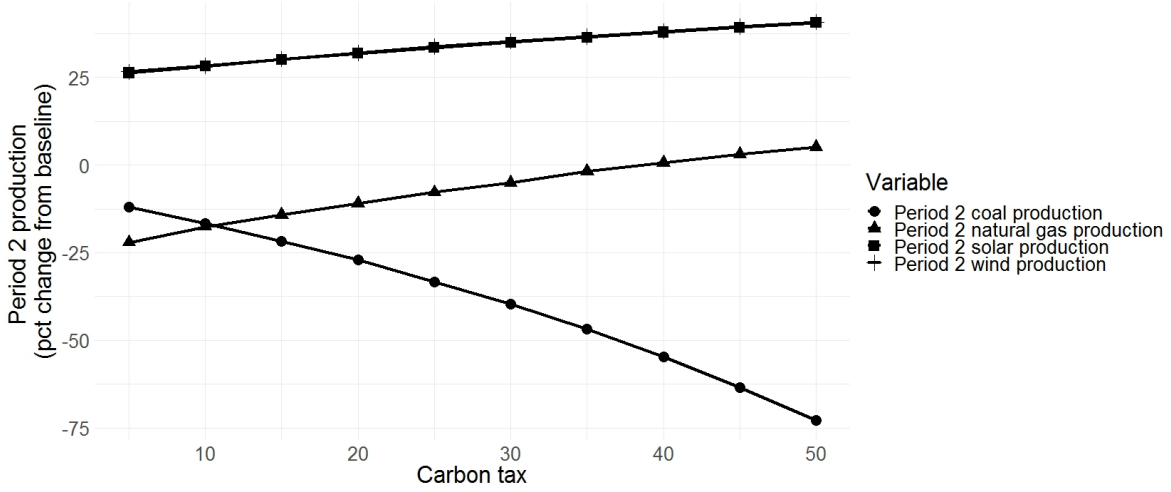

Note: This plot shows the change in period two production that is chosen when a carbon tax is imposed and subsidies equal to 100% of the cost of clean technologies are given to clean technologies, during that period. Results are presented as a percent from the baseline, no-tax scenario. Results are shown by carbon tax. These results are obtained with an annual discount rate of 3%.

Figure 2.22: *Tax vs. Tax-plus-10% subsidies scenario results, 3% discount rate: Change in research spending*

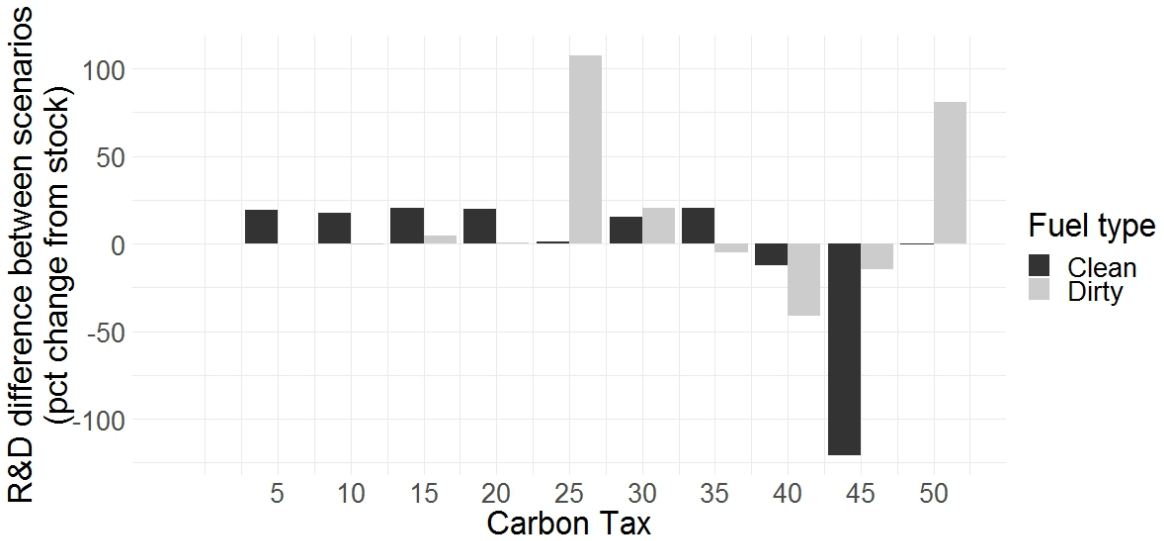

Note: This plot shows the change in research spending between the tax and tax-plus-10% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios. These results are obtained with an annual discount rate of 3%.

One aspect of these results that may warrant further investigation is the large amount of research spending in the tax-plus-100% subsidy case for low carbon taxes (2.24). Intuitively, the emissions cost of small carbon tax is low, and the subsidy should induce enough fuel-switching to reduce that cost even more. One potential explanation is that the optimization started with a large amount of research spending in this case and then did not face enough of an incentive to reduce that parameter before the optimization stopped, given the "slack" of the cost savings from the subsidy.

Overall, results obtained under a lower discount rate match predictions and therefore provide a useful robustness check on qualitative results. They also highlight that a change to inter-temporal social valuation has implications for research spending, which may exacerbate pre-existing distortions in research incentives.

Figure 2.23: *Tax vs. Tax-plus-50% subsidies scenario results, 3% discount rate: Change in research spending*

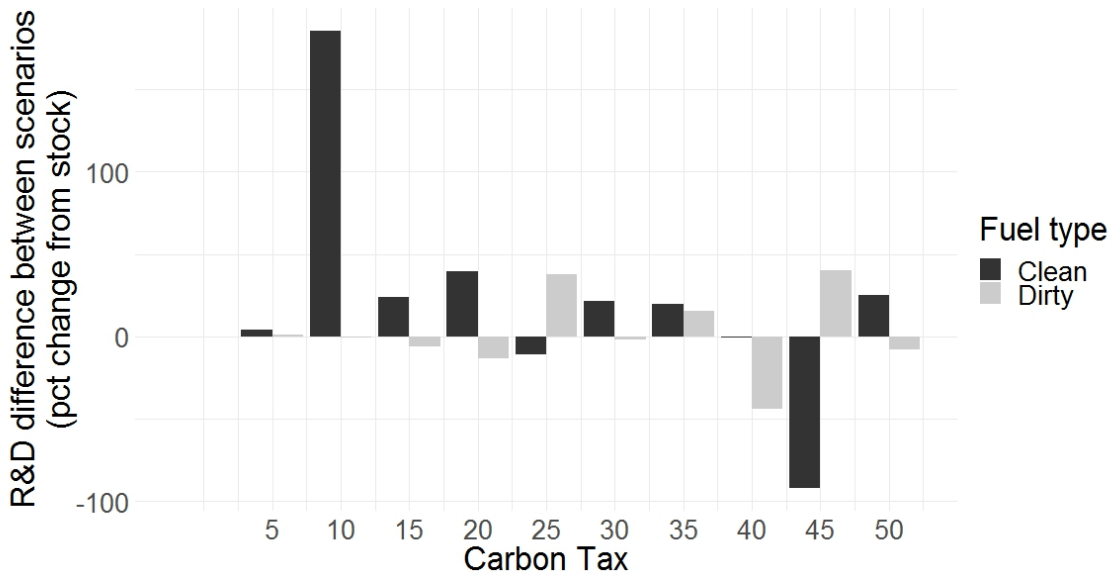

Note: This plot shows the change in research spending between the tax and tax-plus-10% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios. These results are obtained with an annual discount rate of 3%.

2.8 Conclusion

Climate policy as implemented is often far from first-best recommendations. One second-best policy outcome that is quite common is the combination of a carbon price, through a carbon tax or cap-and-trade system, and subsidies for electricity production from low-carbon sources. The justification for this policy combination is that low-carbon technologies are so expensive that they are not competitive with carbon-emitting technologies, and that subsidies will help bring down the cost of those technologies by incentivizing innovation. Empirical literature provides little support for specific market failures that these subsidies correct. A high carbon price should, in itself, incentivize development of low-carbon technologies and bring them to market. There is a known static cost of using such subsidies, as they increase the cost of emissions abatement in the short-term. This paper asks: Do subsidies for specific low-carbon technologies distort incentives for innovation, leading to additional costs in the

Figure 2.24: *Tax vs. Tax-plus-100% subsidies scenario results, 3% discount rate: Change in research spending*

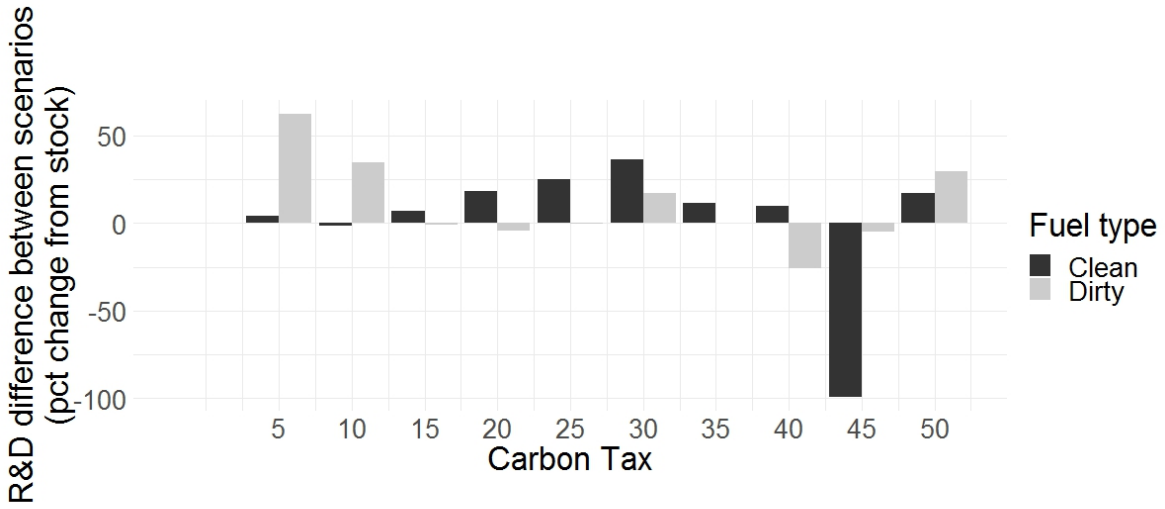

Note: This plot shows the change in research spending between the tax and tax-plus-10% subsidies scenarios. The change in spending is the difference in spending, as a percent of the knowledge stock, between scenarios. These results are obtained with an annual discount rate of 3%.

long-run?

With a two-period analytical model, I show that, in a carbon tax setting, whether subsidies for clean energy will change incentives for unsubsidized low-emissions ("clean") or emissions-intensive ("dirty") technology depends on the marginal returns to innovation in production. In short, subsidies alter production incentives for the electricity sector, which has a tight total production constraint to meet inelastic demand. Long-run innovation changes will be impacted by subsidies if innovation is tied to production in any way. Analytical results for the carbon tax case are therefore ambiguous and reflect recent work showing that innovation incentives are from a combination of input price changes and their resulting impact on input substitution Acemoglu (2003). In a cap-and-trade setting, the incentives for innovation will be changed by subsidies through their impact on production, as in the tax case, but also by their impact on the endogenous carbon price. Ceteris paribus, a subsidy for clean energy will lead to more electricity production from clean and less from dirty sources. This will reduce the

carbon price in the cap-and-trad system, which reflects the marginal cost of abatement. In turn, the incentive for developing technology to further reduce emissions is reduced.

With a simulation of the electricity sector calibrated to EU data, I examine changes in research spending that occur when the sector faces either a carbon tax and clean energy subsidies or a carbon cap constraint and clean energy subsidies. In the case of a carbon tax, small subsidies (on the order of 10% of the cost of production from clean sources) decrease research spending for subsidized clean technology for small carbon taxes ($10 - $20 per ton of CO2). A subsidy reduces the need to reduce technology cost to allow for fuel substitution. Larger subsidies (on the order of 50% of clean technology production cost) cause research spending for dirty technology to be reduced when carbon taxes are higher ($30-$40 per ton of CO2). When the sector must pay for emissions and would normally invest research to develop cleaner technology, subsidies reduce the need for this research. Very high subsidies (100% of the cost of clean technologies) incentivize higher research spending on subsidized, clean technologies and/or reduce the incentive for research spending on dirty technology (which can be interpreted as developing new clean technology), for carbon taxes in the range of $30-$50 per ton of CO2.

A subsidy for clean technology in the presence of a cap-and-trade system reduces the incentive for research, generally. There is less of an incentive to reduce the cost of clean technology or the emissions rate of dirty technology. Interestingly, this effect dissipates in subsidy size, implying that, when large subsidizies are adopted in a cap case, it is more cost-effective, on the margin, to abate emissions through research spending than through fuel-switching. This may be due to the fact that subsidies have forcibly caused fuel-switching to a large extent.

Broadly, analytical and simulation results point to clean energy subsidies, when overlapped with a carbon price, altering the incentives for innovation compared to a carbon-price-only scenario. The potential for discouraging development of less emissions-intensive technology is particularly large when the carbon pricing policy is a cap-and-trade system. Results suggest that policymakers should focus on creating stringent caps and adjust research policy to address

market failures, when identified.

Chapter 3

Induced innovation, market size, and total value to firms: The case of US electricity

3.1 Introduction

Empirical research on induced innovation is often rooted in a famously-quoted remark from (Hicks, 1932, p. 124): "A change in the relative prices of the factors of production is itself a spur to invention, and to invention of a particular kind—directed to economising the use of a factor which has become relatively expensive." Hick's discussion of induced innovation provides a macroeconomic view of innovation in light of changes in the relative supply of factors (inputs) on which national income relies.[1] How that famous statement, called the "induced innovation hypothesis", translates into microeconomic predictions is setting-depending.

Jaffe et al. (2004) discuss that the microeconomic theory on induced innovation is poorly developed. Acemoglu (2003) provides some useful guidance. He shows that producers of intermediate goods demand new technology for a given input if their relative reliance on that

[1]The objective of his analysis is to address whether innovation will alter the share of income that is paid to labor (instead of owners of capital), out of concern for income distribution.

input increases. This effect is mitigated by the relative price of factors. Similarly, the incentive for upstream inventors to produce new technology is a function of the size of the market of the factor (an input for downstream producers) for which new technology will be developed. Again, this incentive is mitigated by the price those factors. In short, the direction of innovation should follow *total profit* changes for producers and *total value* potential for innovators.

In this note, I use the setting of the US electricity sector to illustrate that relative input prices may not lead to innovation devoted to the relatively more "expensive" input, defined in terms of prices only. I provide a pared-down version of Acemoglu (2003)'s model to clarify the incentives for technology adoption. The model shows that, in a microeconomic context, the meaning of "expensive" to firms is the total cost of an input, which includes its firms' reliance on the input in addition to its price. (Acemoglu (2003) calls the former the "market size" effect.) I discuss the importance of heterogeneity in firms' elasticities of substitution and its implications for estimation of induced innovation studies. Empirical estimation that does not consider this heterogeneity and relies on input price shocks exclusively can lead to counter-intuitive results. I argue that induced innovation studies should consider the relevant treatment from an input price (or supply) shock to be total profit shocks to producers.

I look at innovation changes that have occurred among electricity-producing firms in the U.S. in the wake of the circa-2008 "fracking boom", which has provided the industry with a marked increase in the supply of natural gas over the last twelve years. In addition to highlighting the shortcoming of framing input price shocks as the sole driver of innovation changes, this setting includes an exogenous shock to the supply of an input, which is important for obtaining causal estimates of the role of input shocks on innovation outcomes (Acemoglu and Linn, 2004).[2]

This note contributes to induced innovation literature, broadly, as well as recent work that highlights the role of "relative prices" that includes the market size effect. Early empirical work on induced innovation includes that by Popp (2002) and Newell et al. (1999), who examine

[2] I assume that the development of technology for fracking natural gas out of shale deposits, which led to an increase in the availability of natural gas, is only correlated with the development of technology to augment the use of natural gas for downstream producers though the gas supply shock.

the role of energy prices in inducing energy-saving innovation. More recently, Lewis (2011) shows that the availability of low-skilled workers reduced the incentive for US manufacturing firms to adopt labor-saving technology. Acemoglu and Linn (2004) find that innovation for new drugs is responsive to their potential market size. Bustos (2011) demonstrates that, after a tariff change, Argentinian firms updated their technology after being exposed to expanded product markets. This note presents a context similar to that of Hanlon (2015), who finds that, after a reduction in supply of US cotton (from the Civil War), the British textile industry directed innovation toward using an alternative cotton source. After the relative input supply (and therefore price) shock, innovation was directed toward using the more abundant (and relatively cheap, in pure price terms) input.

Much empirical attention has been focused on induced innovation for understanding how market-based environmental policies, which send price signals to producers regarding the social value of pollution abatement, will direct innovation for pollution abatement technology.[3] This paper is highly relevant to this literature, as, in the long run of ambitious environmental policy (such as a high carbon price for addressing climate change), relative factors shares will tend toward less-polluting (such as less carbon-intensive) factors. Innovation, and therefore pollution abatement costs, be driven by not just the price incentives from policy but also the change in the relative abundance of polluting factors. Therefore, an understanding of how innovation is driven by total profit changes for producers and total value potential for innovators is important for environmental policy. Because I am specifically looking at the role of a change in the supply of inputs on innovation in a carbon-intensive industry (electricity) here, this note provides a useful step in that direction.

In the remainder of this note, I first provide, in Section 3.2, a simple version of the model in Acemoglu (2003) that clarifies the microeconomic role of input prices and reliance on inputs in innovation incentives. I then discuss the role of heterogeneity in firms' elasticities of substitution in generating innovation outcomes. In Section 3.3, I give an overview of the US

[3]In addition to Popp (2002) and Newell et al. (1999), see, for example, see Jaffe and Stavins (1995), Linn (2008), Noailly and Smeets (2015), Calel and Dechezlepretre (2016), and Aghion et al. (2016).

electricity sector relevant to this discussion. In Sections 3.4 and 3.5, I describe my data sources and then provide some descriptive statistics that suggest that innovation in the sector is more sensitive to the "market size" effect than the change in input prices, due to high elasticity of substitution between coal and natural gas in electricity production. I then discuss estimation strategies for induced innovation more generally and how one may appropriately estimate the role of input price changes on innovation in this setting, in Section 3.6, before concluding.

3.2 Framework

3.2.1 The market size effect and the role of the elasticity of substitution in innovation incentives

Say a profit-maximizing firm has the following objective function:[4]

$$\max_{A,B} \pi = pY - \omega_A A - \omega_B B \tag{3.1}$$

in which p is the price of the firm's output, Y, and ω_A and ω_B are the prices of its two inputs A and B, respectively.

Assume that the production function has a CES form:

$$Y = \alpha \left(\gamma A^{(\epsilon-1)/\epsilon} + (1-\gamma) B^{(\epsilon-1)/\epsilon)} \right)^{\epsilon/(\epsilon-1)} \tag{3.2}$$

in which ϵ is the elasticity of substitution.

Taking first-order conditions of Equation (3.1) with respect to A and B and solving for input prices ω_A and ω_B allows one to easily obtain:

$$\frac{B}{A} = \left(\frac{(1-\gamma)}{\gamma} \frac{\omega_A}{\omega_B} \right)^{\epsilon} \tag{3.3}$$

Equation (3.3) shows that input quantities and prices are inversely related. Additionally, the relative quantity of inputs that the firm chooses is a function of its elasticity of substitution.

[4] I use notation similar to Acemoglu (2003) to easily refer to some of his conclusions.

When an exogenous change to input prices occurs (which can be the result of a supply shock), the firm will adjust its input ratio according to Equation (3.3). Producers for which ϵ is relatively large will be able to respond to input price shocks with input substitution more readily than others.

How does this relate to incentives for innovation? Note that an input price shock will not just require producers to adjust to $\frac{\omega_A}{\omega_B}$, as Hick's induced innovation hypothesis is commonly interpreted as suggesting, but it will require producers to adjust to $\frac{\omega_A}{\omega_B}$ and its impacts on $\frac{B}{A}$, Y, and p (the latter of which depends on producer heterogeneity and the competitiveness of the output market).

Using a more elaborate version of the model presented here, Acemoglu (2003) models inputs as being produced from a primary input and technology that augments the input. Taking the same approach, let:

$$A = \frac{1}{1-\beta}\left(x_A^{(1-\beta)}\right) I_A^{\beta} \tag{3.4}$$

and the "production" of B have an identical function. In Equation (3.4), x_A is the amount of technology chosen to "augment" input I_A (make the use of I_A more efficient or otherwise rely on I_A less). Acemoglu (2003) shows that, if A and B are competitively produced in an intermediate market, the demand for technology that augments input I_A is given by:[5]

$$x_A = \left(\frac{\omega_A}{\chi_A}\right)^{(1/\beta)} I_A \tag{3.5}$$

Equation (3.5) says that the demand for technology to augment I_A, from which A is produced, depends directly on the amount of I_A used and the price of A, ω_A. The higher the price of a factor, ω_A, the more technology will be demanded for that factor's upstream production from a more primary input. (Ultimately, the adoption of this technology translates in a reduction of ω_A for downstream producers.) Hicks (1932) highlights that relative factor prices ($\frac{\omega_A}{\omega_B}$), which are a function of the relative reliance on factors in ultimate production

[5]See Eq. (10) in Acemoglu (2003). Acemoglu's model includes, additionally, many different (j) types of technology ("machines") available to augment each input.

(Equations (3.3) and (3.1)), direct innovation in this way. Equation (3.5) also highlights, as Acemoglu (2003) points out, that factor demand also directs innovation. Greater usage of I_A leads producers to demand more technology to augment I_A, as well as further upstream developers of that technology to develop it for I_A consumers (not shown here). Acemoglu (2003) calls this the "market size" effect. Note that, because (3.5) can be written as a function of A, innovation follows demand for the downstream input.

The bottom line from this analysis is that factor prices alone do not drive the direction of innovation. Innovation is directed both by factor prices and firms' relative reliance on factors. These of course are related, but they are importantly related through the elasticity of substitution for a given producer (or industry). A given input price shock will impact producers' incentive to adopt innovation the less they are able to substitute away from new relatively expensive inputs. There exists, for every β (which captures the role of existing technology in production from a given input), a critical elasticity at which an input price shock will have no impact on innovation incentives.

3.2.2 Heterogeneity in elasticities of substitution

Some innovation studies look at innovation outcomes from many different industries as a function of a given input price shock. For example, Popp (2002) estimates changes in patenting for energy-related technologies as a function of energy prices.

An important determinant of the impact of input price changes on innovation will not only be how within-industry elasticities of substitution dictate the total change in input cost producers feel from the input price shock but also the heterogeneity of these elasticities across industries. As a first-order consideration, differences in these elasticities across industries (and the pre-existing technology available for augmenting inputs for their specific uses in each industry, as shown above) will determine heterogeneity in innovation outcomes across industries.

Additionally, heterogeneity in elasticities of substitution impacts demand adjustments from an input price shock. As a silly example, if the world supply of bananas succumb to

a disease shock that reduced banana availability, consumers with preferences that include appreciation of a wide variety of fruit may easily substitute to other fruits when bananas become rare and expensive. However, bakeries known for their banana bread may demand more banana-augmenting technology in their production process that allows the flavor of a given banana to extend to more loaves. There are important demand effects on the input price due to this heterogeneity: If both consumers and banana bread producers had low elasticities of substitution, the price of bananas would remain high after the onset of the banana disease. However, a large fruit substitution among consumers reduces the equilibrium banana price post-shock, benefiting banana bread producers. The most advantageous position for a producer would involve having a high elasticity of substitution when producers in other industries reliant on a given input cannot easily substitute toward using it after a negative price shock.

3.3 The US electricity industry and the natural gas fracking boom

As of 2018, nearly 35% of the electricity produced from utility-scale power plants was from combustion of natural gas.[6] Electricity produced from coal amounted to just under 30% of utility-scale production. Natural gas and coal are not only the two largest fuel sources for electricity production in the US, but the extent to which production relies on each is dictated by the relative price of these fuels. In general, either natural gas or coal plants are the marginal plants on which production will rely.[7,8] This relationship is unusually sensitive in a give time

[6]The US Energy Information Administration (EIA) defines utility-scale electricity generation as "electricity generation from power plants with at least one megawatt (or 1,000 kilowatts) of total electricity generating capacity." This designation includes the vast majority of electricity produced in the US, excluding only very small power plants, electricity produced on-site at industrial facilities, and cogeneration as a byproduct of burning fuel for heat.

[7]Conditional on production from other sources, as well as demand. The other major sources of electricity, nuclear power and renewable sources (wind, solar, and hydroelectric) have near-zero marginal costs of operation and are used according to a pre-defined schedule (in the case of nuclear and sometimes hydroelectric power), stochastic weather patterns (in the case of wind and solar), or as a supplemental resource to meet "peak" demand (some hydroelectric power).

[8]Cullen and Mansur (2017) provide a useful description of heterogeneity in natural gas and coal plants that determines which plants are marginal.

frame compared to other industries, as the non-storable nature of electricity requires supply to meet demand, instantaneously.

Widespread availability of new technology to extract (or "frack") natural gas from previously unreachable shale deposits generated a natural gas supply shock to the US, starting in 2008. Between January 2007 and 2009, the supply of natural gas increased by 11%. By January 2011, supply had increased 21% from 2007. Importantly, this supply change has been sustained. Between January 2007 and 2019, supply increased 77%.[9] This supply shock resulted in about a 33% drop in the price of natural gas delivered to electricity producers between just 2008 and 2009. Again, the impact was sustained. The price of natural gas continued to fall over time, reaching a 2018 annual average price that was about half of the 2008 value. In contrast, the price of coal fell only about 8%, 2008-2018 (Figure 3.1). Cullen and Mansur (2017) note that, though some of the initial drop in US natural gas prices in 2008 was due to the global recession, US prices remained low for years afterward, whereas European natural gas prices recovered.

This supply shock was due to a technological change exogenous to other industry trends. It is empirically viewed as a major unanticipated industry event; a number of authors have used the change in natural gas prices to understand how carbon pricing would impact emissions outcomes from electricity production (Cullen and Mansur, 2017; Knittel et al., 2015; Fell and Kaffine, 2018). For this paper, I assume the development of technology for using natural gas and coal in electricity production has only been impacted by the advent of fracking technology through the subsequent supply shock.

In tandem with the availability of relatively cheap natural gas, electricity producers have substituted away from using coal and toward using natural gas for electricity production. Figure 3.2 shows that, though the industry was making this substitution in the eight years before the supply shock, this substitution became more dramatic in the subsequent decade: The average annual change in the share of natural gas used in electricity production is 1.4% in

[9] Author's calculations from EIA *Short Term Energy Outlook* data on US natural gas dry production, obtained in panel form via personal request from the EIA.

Figure 3.1: *Prices of inputs delivered to US electricity producers*

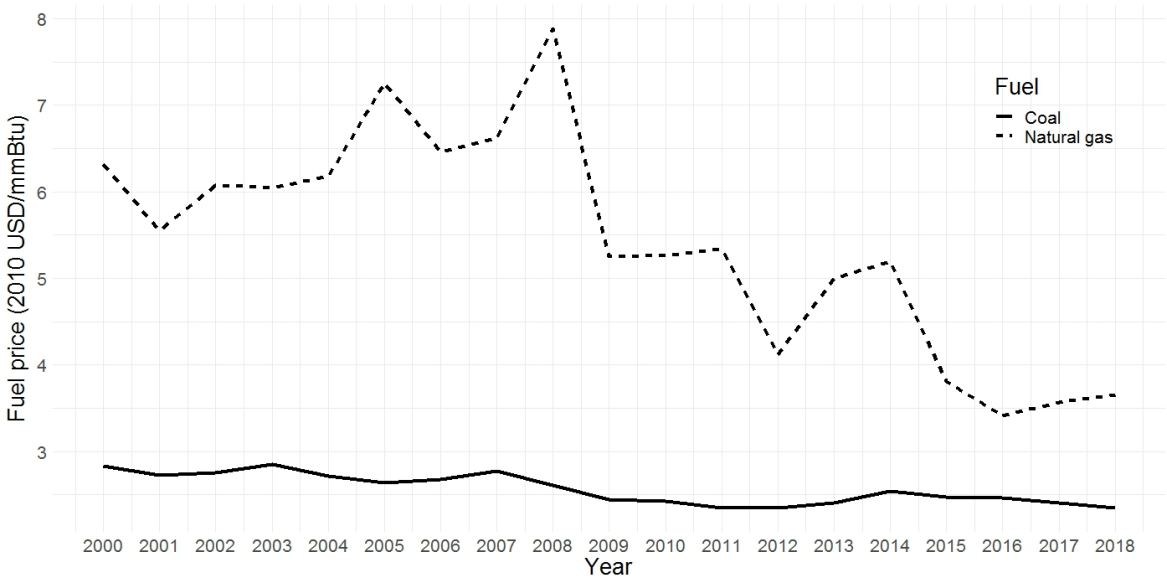

Note: This plot shows the average annual price of coal and natural gas delivered to electricity producers, 2000-2018.

the years 2008-2018, compared to 0.75% in the years 2000-2008. Knittel et al. (2015) estimate that the elasticity of the share of natural gas in fuel consumption for electricity production is about -1.2 for plants and -1.1 for firms that operate in traditionally regulated markets, indicating large input substitution in this industry. These results also indicate that input substitution occurs both within plants and within firms. As described in Section 3.2, the direction of innovation directed toward augmenting inputs will be determined by relative changes in input prices and reliance on each input. For production from natural gas, any change in the adoption of practices or equipment to make natural gas combustion more efficient, for example, will be determined by the net effect of the decline in the natural gas price and the shift toward increased reliance on natural gas. For electricity production from coal, any change in innovation will be determined by the net effect of the relative increase in the price of coal and the reduced reliance on coal. Given the large input substitution in response to the input price shock in this industry, I anticipate the "market size" effect to drive innovation more than the the price effect in this industry. This means that innovative effort should be directed toward the use of natural gas after the natural gas price shock.

Figure 3.2: *Share of US electricity produced by coal and natural gas*

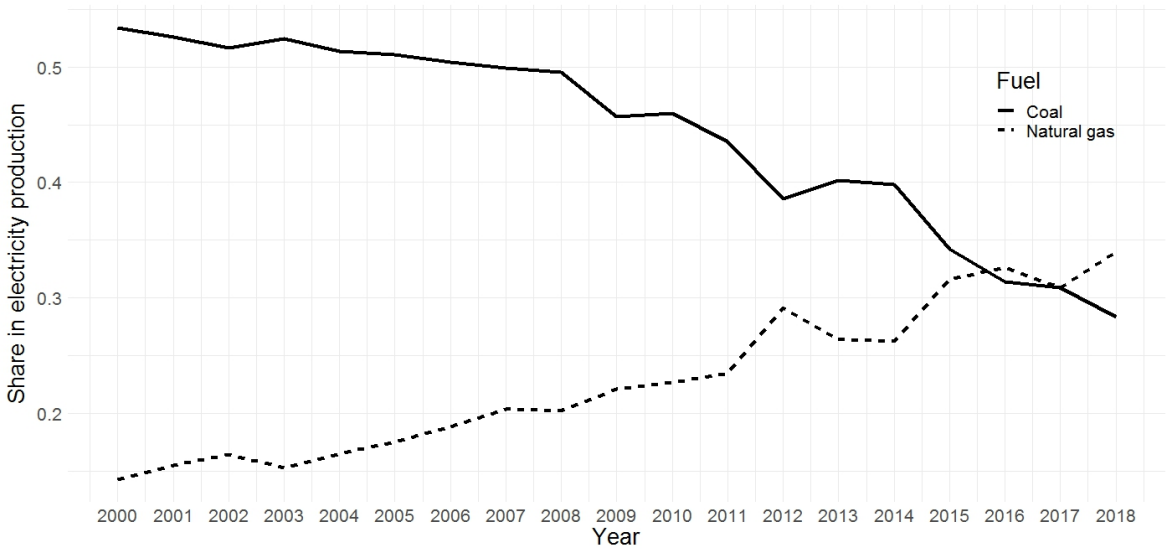

Note: This plot shows the annual share of coal and natural gas in electricity production, 2000-2018.

The electricity industry was responsible for 35% of the consumption of natural gas as of 2018. Natural gas is also consumed for other uses: heating, on-site electricity production, and as an input in other industries; residential heating; commercial heating; and transportation. Aside from consumption for electricity production, which grew about 59% 2008-2018, consumption of natural gas for industrial purposes, the other major area of consumption growth, grew about 28%. This indicates that electricity producers have been able to uniquely benefit from the reduction in the price of natural gas after the supply shock; the industry was able to substitute toward using the fuel while having a limited demand-side, upward impact on prices.[10]

The US electricity industry was historically vertically integrated and regulated until deregulation became politically popular in the 1990's. Through the 1990's and until 2002, many US states unbundled electricity production from transmission, distribution, and retail sale, forming wholesale electricity markets (Borenstein and Bushnell, 2015). The industry currently operates with wide geographic heterogeneity in regulatory status. Firms that own

[10]Author's own calculations from EIA Natural Gas Consumption by Sector data, Table 4.3 in the *Electric Power Monthly*, available at https://www.eia.gov/totalenergy/data/browser/index.php?tbl=T04.03#/?f=A&start=2000&end=2019&charted=1-2-9-13-14

and operate plants that participate in wholesale markets make operational and investment decisions based on a profit-maximizing objective function. Firms that own and operate plants that are regulated by public utility commissions (PUCs) are compensated according to justified costs. When it comes to technological improvement, owners of deregulated plants face a greater incentive to operate at the technological frontier to minimize costs, yet expected revenue—particularly from wholesale markets that have become increasingly dominated by very low marginal cost sources such as wind in the last decade—may reduce owners' financial ability to improve plants. While owners of regulated plants face less of an incentive to minimize costs than those of deregulated plants, they may be able to adopt new technology more readily than owners of deregulated plants if they are able to convince their respective PUCs that the costs are justified to meet the PUCs' goals. Fowlie (2010) found, for example, that owners of regulated coal-fired power plants responded to NOx regulations with more capital-intensive forms of abatement than owners of deregulated plants.

Innovation for both coal- and natural gas-fired power plants can be captured by their heat rate, a measure of their efficiency.[11] Coal plant operators can alter the heat rate of their plants by directing personnel to utilize practices that fine-tune plant operation (Bushnell and Wolfram, 2009). Owners of coal plants can update or install equipment that improves the heat rate of their plants, such new turbines or automated systems that control plant conditions (Linn et al., 2014; Nowling, 2015). Heat rate improvement is considered a first-order change that coal plant owners can take in response to new costs (Nowling, 2015). (Linn et al., 2014) find that coal plant owners adjust plant heat rates through capital-intensive means in response to coal price changes. Owners of natural gas plants can adopt combined-cycle technology to improve their efficiency (McGrath, 2017). For both types of plants, the heat rate directly reflects efficiency changes in input use. These plants have very simple production functions that only contain fuel and operations and maintenance costs as the major variable costs.

[11] A plant's heat rate gives the amount of heat energy from fuel required to produce a unit of electricity, so it is the inverse of efficiency.

3.4 Data

For this note, I rely on a number of publicly-available data sources on the US electricity sector. I gathered data for 2000-2018. The availability of plant regulatory status limits my panel to years 2003-2018.

I use the EIA's State Energy Data System (SEDS) for annual state average prices of coal and natural gas delivered to the electric generation sector.[12].

I calculate the annual share of coal and natural gas used in total electricity production from the data available in the EIA *Monthly Energy Reivew* Table 7.2b Electricity Net Generation: Electric Power Sector.[13]

I rely on the EIA Form-923 ("923") for annual data on fuel used and net electricity generated by natural gas and coal plants.[14] I merge the power plant operational data to plant characteristics data available from the EIA Form-860 ("860"), including plant capacity, age of units, and retirement status. I obtain regulatory status, however, from the 923 fuel receipts data, 2008-2018, and the 860 data, 2003-2007.[15]

I define natural gas plants as having at least one natural gas unit and no coal units. This ensures that I can assess how plant owners uniquely respond to input price changes in terms of updating their plants, as one cannot reliably calculate unit-level heat rates with EIA data.[16] I define coal plants similarly. I define the fuel type of units by their "primary fuel source", the fuel source on which they most rely for production, per year, according to the EIA 860 designation.(Chan et al. (2017) use the same rule.)

---

[12]https://www.eia.gov/state/seds/seds-data-complete.php?sid=US#CompleteDataFile

[13]Available at https://www.eia.gov/totalenergy/data/browser/index.php?tbl=T07.02B#/?f=M

[14]All EIA form data is available for public download from the agency's website. Currently each form has a page with historical data; the 923 form data, for example, can be found at https://www.eia.gov/electricity/data/eia923/.

[15]A "unit" is a boiler-generator pair at an electricity plant. A plant typically has several units.

[16]Based on personal communication with EIA personnel.

3.5 Some descriptive statistics

What do trends in US natural gas and coal plant heat rates (efficiency) indicate about the direction of innovation, post-fracking boom?

Figures 3.3 and 3.4 show the annual distributions of the heat rates of US natural gas plants, 2003-2018, for deregulated and regulated plants, respectively. Evidence of technological improvement would be declining heat rates over time. Interestingly, while heat rates values are declining over time in both groups, there is more of a substantial reduction in heat rate values for regulated plants. As discussed in Section 3.3, improvement of natural gas plant efficiency typically requires a large investment in an updated unit, which may not be profitable for plant owners competing in contemporary wholesale electricity markets. However, these investments could be justified under cost-of-service regulatory review. Figures 3.5 and 3.6, which exclude entrants after the natural gas boom[17], suggest that any technological improvement among deregulated plants takes the form of new, more efficient plants. Among regulated plants, in contrast, there is evidence of plant-level technological improvement.

Figure 3.3: *Heat rates of deregulated natural gas plants*

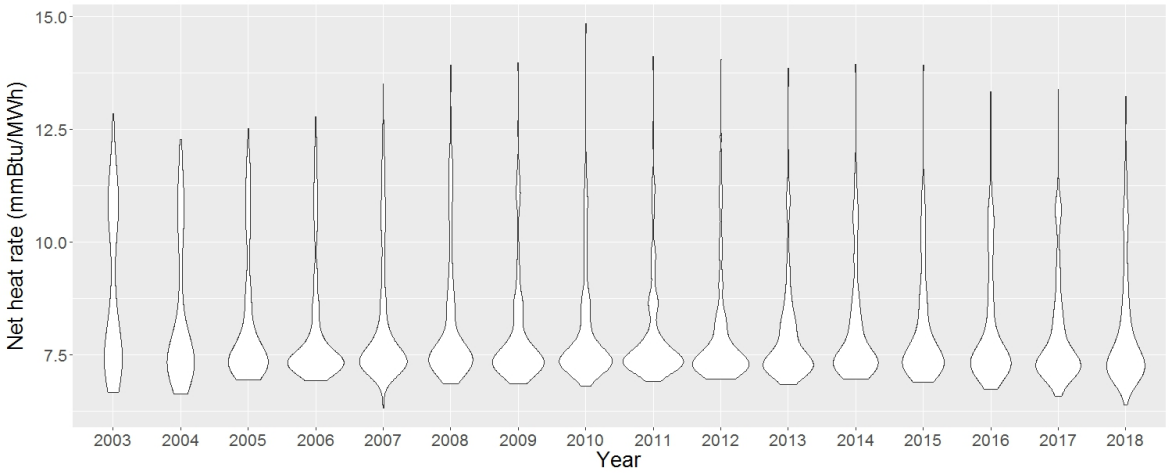

Note: This plot shows the mirrored distribution of annual heat rate observations for deregulated natural gas plants, 2003-2018.

[17] I define "new entrants" as plants at which the oldest generation unit was listed as operational in 2008 or later.

Figure 3.4: *Heat rates of regulated natural gas plants*

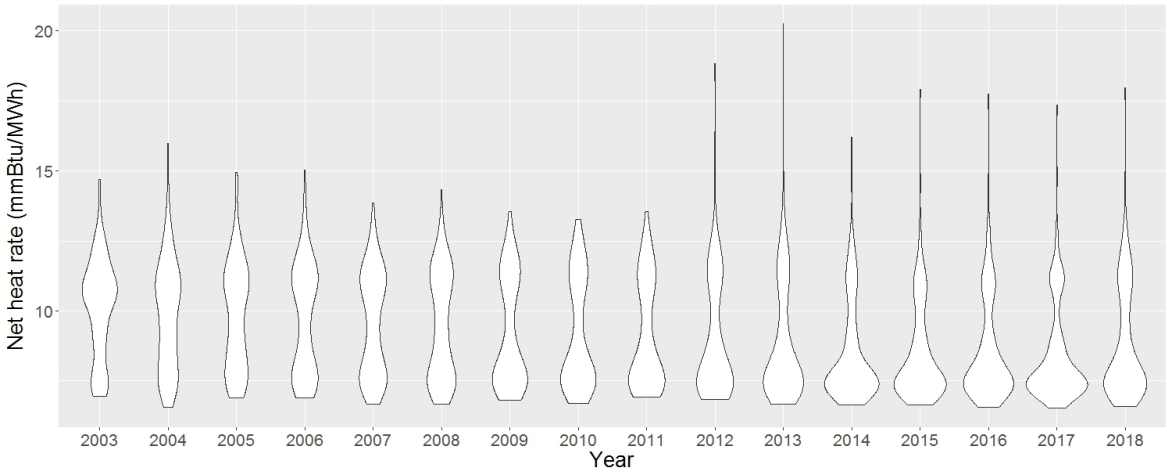

Note: This plot shows the mirrored distribution of annual heat rate observations for regulated natural gas plants, 2003-2018.

Figure 3.5: *Heat rates of deregulated natural gas plants, without post-2007 entrants*

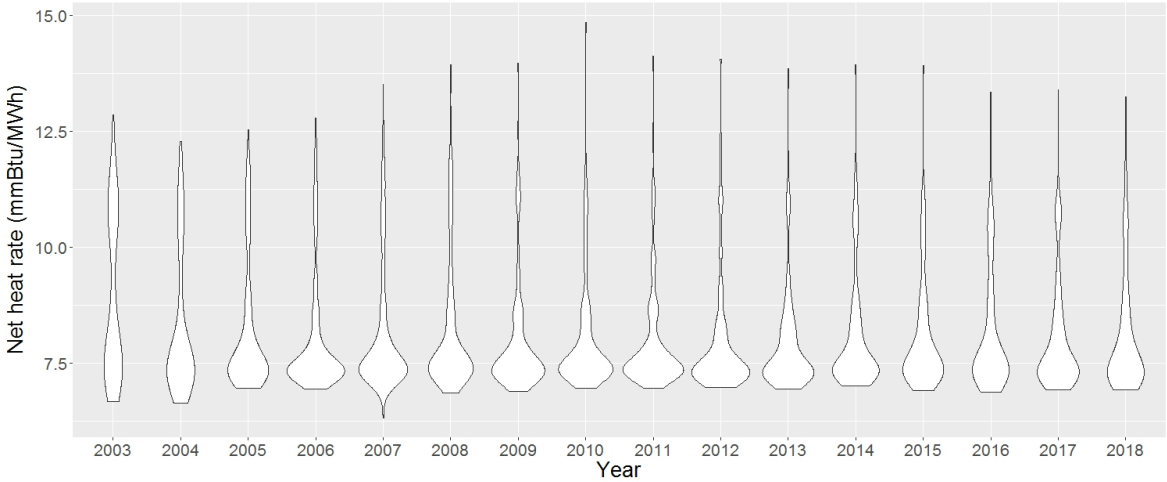

Note: This plot shows the mirrored distribution of annual heat rate observations for deregulated natural gas plants, 2003-2018. I exclude new entrants, plants for which the oldest unit has an operational year of 2008 or later.

Are there indeed plant-level trends in heat rates at natural gas plants that may be indicative of plant owners updating their units? Table 3.1 shows the results of estimating the following

Figure 3.6: *Heat rates of regulated natural gas plants, without post-2007 entrants*

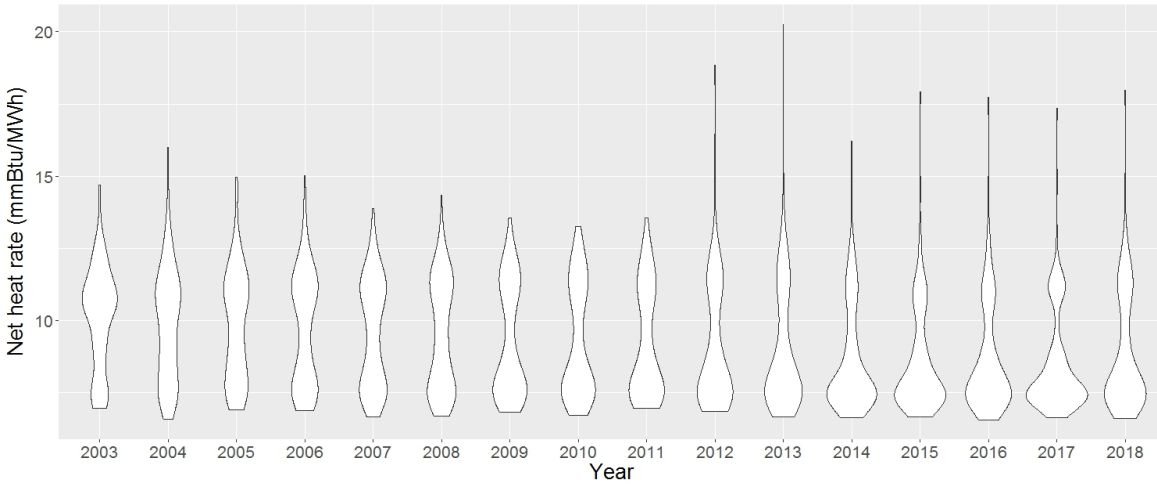

Note: This plot shows the mirrored distribution of annual heat rate observations for regulated natural gas plants, 2003-2018. I exclude new entrants, plants for which the oldest unit has an operational year of 2008 or later.

simple panel model of natural gas plant heat rates:

$$log(HR_{it}) = \beta_0 + \gamma_t + \beta_{CF} log(CF_{it}) + \omega_i + \epsilon_{it} \tag{3.6}$$

To estimate Equation (3.6), I regress the log heat rate of natural gas plant i in year t on a time trend (γ_t); the plant's log capacity factor ($log(CF_{it})$), or output; and a plant fixed effect (ω_i).[18] The coefficient of interest is on the time trend, as I am interested in whether there is a significant, average plant-level change in the heat rate of natural gas plants during this time period. I estimate Equation (3.6) with regulated plants only, excluding new entrants, as this group displays the strongest trend in heat rates. I include the capacity factor of plants because natural gas and coal plant heat rates are mechanically related to their output rate. Plants are designed to operate at an "optimal" output rate below or above which their efficiency suffers. Increased usage of natural gas plants after the post-2008 supply shock could drive efficiency

[18] I calculate the heat rate of plants in terms of mmBTUs of heat input over MWh of output during t. The capacity factor is a standard measurement of the output rate of power plants. I calculate capacity factor by dividing a plant's realized electricity production in t in MWh by its potential output in t. Potential output is the product of its maximum rated capacity, or electrical output, for a given hour in MW and the number of hours in t. Therefore $CF_{it} = \frac{MWh}{(MW)*hours}$.

trends. I cluster standard errors by plant to account for observations being correlated within plant, over time.

Table 3.1: *Regulated natural gas plants heat rate trends*

	Dependent variable:
	Log net heat rate*100
Time trend	-0.25^{***}
	(0.09)
Log capacity factor	-5.00^{***}
	(1.21)
Observations	2,335
Adjusted R^2	0.92

Note: *p<0.1; **p<0.05; ***p<0.01. Results of estimating Equation (3.6), with the log net heat rate multiplied by 100. New entrants, plants for which the oldest unit has an operational year of 2008 or later, are excluded from the panel. Standard errors are clustered by plant.

The negative time trend in Table 3.1 indicates that average regulated natural gas plant-level heat rates are decreasing over the sample period, conditional on plants' output rates. The trend indicates there was an average of a 0.0025% improvement in plant heat rates for regulated plants each year 2003-2018, implying a mean change of 0.0375% over the 15-year period. This improvement is small on average, given that replacement of single-cycle with combined-cycle units, one option for reducing natural gas plant heat rates, can yield a heat rate improvement of 20%.[19] Still, some within-plant improvements are suggested by this result.

Figure 3.7 compares the heat rates of coal plants by regulatory status, over time, excluding plants slated for retirement. There is no indication that there is trend in the heat rates of coal plants over time, in either group.

In summary, in an industry setting in which there has been a large and sustained relative

[19]The EIA provides a quick reference on power unit heat rates by type here: https://www.eia.gov/electricity/annual/html/epa_08_02.html

Figure 3.7: *Heat rates of coal plants not slated to retire, by regulatory status*

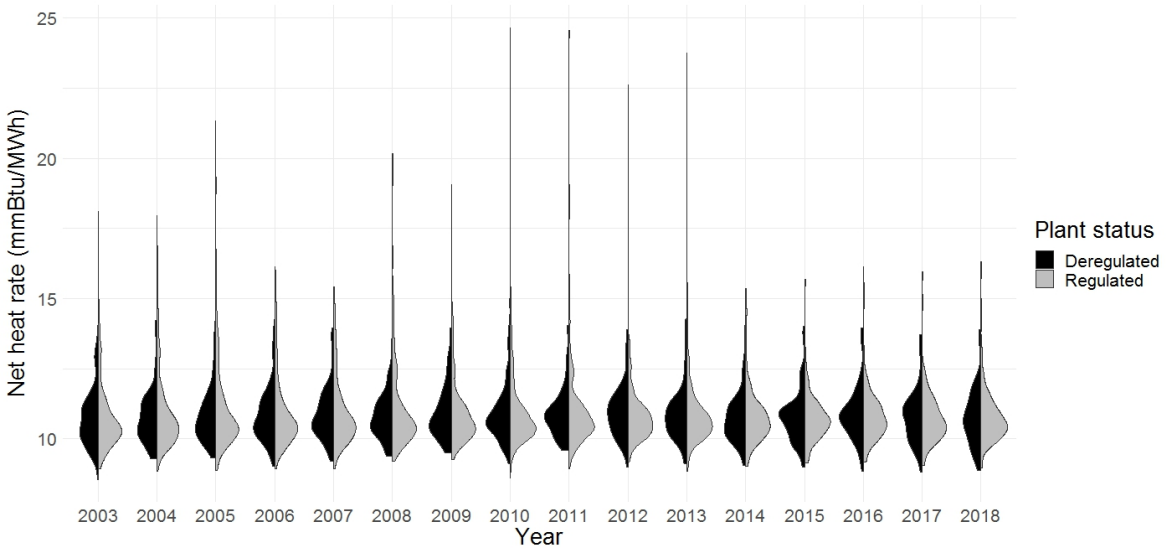

Note: This plot shows the distribution of annual heat rate observations for deregulated and regulated coal plants, 2003-2018. I exclude plants that have reported plans to retire or that retire without announcement.

input price shock, there is little evidence of technological improvement in response. A simple interpretation of Hicks (1932) leads one to suspect that coal plant owners, who have been exposed to an increase in the price of their input relative to the price of the competing input, natural gas, would adopt "coal-augmenting technology". In this setting, such technology adoption would likely be evident in plants' heat rates, which measure the efficiency of coal use.

The extension of Hicks (1932) that Acemoglu (2003) provides—that it is a combination of a price shock and the change in the market size of inputs that directs innovation—makes these trends less surprising. While coal has become relatively more expensive, the reduced reliance on coal, through both the intensive and extensive margins of, respectively, coal plant owners running plants running less and retiring units or plants, may dominate the incentives for adopting new technology. Looking at Equation (3.5), the incentive for coal plant owners to adopt "coal-augmenting technology" should actually be reduced, post-2008. Though coal became relatively more expensive post-2008, it became, absolutely, cheaper. And, the use of coal declined. The combination of the absolute reduction in the price of coal (through ω_A) and reduced reliance on coal (through I_A) leads to an overall reduction in the magnitude of

the right-hand side of the expression.

A simple interpretation of Hicks (1932) would also lead one to hypothesize that natural gas plant owners, facing a reduction in price of natural gas compared to the competing fuel, would reduce demand for "natural gas-augmenting technology". However, despite this relative price change (Figure 3.1), Figures 3.3 and 3.4 show that the fleet of natural gas plants have become more fuel-efficient during this time period. In a recent article, the EIA notes that a significant share of new natural gas units installed after 2013 are advanced combined cycle units, which are more efficient than combined cycle units.[20] This suggests that natural gas plant owners are moving toward the technological frontier along the extensive margin. With a simple OLS model of natural gas plant heat rates for plants that were built before the supply shock, I show suggestive evidence that plant owners have been updating plant efficiency along the intensive margin as well. Given the increased use of natural gas during this time period (Figure 3.2), the "market size" effect likely dominates the incentive for demand of natural gas technology.

3.6 Toward an estimation strategy

Based on the discussion above in Section 3.2, the ideal estimation strategy for estimating the causal effect of an input price shock on firm innovation would use, as the identifying right-hand-side variable, a firm-level measure of the total profit change that occurs from the price shock. A simple interpretation of Hicks (1932) leads one to estimate the elasticity of an innovation outcome with respect to input price changes, but this estimation strategy misses important firm- or industry-level heterogeneity in the elasticity of substitution between inputs, which determines the total profit changes on which innovation is based. For example, an estimation strategy of the form:

$$Y_{j,t} = f(\lambda(\frac{p_1}{p_2}), \mathbf{X}_{jt}, \theta_j) \tag{3.7}$$

in which $Y_{j,t}$ is innovation outcome Y for firm j at time t, $\lambda(\frac{p_1}{p_2})$ is a function of lagged

[20] See the June 19, 2019 "Today in Energy" post at https://www.eia.gov/todayinenergy/detail.php?id=39912.

input price ratios, $\mathbf{X}_{j,t}$ is important firm- and/or time-varying controls, and θ_j is firm fixed effects, will estimate the average impact of the change in input prices on the innovation outcome, across all firms. If a large proportion of firms in the sample have small elasticities of substitution, than the effect of the relative price shock will be to increase the cost of operation to most of the firms. Conditional on the ability of these firms to pass on this cost shock to consumers through the output price(s) of their goods, their innovative efforts will be directed toward reducing the cost of the new relatively "expensive" input, in pure price terms. If, however, a larger proportion of firms in this sample have large elasticities of substitution, then the "market size" effect of them relying more heavily on new, relatively less expensive inputs may dominate incentives for innovation. (This is also conditional on pass-through options.) The result will be innovation directed toward the new, relatively "inexpensive" input. Estimation of (3.7) with a heterogeneous set of firms in terms of their elasticities of substitution will yield a null result.

A blunt improvement to a naive estimation of Equation (3.7) is to subset estimation by firms that have similar elasticities of substitution. For example, if one suspects that US electricity producers have a higher elasticity of substitution than other producers, one should estimate Equation (3.7) for only those producers. The suggestive evidence above indicates that the "market size" effect would dominate innovation incentives for this industry, and the proper sign on the price coefficient from estimating Equation (3.7) should be negative.

A more precise improvement would be use Equation (3.5) to craft a firm-specific measure of the impact of an input price shock on innovation incentives, given how each firm adjusts its relative input use.[21] The simplist version would simply be to multiply input price by firm-level input quantity and control for other drivers of input substitution. Recent papers have used setting-specific proxies to capture the total profitability of technological development or adoption. Acemoglu and Linn (2004), who look at upstream innovation incentives (which are analogous to those in Equation (3.5)), examine how drug manufacturers bring new drugs to market in anticipation of the market size for those drugs. They construct a drug-specific value

[21] Equation (3.5) also highlights the importance of controlling for changing technology costs.

variable for identification, from changes in the demographics of likely drug consumers and their incomes. Bustos (2011) shows that firm entry into export markets after tariff reduction is a function of firm size. She is then able to estimate how firms adopted technology in response to changes in market size, using heterogeneity in firm size for variation in market size.

3.7 Conclusion

In this note, I've discussed the difference between Hick's "induced innovation hypothesis" as commonly referenced and innovation hypotheses that can be made in light of how input price changes translate into total profit changes for firms. Acemoglu (2003) provides a useful interpretation of Hicks (1932) and model to show that it is not just input prices but the total effect of input prices and input substitution that determines demand for new technology. This incentive structure has implications for econometric identification. In particular, I show how the estimation of an induced innovation model that uses just input price changes as the identifying variation may lead to results having the "wrong" sign or null results in various settings, depending on the elasticity of substitution of firms in the sample. Correct identification will use an understanding of elasticities of substitution of firms in the sample and/or a construction of right hand-side variables that capture the total profit shock experienced by firms after an input price shock.

I illustrate that elasticities of substitution matter through some descriptive statistics of trends in the US electricity production in the last decade. Since 2008, the price of natural gas has fallen relative to the price of coal, a competing input, due to an exogenous supply shock. This has led to the acceleration of input substitution in the industry, with the use of natural gas eclipsing that of coal. A look at the fuel efficiency, or heat rate, of natural gas and coal plants, which reflects electricity firms' adoption of processes and technologies to augment these inputs, suggests that the direction of innovation has not been toward the new, relatively "expensive" input (coal), in pure price terms, in the last decade. If any firms have been adopting new technology, it has been those managing natural gas plants. The high elasticity of substitution in this setting suggests that this may be due to the "market size"

effect driving innovation incentives.

A fruitful next step is to craft an estimation strategy to apply to the US electricity sector, to explore whether the input supply shock has led to innovation changes. The key challenge is crafting a measure of total profit changes attributable to the supply shock.

References

ABB (2009). Power generation: Energy efficient design of auxiliary systems in fossil-fuel power plants. Available at https://library.e.abb.com/public/5e627b842a63d389c1257b2f002c7e77/Energy%20Efficiency%20for%20Power%20Plant%20Auxiliaries-V2_0.pdf.

Acemoglu, D. (2003). Factor prices and technical change: from induced innovations to recent debates. *Knowledge, Information and Expectations in Modern Macroeconomics*, page 464–491.

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The environment and directed technical change. *The American Economic Review*, 102(1):131–166.

Acemoglu, D. and Linn, J. (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *The Quarterly Journal of Economics*, 119(3):1049–1090.

Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51.

Borenstein, S. and Bushnell, J. (2015). The us electricity industry after 20 years of restructuring. *Annu. Rev. Econ.*, 7(1):437–463.

Borenstein, S., Bushnell, J., Wolak, F. A., and Zaragoza-Watkins, M. (2019). Expecting the unexpected: Emissions uncertainty and environmental market design. *American Economic Review*, 109(11):3953–77.

Bushnell, J. B. and Wolfram, C. (2009). *The guy at the controls: Labor quality and power plant efficiency*, page 79–102. University of Chicago Press.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *American Economic Review*, 101(1):304–40.

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.

Chan, H. R., Fell, H., Lange, I., and Li, S. (2017). Efficiency and environmental impacts of electricity restructuring on coal-fired power plants. *Journal of Environmental Economics and Management*, 81:1–18.

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

Cicala, S. (2019). Imperfect markets versus imperfect regulation in u.s. electricity generation. Available at http://home.uchicago.edu/~scicala/papers/elec_gov_v_mkt/elec_gov_v_mkt_draft_2.pdf.

Coglianese, J., Gerarden, T., and Stock, J. H. (2018). The effects of fuel prices, regulations, and other factors on u.s. coal production, 2008-2016. *Working Paper*.

Cullen, J. A. and Mansur, E. T. (2017). Inferring carbon abatement costs in electricity markets: A revealed preference approach using the shale revolution. *American Economic Journal: Economic Policy*, 9(3):106–33.

EPA (2018). Documentation for epa platform v6 - chapter 5: Emission control technologies. Available at https://www.epa.gov/airmarkets/documentation-epa-platform-v6-chapter-5-emission-control-technologies.

EPRI (2018). Flexible operations and heat rate. Technical Brief 3002013992.

Fankhauser, S., Hepburn, C., and Park, J. (2010). Combining multiple climate policy instruments: how not to do it. *Climate Change Economics*, 1(03):209–225.

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.

Filsinger, T. (2010). Green energy outlook. *Public Utilities Fortnightly*, 148(1):54–56,58,66.

Fischer, C. and Newell, R. G. (2008). Environmental and technology policies for climate mitigation. *Journal of environmental economics and management*, 55(2):142–162.

Fischer, C. and Preonas, L. (2010). Combining Policies for Renewable Energy: Is the Whole Less Than the Sum of Its Parts? *International Review of Environmental and Resource Economics*, 4(1):51–92.

Fischer, C., Preonas, L. D., and Newell, R. G. (2017). Environmental and technology policy options in the electricity sector: Are we deploying too many?

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

Goulder, L. H. and Mathai, K. (2000). Optimal co2 abatement in the presence of induced technological change. *Journal of Environmental Economics and management*, 39(1):1–38.

Goulder, L. H. and Schneider, S. H. (1999). Induced technological change and the attractiveness of co 2 abatement policies. *Resource and energy economics*, 21(3):211–253.

Hanlon, W. W. (2015). Necessity is the mother of invention: Input supplies and directed technical change. *Econometrica*, 83(1):67–100.

Hicks, J. (1932). *The theory of wages,*. Macmillan.

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.

Holland, S. P., Mansur, E. T., Muller, N., and Yates, A. J. (2018). *Decompositions and Policy Consequences of an Extraordinary Decline in Air Pollution from Electricity Generation*. Number 25339.

Hubler, M., Schenker, O., and Fischer, C. (2015). Second-best analysis of european energy policy: is one bird in the hand worth two in the bush? *Center for European Economic Research (ZEW)*, Discussion Paper No. 15-079.

IEA, NEA, and OECD (2015). Projected costs of generating electricity.

Interagency Working Group on Social Cost of Greenhouse Gases, U. S. (2016). Technical support document: Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866. Available at https://www.epa.gov/sites/production/files/2016-12/documents/sc_co2_tsd_august_2016.pdf.

Jaffe, A. B., Newell, R. G., and Stavins, R. N. (2004). *Technology Policy for Energy and the Environment*. NBER Innovation Policy and the Economy. The MIT Press.

Jaffe, A. B. and Stavins, R. N. (1995). Dynamic incentives of environmental regulations: The effects of alternative policy instruments on technology diffusion. *Journal of environmental economics and management*, 29(3):S43–S63.

Knittel, C. R. (2002). Alternative regulatory methods and firm efficiency: Stochastic frontier evidence from the u.s. electricity industry. *The Review of Economics and Statistics*, 84(3):530–540.

Knittel, C. R., Metaxoglou, K., and Trindade, A. (2015). *Natural Gas Prices and Coal Displacement: Evidence from Electricity Markets*. Number 21627.

Korellis, S. (2014). Coal-fired power plant heat rate improvement options, part 1. Available at https://www.powermag.com/coal-fired-power-plant-heat-rate-improvement-options-part-1/.

Kotchen, M. J. and Mansur, E. T. (2014). How stringent are the us epa's proposed carbon pollution standards for new power plants? *Review of Environmental Economics and Policy*, 8(2):290–306.

Lemoine, D. (2017). *Innovation-led transitions in energy supply*.

Lew, D., Brinkman, G., Kumar, N., Besuner, P., Agan, D., and Lefton, S. (2012). Impacts of wind and solar on fossil-fueled generators: Preprint.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics*, 126(2):1029–1069.

Linn, J. (2008). Technological modifications in the nitrogen oxides tradable permit program. *The Energy Journal*, page 153–176.

Linn, J., Mastrangelo, E., and Burtraw, D. (2014). Regulating greenhouse gases from coal power plants under the clean air act. *Journal of the Association of Environmental and Resource Economists*, 1(1/2):97–134.

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.

McGrath, G. (2017). Natural gas-fired electricity conversion efficiency grows as coal remains stable. Available at https://www.eia.gov/todayinenergy/detail.php?id=32572.

Milliman, S. R. and Prince, R. (1989). Firm incentives to promote technological change in pollution control. *Journal of Environmental economics and Management*, 17(3):247–265.

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.

NAS (2016). *Appendix C: The Role of Research, Learning, and Technology Deployment in Clean Energy Innovation*. The National Academies Press.

Newell, R. G., Jaffe, A. B., and Stavins, R. N. (1999). The induced innovation hypothesis and energy-saving technological change. *The Quarterly Journal of Economics*, 114(3):941–975.

Noailly, J. and Smeets, R. (2015). Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data. *Journal of Environmental Economics and Management*, 72:15–37.

Nordhaus, W. D. (2014). The perils of the learning model for modeling endogenous technological change. *Energy Journal*, 35(1):1–13.

Nowling, U. (2015). Understanding coal power plant heat rate and efficiency. Available at https://www.powermag.com/understanding-coal-power-plant-heat-rate-and-efficiency/.

Papke, L. E. and Wooldridge, J. M. (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics*, 11(6):619–632.

Popp, D. (2002). Induced innovation and energy prices. *The American Economic Review*, 92(1):160–180.

Preonas, L. (2019). Market power in coal shipping and implications for u . s . climate policy.

Rubin, E. S., Azevedo, I. M., Jaramillo, P., and Yeh, S. (2015). A review of learning rates for electricity supply technologies. *Energy Policy*, 86:198–218.

Schumpeter, J. A. (1939). *Business cycles*, volume 1. McGraw-Hill New York.

Appendix A

US Carbon Pricing and Coal Power Plant Productivity

A.1 Data appendix

In this appendix, I describe important data handling procedures for this project.

Net and gross heat rate variables and estimation

I construct monthly plant-level net heat rates from the EIA data by combining heat input and generation over all generation units and fuels at a plant, regardless of primary fuel. (Therefore, I include the natural gas that coal plants use as starter fuel and any other fuels they burn in my heat rate calculations.)[1] Linn et al. (2014) discuss the propensity for reporting issues in the EIA data that lead to heat rates that are outside the bounds of mechanical possibility and develop a rule for eliminating outliers. They remove monthly observations associated with plant-level heat rates that are outside the bounds of 2 standard deviations of a plant's mean across all years. I apply their outlier rule to remove observations that are inconsistent with plants' historical values. I also restrict my panel to observations with heat

[1] Through personal communication with EIA staff, I was advised that this is most reliable way to calculate heat rates from the agency's data.

rates less than 16 mmBtu/MWh and greater than 6 mmBtu/MWh. (My panel is shorter than that of Linn et al. (2014), and the outlier I chose less effective in my case.) I chose these bounds based on summary statistics reported from a variety of academic and industry sources that document historically "normal" heat rates for coal-fired power plants. This eliminates about 8% of my observations from my coal plant panel. I apply the same rule to my unit-level panel, and this eliminates about 1% of observations.

To construct the gross heat rate variable in my unit, monthly panel, I sum all heat input and generation reported by unit by month in the EPA CEMS data. I restrict my panel to units that have coal listed as their primary fuel in the EPA CEMS data, so that I do not observe units after they are converted to another fuel type.[2] For my gross heat rate model, I exclude observations with adjusted capacity factors below 5. These observations indicate that units were hardly running for the month and are associated with heat input values that are very low or very high, meaning that they are unreliable in terms of calculating a monthly heat rate and bias estimation.

Auxiliary power consumption

Construction of auxiliary power consumption requires that net and gross generation data, which are separately collected by the EIA and EPA, be comparable between data sets. The EPA CEMS data does not include fossil-fired units smaller than 25 MW in size, so I cannot compute auxiliary power consumption for coal plants that have small units on-site. Additionally, I cannot compute auxiliary power consumption for plants at which the number of units differs in the EIA and EPA data. Therefore, for my net heat rate decomposition and auxiliary power consumption estimation and decomposition (which relies on the ratio of net to gross generation at plants), I exclude coal plants that have a different number of units reported in each data set. I lose 60 out of 390 coal plants from this. Conditional on this subsetting there is a high correlation (Pearson correlation coefficient = 0.99) between data sets in the monthly

[2]Because unit fuel type is an annual variable, this means that my variable for being within six months of effective "retirement" includes the six months in the latter half of the year before units started the conversion process to a new fuel, if they are converted to new fuel.

fuel consumed by plant, indicating that the monthly, plant-specific generation values can be compared. After this subsetting I observe that 1.9% of observations that have auxiliary power consumption values of less than 0 (which is mechanically impossible) or greater than 20% of gross generation (which is large, based on (ABB, 2009)). I remove these observations for my auxiliary power consumption estimation. Results are robust to including observations with auxiliary power consumption values that are greater than 20% of gross generation.

A.2 Model appendix

In this section, I show that the results of Section 1.3 are robust to modelling a coal plant's efficiency as a function of its output. This is important to consider, as coal plant heat rates and output are mechanically related and may be simultaneously determined.[3]

To model coal plant efficiency as a function of output, I include a new term, u_t, the utilization rate of the plant, in the efficiency equation:[4]

$$e_t = g(\omega_t, \gamma_t, u_t) \tag{A.1}$$

The plant's fuel use is also a function of its utilization rate:

$$f_t = \zeta(u_t) \tag{A.2}$$

Now the output choice is given by u_t instead of f_t (with similar lower and upper bound

[3]As described above in Section 1.2, coal plant owners choose their plants' heat rate in the short term by maintaining their plants and modifying plant operations. Their heat rates are also a function of their output. And, output is determined in part based on plants' heat rates. Practically, coal plants in RGGI states submit individual, detailed supply cost information to their respective wholesale electricity market independent system operator (ISO). The ISO determines each plant's production by meeting a number of objectives, such as minimizing the cost of production from all plants in the system as well as ensuring reliability of service. Therefore, for a given hour of operation, a plant's output is both based on its baseline heat rate and then determined by an an external body. Then, plant managers can adjust plant heat rates real-time.

[4]The output of a plant can be written as a product of its efficiency and fuel use (as in the main text), $q_t = e_t f_t$, or it can be written as a product of the utilization rate of the plant (typically measured in MW) and the number of hours the plant is running, h_t: $q_t = u_t h_t$. Each yields output in MWh of electricity. It is the utilization rate, the output per hour, to which a plant's efficiency is sensitive.

constraints), and the first FOC becomes:

$$\frac{\partial \pi}{\partial u_t} = p_t^w \left[\frac{\partial e_t(\cdot)}{\partial u_t} f_t + \frac{\partial f_t}{\partial u_t} e_t(\cdot) \right] - p_t^f \frac{\partial f_t}{\partial u_t} - m p_t^c \frac{\partial f_t}{\partial u_t} \tag{A.3}$$

This leads to a new version of Equation (1.9):

$$\frac{e_t(\cdot)}{f_t} + f_t p_t^w \frac{\partial e_t(\cdot)}{\partial u_t} \frac{\frac{\partial e_t(\cdot)}{\partial \gamma_t}}{\frac{\partial f_t}{\partial u_t}} = (p^f + m p_t^c) \frac{\partial e(\cdot)}{\partial \gamma_t} \tag{A.4}$$

Compared to Equation (1.9), Equation (A.4) has a new term on the left hand side, $f_t p_t^w \frac{\partial e_t(\cdot)}{\partial u_t} \frac{\frac{\partial e_t(\cdot)}{\partial \gamma_t}}{\frac{\partial f_t}{\partial u_t}}$. This term indicates that, if a coal plant owner reduces output in response to a carbon price, the impact of the output adjustment on increasing the value of the left hand side of Equation (A.4) will be mitigated by the efficiency impact of that output adjustment, $\frac{\partial e_t(\cdot)}{\partial u_t}$. (Though the change in efficiency with respect to utilization is nonlinear as coal plants have an optimal utilization rate, US coal plants have been operating below this optimal rate in recent history. Therefore, $\frac{\partial e_t(\cdot)}{\partial u_t}$ is positive.) The importance of this effect depends on the magnitude of the wholesale electricity price, p_t^w. It also depends on the sensitivity of the plant's efficiency with respect to within-period efficiency expenditures, $\frac{\partial e_t(\cdot)}{\partial \gamma_t}$, and the change in fuel use with the utilization rate, $\frac{\partial f_t}{\partial u_t}$, though these terms are relatively small, positive terms. Equation (A.4) shows that the incentives for responding to a carbon price in terms of efficiency and output adjustments are largely the same when one includes the fact that coal plant efficiency is a function of plant output. This mechanical relationship must be taken into account by coal plant owners. Large efficiency loss means that owners will reduce output even more in the face of a carbon price or work to combat the loss with efficiency-improvement measures.

The between-period efficiency investment choice (given by Equation (1.10)) does not change when e_t is expressed as a function of u_t.

Appendix B

What do we lose by picking winners?

For fuel data, I use the "Fuel and waste costs" and "O&M costs" by technology in Tables 3.11-3.16 of IEA et al. (2015). I construct capacity-weighted averages across all EU countries reporting by technology, using the capacity values reported in Tables 3.9 and 3.10. For wind data, I use onshore wind values only. I remove one reported hydro cost for Spain that has a missing O&M cost value.