Essays in Industrial Organization and Environmental Economics

by

Gabor Palinko

Department of Economics
Duke University

Date: ____________________
Approved:

____
James W. Roberts, Advisor

____
Allan Collard-Wexler

____
Yi (Daniel) Xu

____
Christopher D. Timmins

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2021

Abstract

Essays in Industrial Organization and Environmental Economics

by

Gabor Palinko

Department of Economics
Duke University

Date: ____________________
Approved:

____
James W. Roberts, Advisor

____
Allan Collard-Wexler

____
Yi (Daniel) Xu

____
Christopher D. Timmins

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2021

Abstract

This dissertation is comprised of three chapters in industrial organization, environmental economics and energy economics. In Chapter 2, I study carbon dioxide emission abatement technology for industries participating in the world's biggest emissions market, the European Union Emission Trading System (EU ETS). I propose a production and abatement model to motivate the use of emissions as an input in a production function. I build on recent methods of the production function literature and propose an estimator for the production function that is consistent with my model. Using data from the EU ETS and Orbis, I estimate the elasticity of emissions to abatement expenditures for different manufacturing industries. Increasing the share of abatement expenditures of revenues by 1% is expected to reduce emissions by 8% in cement and 67% in chemicals, with other industries between these two extremes. I use the model's implications to translate estimated abatement elasticities to marginal abatement costs at the individual firm level. My findings show enormous differences both within and across industries. My estimates for the 25th, 50th and 75th percentile cement firms are 15, 22 and 36 €/t respectively. In contrast, these estimates are 22, 48 and 363 €/t for oil refineries. My findings suggests that, cement, chemicals and power firms are the most likely to decrease emissions as the EU ETS market price rises to levels close to the social cost of carbon.

In Chapter 3, I analyze the impact of different policy instruments on the speed of transition to cleaner electricity generation. I develop a non-stationary fully dynamic entry and exit model of power generation. My model includes multiple technologies and hourly spot markets, both key features of the power generation market. I use

the calibrated model to analyze the speed of transition away from coal power plants in PJM, the biggest electricity system in the United States. Correcting the negative externality of carbon dioxide emissions requires environmental regulation. My findings highlight the importance of analyzing the full transition path when comparing environmental policy instruments. Policies that lead to similar long-term outcomes induce vastly different transition dynamics. A carbon tax (the efficient instrument) set to $\$30/tCO_2$ is associated with an almost immediate entry of the long-run gas capacity. In contrast, gas entry and coal exit result in a slower and smooth transition. Welfare differences are significant. Both of these instruments improve only marginally on the baseline scenario and do not come close to the improvement possible by the carbon tax.

In Chapter 4, I study bidding behavior in the New England frequency regulation market. Since 2015, this product is procured through a multi-dimensional Vickrey-Clarke-Groves (VCG) auction. Bidding under a VCG design is simple since truthful bidding is optimal. However, I find that participants bid higher when relative market power increases. This is indirect evidence against truthful bidding. Taking VCG bids as estimates for true marginal cost can be misleading. A combination of a complicated clearing mechanism and low stakes might prevent players to learn the optimal bidding strategy. My results suggest that switching from a uniform price to a VCG auction does not resolve the underlying strategic complexity.

Acknowledgements

I would like to express my gratitude to Allan Collard-Wexler and James Roberts for their guidance, mentorship and support throughout the PhD program. I am indebted to Christopher Timmins for guiding my journey in environmental and energy economics. Daniel Xu always had a paper in mind that provided the solution to the issues I was struggling with. I could not have asked for a better doctoral committee and this dissertation could not have been possible without their support.

The Duke Economics Public-IO Lunch group provided a great environment to present and discuss the work that eventually became this dissertation. I am also grateful for being part of the Duke Energy Initiative. The conversations and resources helped me better understand research in environmental and energy economics. I appreciate the useful feedback from Brian Murray, Billy Pizer and Steve Sexton on several chapters of this dissertation.

Contents

List of Figures

List of Tables

1 Introduction

Now it is well established that human activity contributed significantly to global warming and the changing climate of our planet. Industrial activity such as cement production and power generation emit greenhouse gases to the atmosphere. Greenhouse gases help absorb more of the energy of the Sun. The resulting warming fundamentally changes the climate. Slowing down climate change is likely the biggest challenge for humanity in the 21st century. Any solution includes drastically decreasing the amount of carbon dioxide emissions, the most important greenhouse gas. Since the costs are expected to be enormous, it is essential that we find the most efficient ways to achieve this objective. My dissertation contributes to our understanding of which industries can decrease emissions the most cost-effectively. Economic theory suggests that a carbon tax or a market for emission allowances can be used reach the most efficient outcome. However, introducing these policy tools might not be politically feasible in many countries around the world including the United States. My dissertation provides guidance on what other policy tools can work in this case.

Marginal abatement costs measure the cost of reducing carbon dioxide emissions. As such, they can be used to compare the cost-effectiveness of decreasing emissions through various alternatives. Marginal abatement costs also provide useful guidance for environmental policy, especially when a carbon tax is not politically feasible. In this case, subsidizing abatement in industries with the lowest cost can be welfare improving. However, our understanding of marginal abatement costs is limited in many of the most emission heavy industries. In Chapter 2, I estimate marginal abatement costs for the most emission heavy industries in the European Union. These include cement production, oil refining, power generation and steel manufacturing which combined are responsible for around 40% of total emissions. Since technology is likely similar for the same industry in different locations, my results are likely relevant for a wide range of countries.

The transition to clean electricity generation is a prerequisite for a future with low levels of carbon dioxide emissions. Replacing coal fired power plants with newer, cleaner technologies is likely the most important lever at our disposal. However, it is not clear how this transition should take place and how policy can accelerate it. Economic theory suggests the market inefficiency due to the negative emission externality can be corrected with a well calibrated carbon tax. In Chapter 3, I ask how close alternative policy instruments such as entry and exit subsidies can get to the efficient outcome. Again, this question is relevant, since carbon taxes are often not politically feasible. I argue that because the transition is slow, it is important to analyze the entire transition path. In order to do so, I develop a fully dynamic non-stationary model of power generation. I calibrate the model using data from PJM, the biggest electricity system in the United States.

Finally, in Chapter 4, I study auctions to procure frequency regulation for the New England power generation market. Frequency regulation helps keep supply and demand balance of the power system in a minute to minute basis. My interest in this product is twofold. First, with the transition to clean energy, the product is becoming more and more important. A system with high levels of solar and wind generation requires more balancing. Second, studying the procurement mechanism is interesting in its own right. As a response to a regulatory intervention, the market design changed from a single dimensional uniform price auction to a multidimensional

Vickrey-Clarke-Groves (VCG) auction. The setup is unique as there is virtually no observational evidence on how the VCG mechanism performs for multiunit auctions. The topic is relevant today as both Facebook and Google uses VCG auctions to sell online advertisements.

2

Estimating Marginal Abatement Costs for Carbon Dioxide Emissions in the EU ETS

2.1 Introduction

Now it is well established, that human activity, and in particular, carbon dioxide emissions contributed significantly to global warming and the changing climate of our planet. Tackling these issues will most likely include drastically decreasing the amount of carbon dioxide emissions. The costs are likely to be enormous; therefore, it is essential that we find the most efficient ways of achieving this objective. Economic theory suggests comparing alternatives by their marginal cost of abatement, the cost of reducing an additional unit of emissions. Researchers made significant progress in understanding the cost-effectiveness of a wide range of available options; however, there is surprisingly little observational evidence on how costly it was to reduce emissions in manufacturing industries in the past. For example, Gillingham and Stock (2018), a recent review of marginal cost of abatement estimates, only cites engineering estimates for manufacturing industries. This is particularly striking, since these industries can be responsible for around 20-30% of carbon dioxide emissions in developed countries[1].

The goal of this paper is to estimate the cost of reducing carbon dioxide emissions in manufacturing industries in the European Union. I do this in three main steps. First, I propose a production and abatement model which provides a foundation for using emissions as an input in a production function. Under the assumptions of the model, there is a one-to-one mapping between the output elasticity of emissions and the elasticity of emissions to abatement expenditures. Second, building on recent methods in the production function literature, I estimate these elasticities for different industries. The average output elasticity of emissions across all industries is around 0.07. However, these estimates mask significant heterogeneity across industries with the highest estimates in cement (0.127) and the lowest in chemicals (0.015). Alternatively, if the average firm spends 1% higher share of revenues on abatement, emissions are expected to fall by 14%. Finally, I use the model's implications to translate estimated abatement elasticities to marginal abatement costs at the individual firm level. I find that marginal abatement costs are substantially above market prices for emission allowances. A combination of asymmetric adjustment costs, long-run expectations and the dynamic nature of abatement decisions can explain this result. Differences both within and across industries are substantial. My estimate for the 25th, 50th and 75th percentile cement firms are 15, 22 and 36 €/t respectively. In contrast, these estimates are 48, 363 and 1053 €/t for oil refineries.

I make three specific contributions in this paper. First, I estimate carbon dioxide marginal abatement costs for various manufacturing industries using observational data for the first time. Understanding these costs in the past in the European Union is likely to be informative about expected future costs in other parts of the world. My results complement existing observational evidence in other environments, such

[1] For instance, the Environmental Protection Agency calculates a 22% share of greenhouse gas emissions of industrial activity in the United States in 2018. Eurostat estimates a 29% share for manufacturing industries in the European Union in 2018. Neither of these include emissions from the power sector.

as replacing light bulbs, switching to renewable electricity generation or moving toward a car fleet powered by electricity[2]. In addition, I estimate the CO_2 abatement elasticity for a wide-range of manufacturing industries. This elasticity plays a key role in environmental economics models. For instance, in the model of Copeland and Taylor (2003), it determines the trade-off a firm's faces between production and abatement. Shapiro and Walker (2018) estimates similar elasticities for a variety of pollution outcomes but not for carbon dioxide in US manufacturing industries. They observe abatement expenditures and use exogenous variation due to regulation to identify abatement elasticities. My results complement their findings and provide a new methodology to estimate the same object.

Second, the empirical strategy I propose is new. I start from the production and abatement model of Copeland and Taylor (2003). I show that under certain assumptions[3] for the abatement technology, it is possible to represent both production and abatement technologies in a single, transformed production function with emissions as an input. An implication of the model is that estimating the output elasticity to emissions as an input is equivalent to estimating the elasticity of emissions to abatement expenditures. Following recent approaches in the literature to control for unobserved productivity (Ackerberg, Caves and Frazer (2015)), I propose an empirical strategy to estimate the production function that is consistent with my setting. I extend the production and abatement model by adding timing and information assumptions on the firms' dynamic optimization problem. I show that a modified intermediate input demand equation can be used to control for unobserved productivity. The key advantage of my methodology is that it can be used with widely available data on production variables and emissions. In particular, I do not have to observe CO_2 specific abatement expenditures[4]. As such, my method can be applied to estimate

[2]For a review of the literature and alternatives, see Gillingham and Stock (2018).

[3]The key assumptions is that the abatement technology to decrease the emission intensity of production inhibits constant elasticity. See Section 2.3 for the detailed assumptions.

[4]Even when this data is available, an instrument is needed to deal with the endogeneity of

abatement elasticities for other types of pollution or for CO_2 in different locations or industries. The key drawback is the relatively strong functional assumptions on the abatement technology.

Finally, I contribute to the discussion of market design in the European Union Emission Trading System (EU ETS). The EU ETS is the biggest cap and trade market for carbon dioxide emissions in the world. Total emission allowances are set and allocated to firms who are free to trade to cover their emissions. The resulting price helps align the private and social costs of pollution. My findings can help predict abatement at different levels of market prices. Cement, chemicals and power firms are the most likely to decrease emissions as the carbon price rises to levels close to the social cost of carbon. My results suggest that the EU ETS fails to equate marginal abatement costs across participating firms.

In the next section, I describe the setup, introduce my data and discuss key characteristics of the industries in my data. Section 2.3 presents my production and abatement model for carbon emitting manufacturing industries. In Section 2.4, I propose an empirical strategy to consistently estimate abatement elasticities and marginal abatement costs. Section 2.5 presents the estimates. In Section 2.6, I discuss the implications of my results for the EU ETS and environmental policy. Section 2.7 concludes.

2.2 Background, Data and Descriptive Evidence

2.2.1 The Role of the EU ETS in Limiting Greenhouse Gas Emissions

Greenhouse gas (GHG) emissions is one of the primary examples of how industrial activity can result in unwanted consequences for our society and our planet. It is well established, that greenhouse gas emissions contributes significantly to climate change and global warming. Without regulation, firms are unlikely to internalize the true

abatement. Additionally, measurement issues arise as it is hard separate actions directly taken to decrease carbon dioxide emissions making measurement difficult. For instance, energy efficiency improvements both abate emissions and decrease input costs.

cost of emissions on society. One possible solution is to introduce a market for this negative externality that helps align private incentives with what is socially optimal. The European Union Emissions Trading System (EU ETS) is the world's biggest such market for greenhouse gas emissions. The most important GHG in terms of its contribution to global warming is carbon dioxide and it is so dominant, that the unit of regulation is carbon dioxide equivalent. As a consequence, the EU ETS is usually referred to as a carbon (dioxide) market. It is organized as a cap and trade market. Total emission allowances are set and allocated to firms who are free to trade to cover their emissions. The resulting price helps align the private and social costs of pollution.

The EU ETS covers around 50% of total EU carbon dioxide emissions and almost all emissions in the highest emitting manufacturing industries and the power sector[5]. The biggest contributor is the power sector responsible for around 60% of emissions, while cement, oil refining and steel has a share of around 10%. Chemicals, ceramics, glass, non-ferrous metal and paper production combined is responsible for the remaining 10%. It is not clear how much incentives the program provides to decrease emissions. In Phase I (2005-2007) and Phase II (2008-2012) the majority of firms received allowance allocations that covered most of their emissions for free. After Phase III (2013-2020), allocations are determined based on historical production and product benchmarks. The average price in 2005-2016 was 11.6 €/ton CO_2 and prices generally remained below 20 €/ton CO_2 (see Figure 2.1)[6]. Predetermined quantity caps and low output levels due to the Great Recession are the two most likely reasons to lower than expected prices.

[5]Remaining sources of emissions include transportation and residential energy use. Small power plants are automatically excluded. Small plants in other industries can acquire an exemption but it is rare in practice.

[6]As a comparison, the current estimate of the Biden administration for the social cost of carbon is \$51/ton CO_2

Source: Point Carbon (Thompson Reuters).
Figure 2.1: EU ETS Monthly Price Averages (2005-2018)

2.2.2 Data

For my empirical analysis, I comply data on the emissions, allowance trading and financial statements of firms regulated by the EU ETS. I use data from 2005 to 2016 covering Phase I-III of the EU ETS. The EU ETS registry administers data on regulated entities and tracks allowances used for complying with the program. I observe the identifier, company registration number, location and main activity of every installation (plant). Emissions and allowance allocation are reported at the plant level. In fact, the existence of the EU ETS is the main reason why emissions are measured. The EU ETS registry also records the exchanges of allowances at the firm level. For each transaction, I observe the quantities exchanged, the time and the parties involved[7]. Financial statements of firms come from the Orbis dataset by Bureau van Dijk. The coverage of Orbis in EU countries is substantial. For instance, Kalemli-Ozcan et al. (2015) is able to cover around 70% of total economic activity in the EU and reproduce firm distributions similar to produced by Eurostat. The data allows me to measure production variables in monetary values at the firm level. I observe revenues, fixed assets, the number of employees and materials costs

[7]Price is not available. This is not the universe of allowance trading. Financial transactions can take place outside of the registry (Commission et al. (2015).)

as measures of output, capital, labor and intermediate inputs respectively. I link the EU ETS and Orbis data using country level firm identifier codes. I am able to match around 80% of EU ETS accounts to firms in the Orbis dataset[8]. Appendix A.1 provides more details on data construction.

A couple of notes about data quality and selection are in order. First, emissions data is missing for around 30% of EU ETS accounts and I drop these observations from my analysis. Second, although Orbis covers the range of firms well, there is significant attrition as firms get in and out of sample even when they keep operating[9]. I lose an additional 20% of firm-year observations because of attrition. Third, key production function variables (especially materials costs) are available for only around 70% of matched firm-year observations. My final sample that I can use for production function estimation has 31782 firm-year observations over the 2005 to 2016 time period. The final sample covers around 30% of the total emissions in the EU ETS program in any given year. Firms in the final sample tend to be bigger by 20-30% when measured by revenues or capital. Appendix A.1.4 provides more details on data quality and selection.

The key benefit of the Orbis data is that it allows me to analyze the behavior of EU ETS entities all around Europe with a relatively large sample size. Using Orbis makes the unit of observation the firm. This is reasonable for analyzing allowance trading decisions that are executed at the firm level. However, the relevant unit for production technology is the plant. Around 75% of firms in my data have only one plant under the regulation of the EU ETS. Although, this might suggest that firm level measures are representative of the plant this is unfortunately not the case. Firms might own small plants in EU ETS industries or have activities outside the EU ETS.

---

[8]I find the quality of the merge to be comparable to other papers in the literature using similar data (e.g. Marin, Marino and Pellegrin (2018))

[9]As such, Orbis is not well-suited for entry and exit analysis. In addition, a balanced sample of Orbis for ten year time period is very restrictive. I find that I lose around 70% of observations when using a balanced sample.

Vertically integrated firms are common in several industries in my data. I discuss the implications of these issues on my estimates in Section 2.5.3 and perform several robustness checks to overcome the data limitation. The alternative to Orbis is to use manufacturing census data at the plant level. However, this type of data is available for researchers in only a few larger EU countries[10]. Additionally, quantity based output measures are usually available instead of purely monetary measures that help overcome unobserved price limitations in productivity estimation.

2.2.3 Industry Classification

I use the main activity of a plant as reported by the EU ETS to categorize firms. These are the same categories that determine whether an installation is included in the program. I group firms to ten industries based on their main activity according to the EU ETS. The industries are cement, ceramics, chemicals, glass, non-ferrous metals, paper, power, refineries, steel, and other. The EU ETS lists activities in a more refined way but I group them to ease estimation. The final groups are relatively homogenous[11]. Roughly 90% of firms have only one type of plant and I assign the modal activity type for multiplant firms. The exercise is relatively straightforward with two complications. First, multiple stages of the production of a final product can be carbon emitting. In this case, I group together the entire chain even though some parts might be more similar to other products. I do this because firms are often vertically integrated and I want to be able to keep close competitors in the same activity category[12]. Second, the data does not differentiate well independent power plants from small on-site power generators. I keep only power generators that belong

[10]The French manufacturing census is relatively easily available for EU citizens during times other than a global pandemic. The German census must be accessed through data centers located in Germany. Other manufacturing censuses at the plant level are not widely available (at least to my knowledge). The Spanish survey ESEE widely used in the production function literature (e.g. Doraszelski and Jaumandreu (2018)) is at the firm level.

[11]See a detailed description of definitions in Appendix A.2.

[12]For instance, a firm might own a steel and an iron coke plant. The firm is essentially a steel manufacturer even though it might have more iron coke plants. Iron coke production might be closer by technology to oil refining but it would be a mistake to group the firm together with oil refineries.

to firms classified as primarily electricity producers in category power[13]. I include the emissions of on-site generators in industries regulated by the EU ETS in the final emissions of the firm. I assign power generators of firms in unregulated industries to category other[14].

I prefer using EU ETS activity types to standard industry classifications for several reasons. First, these are plant level whereas I have access to firm level NACE and NAICS codes. Second, the EU ETS categories help focus attention on industries from an emissions viewpoint. Therefore, they are likely to represent emission technologies that are more homogenous. Finally, they are likely to represent vertically integrated industries better. Using categories based on standard industry classification would likely lead to similar outcomes. Most activity types have a primary 3-digit NACE code that is responsible for around or above 70% of firms (see Appendix A.2 for details). Vertical integration usually explains the remaining codes. For instance, the second and third most common for cement are limestone mining and ready-to-mix concrete production.

2.2.4 Sources of Emissions and Abatement Options

The industries under regulation of the EU ETS are capital heavy and use physical and chemical processes to transform raw materials to final products. The chemical process of cement, ceramics, chemicals, glass, paper, and steel production requires very high temperatures to create compounds with more desirable properties. For instance, glass production requires mixing silicates, soda-lime and additives at temperatures above 1500 Celsius. The required heat is generated by burning coal which is the main source of carbon dioxide emissions. In some cases, most notably cement, the

[13]NACE code 3511. I am trying to be conservative and even exclude companies whose primary codes are electricity distribution and transmission.

[14]Any power generator above 20MW capacity is regulated by the EU ETS. In most cases (above 2000 observations in a year), on-site generators belong to firms in industries that are otherwise not regulated by the EU ETS. Industries range from food processing facilities to universities. When a firm has plants with multiple EU ETS activity types including power, I assign the firm to another activity type if observed even the mode is power.

chemical process itself results in carbon dioxide as a byproduct. Oil refineries heat crude oil and separate hydrocarbons with different boiling points. The required heat comes from burning fossil fuels (oil or coal in this case). The most common process of manufacturing non-ferrous metals is electrolysis which uses electricity to separate the desired metal from its ore. Many times, this comes from on-site power generation emitting CO_2. Finally, traditional power plants burn fossil fuels to produce electricity emitting carbon dioxide. As a summary, CO_2 emissions come from three main sources: heat generation, chemical reactions and power generation[15].

Abatement options inside a plant vary across industries but generally can be divided into three distinct categories: energy efficiency improvements, alternative inputs and fuel switching. First, energy efficiency improvements decrease carbon emissions by decreasing the amount of energy required for producing the same output. For most industries, there are many different levers along the production process all resulting in small improvement of carbon efficiency. For instance, the Environmental Protection Agency lists above 70 different ways to improve energy efficiency and decrease emissions for integrated steel plants (Jones (2012)). Most of these require upfront investment to boost energy efficiency permanently. When plants generate power onsite, decreasing the power use of office buildings or switching to solar panels decreases emissions through this channel as well. Capacity utilization and external factors such as weather might also influence emission efficiency. Second, the use of alternative inputs can decrease emission in several industries. For instance, substituting clinker with alternatives (coal ash, slag) in cement production can reduce carbon emissions by a third (Mahasenan, Dahowski and Davidson (2005)). Third, switching to less carbon intensive fuels is available in some industries. Using gas instead of coal to generate heat is possible in industries when the required heat is not too high. Plants might also decrease onsite power generation and increase electricity inputs to

[15] Appendix A.2 provides a more detailed review of the manufacturing processes and sources of carbon emissions.

cut direct carbon emissions. Coal power plants and paper producers can also switch to biomass. Finally, end of pipe technologies (filters) are not available for carbon dioxide. Although carbon capture is a rising technology it is far from being competitive at current carbon prices. In addition to plant level improvements, firm level emission can be decreased by reshuffling production to more emission efficient plants. Examples include substituting towards gas in power, dry process kilns in cement or towards newer plants in any industry. Similar channels might exists within plants with multiple production lines as well[16].

The few papers that are able to study abatement channels suggests that energy efficiency improvements are the most important channel within a plant (Petrick and Wagner (2014), Colmer et al. (2020)[17]). Surveys tend to find that decreasing emissions are considered only a secondary benefit of energy efficiency improvements (Martin, Muûls and Wagner (2016)). There is some evidence for fuel switching in manufacturing in Phase I (Petrick and Wagner (2014)), composition effects across plants in cement (Branger et al. (2015) and fuel switching across plants in power (Ellerman and McGuinness (2008)). Finally, it is likely that at least some of the gains during the low-demand environment of the financial crisis were due to composition effects. To summarize, it is not entirely clear what abatement mechanisms are the most important and how firms make abatement decisions.

[16] Cement plants usually have multiple kilns likely with different energy efficiency. Power plants generally have several generators at the same site often using different technologies.

[17] Petrick and Wagner (2014) uses German microdata and finds that firms they reduced fuel intake in Phase II while not increasing power consumption from the grid. They interpret Phase II findings as evidence for improvements in the use of process heat. They supplement this with interview evidence suggesting that improving capital stock to improve energy efficiency is the most important channel for firms. Colmer et al. (2020) studies abatement using French manufacturing microdata. They provide evidence that plants did cut natural gas fuel consumption but did not increase electricity consumption from the grid. They also see increases in capital investments likely toward more energy efficient capital. They interpret these findings as evidence for capital investment for energy efficiency being the key channel of abatement.

2.2.5 Descriptive Statistics

Table 2.1 presents descriptive statistics by industry in 2010 for a restricted sample of firms with all variables available for production functions estimation. Category other is responsible for around 50% of the total 4743 firm observations in 2010. The number of observations by industry is in the range of 50 (non-ferrous metals) to 741 (ceramics). Cement, power, refineries and steel are responsible for only 16% of total observations but around 75% of emissions. Panel A reports medians of production variables[18]. The median firm in my data is large and capital intensive with revenues and fixed assets of €65 and €34 million respectively. The median firm in every industry owns a single plant under EU ETS regulation[19]. This masks differences across industry in the share of multi-plant firms that is between 15% (paper) and 41% (cement). The high share of single-plant firms suggests that firm level data might be informative of abatement possibilities at the plant level. The Great Recession hit some industries harder than others. Cement, ceramics, non-ferrous metals and steel were among the hardest hit, revenues between 2007 and 2010 declined by 12, 34 and 20% respectively. These industries were likely operating at lower levels of capacity during and shortly after the Great Recession. In contrast, others either were hardly hit at all or even improved (power, refineries). The high share of multi-plant firms and a big hit in demand suggests that production reshuffling might be an important abatement channel at the firm level in cement or steel but less so in paper.

Panel B of Table 2.1 presents emissions statistics by industry. The median plant would spend from a few hundred thousand to a few million euros on allowances if it

[18] I prefer medians as most industries are characterized by a long right tale of firms that would skew averages.

[19] Note that this does not exclude ownership relations across firms. Multi-nationals are likely relatively common in my data. These firms could coordinate emissions choices and trading. The EU ETS literature usually uses firms at a country level as a unit for analysis. Survey evidence suggests that even when this is the case, firms make emissions and trading decisions locally (Martin, Muûls and Wagner (2014)).

were to pay market prices[20]. This corresponds to less than one percent of revenues. However, these numbers mask vast heterogeneity both within and across industries. Emission expenditures of the median cement plant would be €2.9m or 9.5% of total revenues whereas they would be negligible for non-ferrous metal plants. Cement both requires higher emissions per unit of output and has a much lower market price. Even in industries with low revenue shares, emission expenditures would account for at least a few percentage of profits[21]. Additionally, emission costs are substantial for a high share of firms in all industries. For instance, the 90th percentile of refineries has a cost share of emissions around twenty times of the median. These numbers suggest that emission expenditures are likely important for some but not all firms.

Aggregate descriptive statistics in Table 2.1 Panel C suggest substantial differences across industries in allowance allocations and trading behavior. Output decline due the Great Recession lead to a large surplus of allowance allocations in all industries other than power and refineries. This is consistent with the intentions of the EU ETS to make industries with high emitters short to enhance trading. As a consequence, power and refineries firms are more likely to participate in trading allowances. Overall, only 37.5% of the restricted sample of firms trades in 2010.

[20] The average market price of 2010 is €14.4/ton CO_2 which is slightly above the historical average of 10.7 in 2005-2016. Market prices vary substantially between years, see Figure 2.1

[21] Since operating profits can be negative, the median emission costs over profits are close to zero and below the cost share of revenues close for several industries.

Table 2.1: Descriptive Statistics by Industry in 2010

	Cem	Cer	Chem	Glass	Non-ferr	Paper	Power	Ref	Steel	Other	All
Panel A: Production											
Revenue (mEUR)	43.8	4.7	252.7	53.6	347.2	51.9	35.4	454.2	208.6	53.7	64.6
Fixed assets (mEUR)	34.3	5.5	96.2	34.2	64.1	19.0	45.5	348.4	66.1	28.7	33.5
Number of employees	175.0	47.5	425.0	300.0	478.0	166.5	34.5	496.5	594.0	175.0	183.7
Number of EU ETS plants	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Revenues (2007=100%)	87.3	63.6	95.2	90.0	84.0	102.6	125.3	107.6	77.9	107.5	97.8
Panel B: Emissions											
Emissions (th ton)	204.7	0.8	43.5	23.2	0.0	10.6	35.7	612.2	30.2	8.5	26.6
Emission cost (mEUR)	2.9	0.0	0.6	0.3	0.0	0.2	0.5	8.8	0.4	0.1	0.4
90 pct em. cost (mEUR)	25.2	0.3	6.9	2.0	0.6	1.5	30.3	60.5	3.9	1.3	5.0
Em. cost/revenues (%)	9.5	0.6	0.2	0.8	0.0	0.5	2.1	0.4	0.3	0.2	0.8
90 pct em. cost/revenues (%)	16.1	5.7	3.1	2.3	0.2	1.5	10.7	8.4	1.9	5.1	5.3
Em. cost/profits (%)	50.0	0.0	2.0	5.0	0.0	3.0	7.0	3.0	0.0	1.0	3.4
Panel C: Activity totals											
Multi-plant (%)	41.4	16.3	41.0	20.5	16.0	14.9	27.0	33.9	29.4	24.3	23.4
Activity emission share (%)	12.6	0.8	2.6	1.6	0.1	2.4	39.3	9.8	12.3	18.5	13.5
Allocations/emissions	1.4	2.1	1.3	1.3	1.6	1.3	0.9	1.0	1.4	1.0	1.3
Trader (%)	61.5	40.8	32.0	27.2	12.0	35.3	58.4	54.2	27.6	34.6	37.5
Observations	174	741	122	224	50	484	322	59	221	2345	4742

Notes. Medians reported if not otherwise indicated. Sample is restricted to observations with production function variables available. Emission cost calculations assumes firms pay market price for allowances. Profit measure is EBIT. Sample includes only firms with all production function variables (revenues, fixed assets, number of employees and emissions) available in 2010.

2.3 Production and Emissions Model

This section describes a structural model of production and emission abatement. The key modeling assumption is that firms need to allocate productive resources to decrease emissions. The model captures the trade-off between productivity and abatement expenditures[22]. In the model, there is a short term relationship between yearly abatement expenditures and revenue emission intensity. As a consequence, abatement and production technologies can be represented in a transformed production function with emissions as an input. The model allows me to measure marginal abatement costs without explicitly observing abatement expenditures.

2.3.1 Production and Abatement Technology

Firm i combines capital (K), labor (L) and intermediate inputs (M) to produce output (Y) in each time period t using production technology $f(\cdot)$. Intermediate inputs include raw materials (iron ore, limestone) and sources of energy to produce heat (coal, natural gas). The abatement assumptions follow Copeland and Taylor (2003) and Shapiro and Walker (2018). Firms dedicate a share a_{it} of output to decrease emissions intensity:

$$Y_{it} = (1 - a_{it})\bar{Y}_{it} = (1 - a_{it})f(K_{it}, L_{it}, M_{it}). \tag{2.1}$$

Equivalently, the firm uses a mix of inputs in the same proportions as the production technology for abatement. Without abatement, the firm would produce potential output $\bar{Y}$. Emissions are a function of potential output and abatement expenditures:

$$E_{it} = g(a_{it})f(K_{it}, L_{it}, M_{it}), \tag{2.2}$$

where $g(\cdot)$ is monotonically decreasing and $g(0) = 1$. Emissions increase linearly with potential output. The multiplicative functional form implies that abatement

[22] This assumption is relatively standard in the literature but not without controversy. It is widely used in the trade-emissions literature (Copeland and Taylor (2003), Shapiro and Walker (2018)) and assumed in several recent work on emissions and environmental externalities (He, Wang and Zhang (2019), Rafey (2020)). The key empirical studies supporting this view are Greenstone (2002) and Greenstone, List and Syverson (2012). Both use US Census of Manufacturing data and find sizable short term productivity declines to environmental regulation. In contrast, the Porter hypothesis (Porter and Van der Linde (1995)) states that environmental regulation can lead to higher productivity on the long run (see Dechezleprêtre and Sato (2017) for an empirical review).

expenditures decrease emission intensity where the denominator is potential output[23]. The functional form and shape of $g(\cdot)$ determines this relationship. When a_{it} equals zero, there is no abatement and output equals potential output. As $g(\cdot)$ is monotonic, it can be inverted to express a_{it}. Combining (2.1) and (2.2) leads to

$$Y_{it} = \left[1 - g^{-1}\left(\frac{E_{it}}{f(K_{it}, L_{it}, M_{it})}\right)\right] f(K_{it}, L_{it}, M_{it}) = \tilde{f}(K_{it}, L_{it}, M_{it}, E_{it}). \tag{2.3}$$

Equation (2.3) represents both the production and abatement technology of an industry in a single function. This representation provides two main advantages. First, it shows clearly that emissions can be treated as an input to production. A firm can increase output by increasing emissions through cutting abatement expenditures. The key assumption is that abatement expenditures only effect emission intensity in the same year. As a result, there is a one-to-one mapping between emissions and abatement expenditures conditional on other inputs and productivity. Second, it is possible to empirically estimate (2.3) with data on emissions only. As I show in Section 2.4 , existing production function estimators can be applied to estimate $\hat{f}$.

The model and empirical strategy is flexible to accommodate a wide range of production and abatement technologies. I choose the Cobb-Douglas production function and the constant elasticity abatement technology:

$$f(K_{it}, L_{it}, M_{it}) = \Omega_{it} K_{it}^{\alpha_K} L_{it}^{\alpha_L} M_{it}^{\alpha_M} \quad \text{and} \quad g(a) = B_{it}(1 - a_{it})^{\frac{1}{\alpha_E}}.$$

The key benefit of this functional form is that the transformed production function remains Cobb-Douglas[24]. Cobb-Douglas provides a first-order approximation to any production function, is widely used in the literature and represents an easy to interpret abatement technology. B_{it} is a firm-time specific emission intensity parameter that represent emission intensity if $a_{it} = 0$. $\frac{1}{\alpha_E}$ is the elasticity of emission intensity to the share of potential output left for production after spending on abatement $(1 - a_{it})$.

[23] An alternative is to model emissions as a result of input use. Shapiro and Walker (2018) finds the two models comparable both qualitatively and quantitatively.

[24] An alternative is to assume $g(a) = C - a^{\gamma_E}$. With appropriate parameter choices, the two represent similar technology. Therefore, writing the abatement function in the original form is not restrictive.

If α_E is high it is relatively difficult to decrease emissions intensity by spending on abatement. Crucially, this does not depend on initial emission intensity (B_{it}). It is as hard to cut emissions (in percentages) for a firm that starts from a low baseline as for a firm that starts from a high emission intensity level. The transformed production function takes the following form:

$$Y_{it} = (\Omega K_{it}^{\alpha_K} L_{it}^{\alpha_L} M_{it}^{\alpha_M})^{(1-\alpha_E)} \left(B_{it}^{-1} E_{it}\right)^{\alpha_E} = \bar{\Omega}_{it} K_{it}^{\bar{\alpha}_K} L_{it}^{\bar{\alpha}_L} M_{it}^{\bar{\alpha}_M} E_{it}^{\alpha_E}$$

where $\bar{\Omega}_{it} = \Omega_{it}^{1-\alpha_E} B_{it}^{-\alpha_E}$ and B_{it}^{-1} can be interpreted as emission augmenting technology. However, in the Cobb-Douglas case all productivity is factor neutral. Notice that $\bar{\Omega}_{it}$ represents both the initial productivity Ω_{it} and potential emission intensity B_{it} differences across firms. Firms with lower emission intensity are more productive if measured by $\bar{\Omega}_{it}$. Eventually, this allows me to separate the impact of exogenous emission intensity improvements (changes in B_{it}) from the impact of abatement expenditures[25]. The final production function takes the Cobb-Douglas form. The elasticity of substitution of all inputs including emissions is unity.

At this point, it is important to interpret key modeling assumptions and discuss potential limitations. First, a_{it} can also be interpreted as abatement stock. My model and results remain similar as long as a unit of abatement expenditures increases abatement stock by a unit. Second, my model can accommodate a wide range of abatement options. For instance, my model is consistent with capital only abatement technology. An a_{it} share of capital is dedicated towards abatement. When a firms invests in capital to decrease emissions, a_{it} increases. In the Cobb-Douglas case, the transformed production function differs only in the interpretation of the parameter estimates. Third, the relationship between emissions and output might not be linear. For instance, emission intensity might naturally vary with capacity utilization. I can allow emissions to increase in output at a faster ($\bar{Y}^2$) or slower ($\sqrt{\bar{Y}}$) rate as long as the relationship takes the form of a power function. Appendix A.3.4 discusses the detailed arguments for all three cases. Altogether, the modeling framework is flexible

[25] Allowing for B_{it} to vary across firms and time would be more difficult with other functional forms. B_{it} becomes emission augmenting technology and the production function is difficult to identify without additional data. See for instance Doraszelski and Jaumandreu (2013) or Raval (2019)

and is able to incorporate different abatement options from capital intensive energy efficiency improvements to static fuel switching.

2.3.2 Timing and Optimization

Identification of the production function requires assumptions on timing and optimization[26]. The assumptions are standard and allow for a wide range of dynamic optimization and timing. The key assumption is that materials t are chosen after observing productivity $\bar{\Omega}_{it}$. The main benefit of the approach I take for understanding abatement is that I do not need to assume that abatement choices are made optimally[27]. In the context of the EU ETS, survey evidence suggest that the majority of firms do not intentionally optimize abatement decisions (Martin, Muûls and Wagner (2014)).

At the beginning of period $t-1$ firms observe productivity $\bar{\Omega}_{it-1}$. Capital K_{it} is a dynamic and predetermined in time $t-1$. Labor and abatement expenditures are chosen at time $t-b$ where $b \in [0,1]$. This allows for L_{it} and a_{it} to be static or predetermined. Then, firms observe $\bar{\Omega}_{it}$ and choose M_{it}. Finally, consistent with the production model and the timing assumptions, emissions E_{it} are determined as a result of all the choices and productivity at the end of the period. This is because E_{it} is a function of abatement expenditures and output; and, therefore also productivity and all other input choices. It is useful to explicitly define the firm's information set at different points in time. The information set at the end of time $t-1$ includes productivity for the next period, predetermined inputs, prices, allowance allocations, and the full $t-2$ information set:

$$I_{i,t-1} = \{\bar{\Omega}_{it}, K_{it}, M_{it-1}, L_{it-1}, E_{it-1}, a_{it-1}, A_{it}, P_{it}, I_{i,t-2}\}.$$

A_{it} denotes emission allowances allocated for firm i for period t and P_{it} is output price. Then, the firm's information set at time $t-b$ with $b \in [0,1]$

[26]This is true of any production function estimator that hopes to control for endogeneity of input choices due to unobserved productivity.

[27]Most previous structural approaches assume some form of optimization subject to a well-defined model. For instance, Toyama (2018) assumes that firms optimize abatement and trading decisions in a dynamic model with transaction costs in the US Acid Rain Program.

$$I_{i,t-b} = \{F_{i,t-1}, \tau_{it}, P_{it}^L\},$$

where τ_{it} and P_{it}^L denote the allowance price and wage respectively. Finally, at time t intermediate input prices P_t^M are observed and M_{it} is decided.

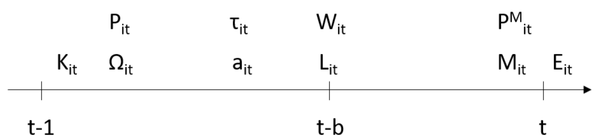

Notes. Placement of variables indicates when they are observed by the firm.
Figure 2.2: One Possible Example for the Timing Assumptions of the Model

The most important timing assumption is that materials are decided at the end of the period. The approach allows slightly different timing assumptions on other inputs but some of these requires small modifications of the estimator[28]. Similarly, when and whether firms observe prices and the allowance allocation is not crucial. I do not need to specify a full dynamic model of optimization for the identification of the production function. The only assumption I make is the static optimality of material input demand. This excludes any dynamic implications including adjustment costs. However, I do not need to assume static optimization of labor and abatement. Transaction costs of trading inputs and adjustment costs are also compatible with the model and estimator. In my context, this is particularly important as the abatement choice is likely not optimal in a static sense. For most firms, emission expenses are a very small share of revenues providing less reason to spend resources to optimize. In addition, transaction costs of trading are likely relatively high creating a wedge between marginal costs and benefits of emissions. Finally, I assume that firms take output and input prices as given. I discuss the consequences of imperfect competition in Section 2.4.3 and Appendix A.3.2.

Under perfect competition and static profit maximization of materials the intermediate input demand takes the following form[29]:

[28] The key is that instruments come from I_{it-1} (or even earlier). Figure 2.2 presents one possibility. If labor is a static input the instrument is L_{it-1}. If it is predetermined L_{it} can also be an instrument for itself.

[29] This directly follows from the first-order condition of static profit maximization. I show this in

$$M_{it} = \mathbb{M}(K_{it}, L_{it}, E_{it}, P_{it}, P_{it}^M, \tau_{it}, \bar{\Omega}_{it}). \tag{2.4}$$

Equation 2.4 is important as it serves the basis of the empirical strategy. The key idea is that conditional on other inputs and prices, materials demand is informative about unobserved productivity $\bar{\Omega}_{it}$. Notice that the prices that enter the conditional intermediate input[30] demand equation are output prices P_{it}, materials prices P_{it}^M and emission prices τ_{it}. Equation 2.4 explicitly shows that if firms differ in output, materials or emission prices, these have to be observed to recover $\bar{\Omega}_{it}$ from materials choices. However, other input prices do not enter the conditional intermediate input demand equation. Capital and labor inputs are decided before materials and only inputs are relevant for output. Therefore, only these inputs and not input prices enter the relevant profit maximization for materials.

The key difference of my model to the literature is that production leads to costly emissions. As a consequence, I have to control for emissions and emissions prices in the intermediate input demand equation. Emissions are a result of production and abatement decisions. As such, emissions are determined after all other inputs are decided. Therefore, the firm has to take into account the contribution of its input decisions to final emissions costs. This is most pronounced for the materials choice, since it eventually decides the scale of production. The firm might find it optimal to decrease the scale of operations when emission prices are high. Notice that only emissions and not abatement enters equation (2.4). In the production function of equation (2.3) emissions is the relevant input. According to my model, abatement is chosen before materials and emissions are determined as a consequence of all input choices and abatement. There is no additional information after the materials choice that would influence emissions.

Appendix A.3.1 and I also provide the parametric form for the Cobb-Douglas case.

[30]Following the terminology of Ackerberg, Caves and Frazer (2015) this equation is the input demand equation conditional on other inputs and in particular on labor and emissions.

2.4 Empirical Strategy

In this section, I introduce my empirical strategy to estimate the production and abatement model. The key object to estimate is the empirical counterpart of the transformed production function that represents both production and abatement technologies. The transformed production function takes the gross-output form:

$$y_{it} = \tilde{f}(k_{it}, l_{it}, m_{it}, e_{it}) + \omega_{it} + \epsilon_{it},$$

where y_{it} is log output, $k_{it}, l_{it}, m_{it}, e_{it}$ are log capital, labor, materials, and emissions respectively. ω_{it} is log productivity and consistently with the timing assumptions is observed by the firm before making decisions. Additionally, ω_{it} is additively separable from the combination of inputs. ϵ_{it} is an additive shock to log output that is not observable by the firm when making input decisions[31]. Both ω_{it} and ϵ_{it} are econometric unobservables. My main specification uses the Cobb-Douglas functional form:

$$y_{it} = \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \beta_e e_{it} + \omega_{it} + \epsilon_{it}$$

The key issue of identification is that input choices depend on the unobservable ω_{it}. In most models of production, more productive firms choose higher levels of inputs creating a positive correlation between $k_{it}, l_{it}, m_{it}, e_{it}$ and ω_{it}. As a consequence, OLS estimates of $\beta_k, \beta_l, \beta_m$ and β_e are not consistent. This endogeneity problem is well-known in the literature of production function estimation since Marschak and Andrews (1944). My proposed estimator solves the endogeneity problem by using the inverse of conditional materials input demand as a control function for ω_{it}[32]. The estimator can be thought of as a modification of Ackerberg, Caves and Frazer (2015) with emissions as an additional input[33].

[31] Alternatively, ϵ_{it} is measurement error in output. Hence, does not influence decisions.

[32] Other solutions include dynamic panel estimators (Blundell and Bond (2000) and estimators using static first-order conditions (Hall (1988) and the following literature).

[33] Other estimators of the same family include Olley and Pakes (1996) using investments and Levinsohn and Petrin (2003) using unconditional materials demand as the basis for the control function. Ackerberg, Caves and Frazer (2015) shows that both suffer from identification issues due to functional dependence. ACF style estimators have the additional advantage of allowing more flexible timing and unobservables related to the capital, labor and (in my case) emission inputs.

Next, I introduce the empirical assumptions necessary for identification. Section 2.4.2 describes the identification strategy. I turn to measurement of key variables in Section 2.4.3. Estimation follows the identification strategy closely and is discussed in Section 2.4.4.

2.4.1 Assumptions

In this section I provide necessary conditions under which my estimator can be applied to consistently estimate the primitives of my production and abatement model. Assumptions 1 and 2 are implied by the production model, whereas assumption 3 is a separate modeling assumption. The empirical assumptions are weaker and might hold for a wider range of production processes. Assumptions 1 and 2 ensure that conditional materials input demand can be used as a control function for unobserved productivity. Assumption 3 combined with the timing assumption produce the moment conditions for estimation[34].

Assumption 1 *Scalar Unobservable. ω_{it} is the only unobservable entering the conditional intermediate input demand equation: $m_{it} = h(k_{it}, l_{it}, e_{it}, \omega_{it}, \tau_{it})$*

Assumption 1 rules out multiple structural unobservables[35]. Note that in my production and abatement model, ω_{it} is a mixture of productivity and initial emission intensity. Therefore, assumptions on ω_{it} are about this combined productivity term and do differ from assumptions on the two terms separately. Because I want to allow for differences in emission prices (τ_{it}), I explicitly include it in the intermediate input demand equation. In contrast, notice that input and output prices do not enter $h(\cdot)$. Assumption 1 does not allow unobserved differences in output and input prices. If output prices differ across firms, they can be included in the composite ω_{it} term. In this case, Assumption 1 would refer to this new productivity term[36]. I discuss the issue of unobserved output and input prices in more detail Section 2.4.3.

[34] Assumptions 1-3 correspond to Assumptions 4, 5 and 2 in ACF. The information set assumptions in Section 2.3.2 correspond to ACF Assumption 1 and 3.

[35] ϵ_{it} is an econometric unobservable also not observed by the firm. Therefore, it does not enter the conditional intermediate input demand equation.

[36] In this case, I would measure a composite of TFPR and emission intensity. Note that this is conceptually correct, Assumption 2 is likely to hold in this case. For instance, in Appendix A.3.1 I show this for the Cobb-Douglas case. In contrast, input prices are conceptually harder to justify in a

Assumption 2 *Strict Monotonicity.* $h(k_{it}, l_{it}, e_{it}, \omega_{it}, \tau_{it})$ *is strictly increasing in* ω_{it}*.*

Assumption 2 is necessary for the inversion of $h(\cdot)$. The assumption is relatively straightforward to prove when materials are a static input without adjustment costs. This is exactly what I assume and in Appendix A.3.1 I explicitly show this for the Cobb-Douglas case.

Assumption 3 *First Order Markov. The distribution of* ω_{it} *evolves according to an exogenous first order Markov Process:* $F(\omega_{it+1}|I_{it}) = F(\omega_{it+1}|\omega_{it})$*.*

Assumption 3 allows for a wide range of exogenous productivity processes. For instance, it does not require ω_{it} to follow an AR(1) process as is usually assumed for dynamic panel estimators. As a consequence, it is possible to decompose $\omega_{it} = \mathbb{E}[\omega_{it}|I_{t-1}] + \xi_{it} = g(\omega_{it-1}) + \xi_{it}$. The productivity innovation ξ_{it} might follow any distribution. Higher order processes such as any predictable cyclicality is excluded. Assumption 3 does not allow firms to have influence over the evolution of their productivity. For instance, this excludes endogenous productivity evolution due to *R&D* investments[37]. Note that Assumption 3 is about the composite productivity term. Again, this is different from assuming the same separately over its individual components.

2.4.2 Identification

Given Assumptions 1 and 2, the intermediate input demand equation can be inverted to express productivity : $\omega_{it} = h^{-1}(k_{it}, l_{it}, e_{it}, m_{it}, \tau_{it})$. Plugging this expression for ω_{it} back to the production function results in the first stage estimating equation

$$y_{it} = f(k_{it}, l_{it}, m_{it}, e_{it}) + h^{-1}(k_{it}, l_{it}, e_{it}, m_{it}, \tau_{it}) + \epsilon_{it} = \phi(k_{it}, l_{it}, m_{it}, e_{it}, \tau_{it}) + \epsilon_{it}. \quad (2.5)$$

Assuming $\mathbb{E}[\epsilon_{it}|I_{it}] = 0$ nonparametrically identifies $\phi(\cdot)$. Timing assumptions of the model and Assumption 3 leads to the moment condition $\mathbb{E}[\xi_{it}|I_{it-1}] = 0$. Productivity

composite productivity term. Therefore, I explicitly have to assume away materials price differences across firms.

[37] It is relatively straightforward to relax this assumption and let the productivity process depend on observables (including choice variables). See De Loecker (2013) for an example in which a firm exports status influences its productivity evolution.

innovation in time t is independent of variables that are decided before the period. Given an estimate $\hat{\phi}(\cdot)$ of $\phi(\cdot)$ this results in the second stage equation

$$y_{it} = f(k_{it}, l_{it}, m_{it}, e_{it}) + g(\phi_{it-1} - f(k_{it-1}, l_{it-1}, m_{it-1}, e_{it-1}) + \xi_{it} + \epsilon_{it}. \tag{2.6}$$

Absent endogeneity issues, the moment condition $\mathbb{E}[\xi_{it} + \epsilon_{it} | k_{it}, m_{it}, l_{it}, e_{it}]$ would identify the parameters of the production function. However, l_{it}, m_{it} and e_{it} are determined after observing the productivity innovation ξ_{it} and therefore are likely correlated with it. The standard solution in the literature is to use lagged input variables that are decided before ξ_{it} is realized as instruments. Input choices are likely to be highly serially correlated due to persistent input prices or adjustment costs. Exogeneity is guaranteed by the timing assumptions. To implement this approach I use the following moment conditions[38]:

$$\begin{aligned}\mathbb{E}[\xi_{it}|I_{it-1}] =& \mathbb{E}[\phi(k_{it}, l_{it}, m_{it}, e_{it}) - f(k_{it}, l_{it}, m_{it}, e_{it}) \\ & - g(\phi(k_{it-1}, m_{it-1}, l_{it-1}, e_{it-1})) - f(k_{it-1}, l_{it-1}, m_{it-1}, e_{it-1})|I_{it-1}] = 0\end{aligned}$$

Identification of the production function parameters comes from both cross-sectional and time-series variation in inputs and output. Input elasticities are identified from the comovement of exogenous variation in inputs and output. Therefore, it is essential to establish that there is independent variation in all input variables after controlling for other input choices. Otherwise, the GMM rank condition does not hold and input elasticities might not be identified. First, as Ackerberg, Caves and Frazer (2015) describes using moments of the productivity innovation solves the potential problem that l_{it} is functionally dependent on k_{it}, m_{it}[39]. Second, Gandhi, Navarro and Rivers (2020) point out that the material input elasticity is not identified from ACF type of moments if only panel data on inputs and outputs is used. This is the

[38]Ackerberg, Caves and Frazer (2015) suggests using $\mathbb{E}[\xi_{it} + \epsilon_{it}|I_{it-1}] = 0$ for identification. They use $\mathbb{E}[\xi_{it}|I_{it-1}] = 0$ in an earlier version, see their footnote 11 for details for why this might be desirable. I follow the method of their Monte Carlo simulations and use their concentrating out procedure to reduce the dimension of the search over parameters as documented in Appendix A.4 of their paper. There they also use $\mathbb{E}[\xi_{it}|I_{it-1}] = 0$.

[39]Several reasonable data generating processes lead to functional dependence of l_{it} on k_{it}, m_{it}, when unconditional intermediate input demand is used to control for unobserved productivity (such as in Levinsohn and Petrin (2003)). Ackerberg, Caves and Frazer (2015)'s solution is to use conditional (on l) intermediate input demands to control for unobserved productivity.

main reason Ackerberg, Caves and Frazer (2015) recommends using their method for estimating value-added production functions only. They also suggest using the static first-order condition to identify the materials elasticity. In my environment, this is not possible to implement since this first-order condition depends on emissions and allowance prices[40]. An alternative solution is to introduce observed heterogeneity in intermediate input prices to identify β_m (Doraszelski and Jaumandreu (2013)). Using proxies for allowance prices introduces variation similar to observed input prices in intermediate input demand equation. Therefore, the materials coefficient is identified in my model[41]. Finally, it is worth discussing the identifying variation of β_e, especially when τ_{it} is included in the control function. Conditional on all other inputs and prices, emissions can still vary across firms due to differences in abatement choices. Adjustment costs in abatement expenditures can provide variation in abatement independently of prices which also translates to serial correlation in emissions. Hence, introducing emissions to the production function does generally not introduce additional functional dependence issues.

2.4.3 Measurement, Unobserved Prices and Firm-Plant Level Data Issues

I use revenues, fixed assets and materials expenditures as measures for output, capital and intermediate inputs respectively. I measure all variables in euros (€) using market exchange rates. Data limitations do not allow me to use quantity measures or the perpetual inventory method for capital[42]. Using monetary values might be preferred when input quality can differ (Ackerberg, Caves and Frazer (2015) or when

[40]Using value-added production functions with emissions is hardly meaningful. Emissions would need to increase in value-added but not materials usage. This could lead to situations when the firm increases its materials share which should result in lower emissions.

[41]Note that it is likely that some of the models assumptions do not hold perfectly in practice. If there is variation in materials prices or adjustment costs in the materials the control function assumptions might fail. However, these would provide identifying variation for the materials coefficient.

[42]There are well known issues with using fixed assets as a measure for capital. Accounting measures of depreciation are based on averages and are likely not accurate. For instance, the US Census Bureau uses a perpetual inventory method based on industry level depreciation rates and investments Raval (2019). Depreciation rates are not available in the EU and investments in Orbis is not available for most firms.

firms manufacture multiple products (De Loecker and Goldberg (2014))[43]. Additionally, since materials is an aggregate of different inputs (e.g. iron-ore, coal, alloying elements for steelmaking) it is best measured in monetary terms. Similarly to most production datasets, output and input prices are unobserved in my data[44]. To control for across time variation, I deflate revenues and materials expenditures with industry specific output price deflators[45]. As is well known in the literature, two key issues arise when there are unobserved differences in output prices across firms (De Loecker (2011)). First, elasticity estimates might be biased if inputs are correlated with prices. For instance, when firms face downward sloping output demand curves, output and prices move in opposite directions. As a result, input choices and output prices are also correlated. Second, productivity estimates contain price and demand variation. That is, lack of data on prices or demand conditions, it is not possible to differentiate the impact of prices, demand shocks and quantity productivity on revenues. The standard solution in the literature is to introduce assumptions on demand and the form of competition. This does not necessarily make the problem go away, but it provides explicit formulas for the bias and how a composite productivity term should be interpreted. When the price variation is exogenous (such as with perfect competition and exogenous demand shocks), the input elasticity estimates are not biased. Generally, less elastic demand (or more market power) results in upward biased coefficient estimates. As a result, scale estimates are informative of the likely extent of the potential bias. In Section 2.5, I show that returns to scale estimates are close to one indicating that perfect competition is likely not an unrealistic assumption. In any

[43]If input quality differs, I need to assume that firms have access to the same menu of prices and qualities. Product level estimation is usually difficult for multiproduct firms even with better data. Intermediate inputs are almost never observed at the product level.

[44]Exceptions include the US Census of Manufacturing or the ESEE dataset from Spain. There are also well-documented issues with prices when using these datasets, see De Loecker and Goldberg (2014) for more details.

[45]Theoretically, one should use a separate input price deflator for all inputs or when this is not available an aggregate materials deflator. However, this is not available for EU datasets. In this case, the literature uses the output deflator for materials too. The usual justification is that the two are likely to be highly positively correlated. Eurostat reports NAICS 3 level output deflators and I assign these to firms based on primary NAICS 3 codes from Orbis. The base year for the output deflators is 2015, so all values are measured in 2015 Euros.

case, ω_{it} should be interpreted as measuring both productivity and demand effects[46]. Using intermediate input demand remains conceptually correct but now it controls for the composite productivity term. These conditions also provide a microfoundation for Assumption 1 which explicitly assumes the two issues away.

Unobserved intermediate input prices and allowance prices can introduce an additional bias through the control function. Since I do not observe intermediate input prices, I have to assume away price differences across firms. For instance, assuming perfect competition in input markets provides a justification for this assumption[47]. Allowing for differences in emissions prices and the possibility that firms do not equate marginal emissions prices to costs is crucial for my empirical approach. At the same time, these unobserved differences can lead to misspecification of the control function. I observe market emission prices and allowance allocation from which I build flexible controls for emission prices. The key assumption is that firms take emissions prices into account in a similar way in intermediate inputs decisions. This is different and weaker than assuming homogenous emissions prices and static optimization of emissions decisions.

Finally, my measures of emissions and allowances are available only for plants under the regulation of the EU ETS. In contrast, financial variables are available at the firm level. Firms in my dataset might have activities that belong to plants outside of my emissions data. For instance, cement firms are often vertically integrated with concrete producers. However, only cement production emits carbon dioxide and is regulated under the EU ETS[48]. Therefore, vertically integrated firms in my data appear to have higher revenues for the same level of emissions than a firm that is only active in cement. Alternatively, the emissions are measured too low for vertically integrated firms. This non-classical measurement error in emissions leads to biased coefficient estimates. The sign of the bias is theoretically ambiguous and empirically

---

[46]In Appendix A.3.2 I derive the elasticity bias and formulas or the composite productivity term. I assume a CES demand system with exogenous shocks, monopolistic competition and a Cobb-Douglas production function.

[47]An alternative is to assume that output and input price biases offset each other. De Loecker and Goldberg (2014) shows provides conditions for when this holds. Note that in this case, the elasticity estimates are unbiased but it is less clear what the productivity term measures.

[48]Other prominent example is power generation and distribution.

hard to predict (see Appendix A.4). In Section 2.5.3, I report a series of robustness exercises to validate that my results are not driven by this measurement error.

2.4.4 Estimation

I estimate the model for the ten industries defined in Section 2.2.3 separately. The underlying assumption is that technology is the same for all firms in an industry. This also excludes differences across countries in technology for instance due different levels of development[49]. Estimation follows the identification strategy closely. As a first stage, I estimate Equation (2.5). In the Cobb-Douglas case, the conditional materials demand equation is linear (in logged variables). Therefore, my main specification includes a linear control function and first-stage equation. I check the robustness of my results to a second-order polynomial control function. I also experiment with different proxies for emission prices (τ_{it}) in the control function. The main specification does not include emission price proxies since I find that including these does not change results significantly[50]. See Section 2.5.3 for details on the consequences of these choices on my empirical results.

To implement the second-stage (Equation 2.6), I use instruments (Z) from I_{it-1} to generate unconditional moments of the form: $\mathbb{E}[\xi_{it} \otimes Z'_{it-1}] = 0$. Standard choices for instruments include all inputs possibly going several periods back. I use variables going back to $t-2$ as instruments:

$$\{k_{it}, k_{it-1}, k_{it-2}, l_{it-1}, l_{it-2}, m_{it-1}, m_{it-2}, e_{it-1}, e_{it-2}\}.$$

After experimenting with estimation I find that adding additional lags significantly decrease estimated standard errors while do not meaningfully alter coefficient estimates. I assume that productivity follows an AR(1) process, with parameter ρ measuring persistence. This is somewhat restrictive but is widely applied in practice[51]. In order

[49]The model could be estimated for country-industry pairs but the data does not allow this. The structural production function estimator requires a bigger sample size and since I need to observe emissions, the sample size is given by the EU ETS.

[50]It is likely that firms do not take into account emissions prices when ordering materials. The impact of emission prices is likely to be limited for most firms.

[51]See the discussion in page 9 of Collard-Wexler and De Loecker (2016) for the trade-offs of this assumption.

to decrease the dimensionality of the search I follow Ackerberg, Caves and Frazer (2015) and concentrate out the constant term and parameter ρ. I find the parameters that minimize the distance of the sample analog of the unconditional moments from zero using a generalized method of moments (GMM) procedure. I estimate standard errors using a firm-level block bootstrap procedure with 100 iterations.

2.5 Results

Now I report the main results of this paper. I start with my production function estimates in Section 2.5.1. I discuss in detail my estimated abatement elasticities, the key empirical parameter of interest. Using my model's implications and the estimated elasticities, I calculate marginal abatement costs for all firms in my data. Section 2.5.2 reports these estimates. I discuss the robustness of my estimates in Section 2.5.3. Finally, Section 2.5.4 provides reasons for why my marginal abatement cost estimates are significantly above market prices for emission allowances.

2.5.1 Production Function Estimates

Table 2.2 reports production function estimates for the main specification: Cobb-Douglas with a linear control function that does not contain emission price proxies. The input elasticity estimates of labor, capital and materials are reasonable and stable across specifications. The low capital coefficient estimates are not unusual for gross output production functions (Collard-Wexler and De Loecker (2015), Gandhi, Navarro and Rivers (2020)). The elasticity estimates are also broadly similar to estimates in the literature for similar industries. I find my estimates of steel, cement and paper similar to Collard-Wexler and De Loecker (2015), Backus (2020), and Doraszelski and Jaumandreu (2018) and Demirer (2020) respectively[52]. The returns to scale

[52] My elasticity estimates need to be divided by 1-β_e to be able to compare but this does not make a meaningful difference for most industries. The returns to scale estimates are essentially the same using both forms of the production function. Note that the above papers use vastly different methdologies and datasets to estimate production functions. Gandhi, Navarro and Rivers (2020) and Demirer (2020) uses non-parametric estimators for the widely used plant level Colombian and Chilean data. Gandhi, Navarro and Rivers (2020) finds average cross-industry capital elasticities of 0.14 (Colombia) and 0.16 (Chile). Demirer (2020) finds average β_k, β_l and β_m are 0.09, 0.36 and 0.59 for the Colombian paper industry. Doraszelski and Jaumandreu (2018) estimates are 0.09, 0.321 and 0.621 for Spanish paper firms. Collard-Wexler and De Loecker (2015) uses Olley-Pakes type of control function and plant level steel data from the US Census of Manufacturers. Their

estimates are statistically not different from one for most industries. The estimates indicate weak market power for chemicals, refineries, power companies and do not signal misspecification[53]. I report statistics for total productivity $\omega_{it}+\epsilon_{it}$. My average persistence estimate is 0.79 while the interquartile range is 38 log points[54]. Altogether, I interpret these estimates as supporting evidence that introducing emissions as inputs to the production function does not change estimates fundamentally. This conclusion is ex-ante implied by my model but empirically did not have to be so. Additionally, these results provide confidence in the data and in particular, against serious issues due to vertical integration.

The average β_e estimate is 0.068[55]. Decreasing emissions by 1% while keeping other inputs constant is associated with a drop in output of 0.068% for the average firm in my data. Generally, a higher β_e estimate indicates that emission intensity is less responsive to abatement expenditures. Based on this metric, it is the most difficult to decrease emissions in cement and the least difficult in chemicals. Decreasing emissions by 1% while keeping other input level constant is associated with a drop in output of 0.127% for the average cement firm in my data. An alternative to interpret my β_e estimates is through the lens of the structural model of production and abatement. Decreasing the share of output that is not dedicated for abatement by 1% is associated with a drop in emission intensity of $\frac{1}{0.127} = 7.9\%$ for the average cement firm. Alternatively, a 1 percentage point increase in the share of abatement expenditures translates to roughly 7.9% drop in emission intensity for the average cement firm[56]. It is somewhat difficult to get a good sense of how realistic the mag-

benchmark estimates for β_k, β_l and β_m are 0.08,0.27 and 0.68 respectively. Backus (2020) uses an input cost-share based method and finds a close to zero capital coefficient estimates for US cement.

[53]The lowest implied demand elasticity is close to 10. I find lower returns to scale estimates when I use a second-order polynomial control function. The lowest implied demand elasticities are in the range of 4 to 6 which is more in line with what the literature considers reasonable (Asker, Collard-Wexler and De Loecker (2014), De Loecker (2011)). In any case, my returns to scale estimates are below or around one and do not indicate misspecification.

[54]This is broadly in line with the literature. For instance, Asker, Collard-Wexler and De Loecker (2014) reports a median productivity persistence in the US census data to be 0.85. Syverson (2004) reports an interquartile range of 28 log points.

[55]I exclude category other from this comparison. The average is weighted by the number of observations in each industry.

[56]Notice that the share of abatement expenditure is likely to be very low for the average firm.

nitude of these estimates are. In Section 2.5.2, I calculate marginal abatement costs which is easier to interpret. Nevertheless, it is clear that estimates of β_e are substantially above the revenue share of emissions (see Table 2.1). The average β_e is 0.068 whereas revenue share is around 0.01. Robustness exercises (see Section 2.5.3) confirm that this difference is unlikely to be caused by measurement error due to vertical integration. In a static, frictionless world, input revenue shares and elasticities are equal. However, firms operate in a dynamic, uncertain environment when both adjustment and transaction costs are important. Therefore, the cost share estimator likely underestimates the input elasticity of emissions[57].

Therefore, a percentage point and a percentage change in $1 - a_{it}$ are numerically similar. However, these are very different for a_{it} itself. For instance, when a_{it} grows from 1 to 2% of output corresponds to doubling abatement expenditures. $1 - a_{it}$ would decrease from 0.99 to 0.98, that is, roughly by 1%.

[57]My marginal abatement cost estimates are above market prices for similar reasons. See the detailed discussion in Section 2.5.4.

Table 2.2: Production Function Estimates

	(1) Cem	(2) Cer	(3) Chem	(4) Glass	(5) Non-ferr	(6) Oth	(7) Paper	(8) Power	(9) Ref	(10) Steel
Other input elasticities										
β_l	0.448	0.463	0.104	0.324	0.044	0.348	0.294	0.163	0.073	0.294
	(0.044)	(0.031)	(0.078)	(0.043)	(0.078)	(0.027)	(0.043)	(0.027)	(0.121)	(0.040)
β_m	0.348	0.360	0.431	0.494	0.708	0.436	0.532	0.404	0.487	0.624
	(0.042)	(0.023)	(0.056)	(0.029)	(0.059)	(0.021)	(0.046)	(0.035)	(0.046)	(0.034)
β_k	0.073	0.122	0.339	0.103	0.111	0.184	0.128	0.294	0.193	0.040
	(0.028)	(0.015)	(0.070)	(0.026)	(0.054)	(0.017)	(0.018)	(0.026)	(0.095)	(0.015)
Emission elasticity										
β_e	0.127	0.120	0.015	0.058	0.068	0.022	0.022	0.051	0.125	0.028
	(0.021)	(0.016)	(0.020)	(0.015)	(0.032)	(0.006)	(0.009)	(0.014)	(0.062)	(0.013)
Returns to scale										
$\beta_k + \beta_l + \beta_m + \beta_e$	0.995	1.064	0.888	0.979	0.931	0.990	0.976	0.912	0.878	0.986
	(0.012)	(0.012)	(0.035)	(0.021)	(0.063)	(0.008)	(0.012)	(0.012)	(0.055)	(0.010)
Productivity										
Persistence	0.867	0.763	0.791	0.808	0.882	0.876	0.802	0.760	0.924	0.709
Median	2.086	0.958	3.220	1.927	2.567	1.582	1.723	2.814	4.325	1.542
Inter-quartile range	0.400	0.388	0.468	0.300	0.359	0.467	0.233	0.653	0.768	0.238
Observations	1713	4534	901	1953	274	17674	3988	2684	506	1785

Notes. An observation is a firm-year, all observations (unbalanced panel) included, 2005-2016. Block bootstrapped (firm level) standard errors reported in parenthesis, 100 iterations. Control function is linear in log inputs, emission price controls not included. Instruments: $\{k_{it}, k_{it-1}, k_{it-2}, l_{it-1}, l_{it-1}, m_{it-1}, m_{it-2}, e_{it-1}, e_{it-2}\}$. Estimated productivity $\hat{\omega}_{it}$ calculated as $\hat{\phi}(\cdot) - \hat{f}(\cdot)$, does not contain measurement error ϵ_{it}. Persistence is calculated as the regression coefficient on $\hat{\omega}_{it} = \rho\omega_{it-1} + \nu_{it}$.

2.5.2 *Marginal Cost of Abatement Estimates*

In this section, I present marginal abatement cost (MAC) estimates for carbon dioxide emissions at the individual firm level. As I show in Appendix A.3.5, under the assumptions of my model marginal cost of abatement is equivalent to the marginal revenue product of emissions (MRPE) in the transformed production function. My measure of MRPE (denoted by Λ) is purged of measurement error[58]:

$$\Lambda_{it} = \mathbb{E}\left[\frac{\partial Y(K,L,M,E)}{\partial E}\right] = e^{\omega_{it}}\mathbb{E}[e^{\epsilon_{it}}]\frac{\partial f(K,L,M,E)}{\partial E}.$$

Λ_{it} measures the expected revenue impact of an additional ton of CO_2 emissions in the beginning of time t at the observed levels of inputs and productivity. Alternatively, it represents additional abatement expenditures (in €) when the decreases CO_2 emissions by one ton[59]. As such, Λ_{it} does not include adjustment costs associated with changing abatement expenditure or emission levels. Λ_{it} best represents the marginal cost of abatement at the observed level of output and emissions without adjustment costs. With the Cobb-Douglas functional form and after taking logs this simplifies to:

$$\Lambda_{it} = \beta_e \frac{Y_{it}}{E_{it}}\mathbb{E}[e^{\epsilon_{it}}]e^{-\epsilon_{it}} \quad \text{and} \quad \lambda_{it} = ln(\mathbb{E}[e^{\epsilon_{it}}]) - \epsilon_{it} + y_{it} - e_{it} + ln(\beta_e).$$

Marginal abatement costs are high when the output elasticity of emissions (β_e) is high, that is, when emissions are not very responsive to abatement expenditures. Low emission intensity of production (high $y_{it} - e_{it}$) translates the estimated elasticity impacts into high nominal values of MACs.

Table 2.3 reports the distribution of marginal abatement cost estimates for the industries in my data. The median estimated MAC across all industries is 90 €/t. There is vast heterogeneity both within and across industry. First, across industries the median MAC ranges from 22 €/t in cement to 717€/t in non-ferrous metals.

[58] I experiment with a measure that includes measurement error as well. The key results do not change, the distribution remains very similar.

[59] Since it is the derivative of abatement expenditures in emissions and hence should be interpreted as a marginal impact. I decided for the simpler interpretation in the main text.

Table 2.3: Marginal Abatement Cost Distribution Across Industries

	Cem	Cer	Chem	Glass	Non-ferr	Paper	Power	Refineries	Steel	All
Cross-industry										
β_e	0.13	0.12	0.01	0.06	0.07	0.02	0.05	0.13	0.03	0.07
Y/E median	149	714	3113	1588	7895	2646	407	2104	3793	1693
Λ_{it} ***distribution***										
10th percentile	11	29	6	41	46	22	9	22	17	22
25th percentile	15	45	16	58	73	36	16	48	54	37
Median	22	98	48	96	717	66	35	363	116	90
75th percentile	36	235	149	170	2319	157	238	1053	199	239
90th percentile	113	523	429	397	6667	499	3593	2554	378	1045
Observations	1713	4534	901	1953	274	3988	2684	506	1785	18338

Notes. An observation is a firm-year, unbalanced panel, 2005-2016. Λ_{it} is measured in €/ton CO_2 equivalent. Λ_{it} represents the expected marginal abatement cost for firm i in time t at the observed level of output and emissions. Control function is linear in log inputs, emission price controls not included. Firms in category other not reported and excluded from "All".

Both, the inverse of emission intensity and abatement elasticities contribute to these estimated differences. For instance, the five times bigger median MAC for ceramics than cement is almost entirely due to emission intensity differences. In contrast, the difference between glass and steel can be best explained by differences in abatement elasticities. My estimates suggest that the lowest costs of decreasing emissions are in cement, power and chemicals. Second, I find that the within industry dispersion is substantial in all industries. The MAC of the 75th percentile plant is 3-10 times bigger than that of the 25th percentile in most industries. Again, this is mostly due to heterogeneity in emission intensity at the firm level.

Median marginal abatement costs are higher than most previous estimates but the relative ranking of industries is similar. The general consensus in the EU ETS considers switching from coal to gas in power generation one of the cheapest abatement options. Several other options in the power sector, such as switching to wind generation, have low estimated abatement costs (Gillingham and Stock (2018)). Consequently, Nauclér and Enkvist (2009) predicts marginal abatement costs in the power sector around 20€/t in 2030. Evidence for other industries is more scarce. Average predictions from Nauclér and Enkvist (2009) for 2030 are around 0€/t for cement, 5€/t for chemicals and 18€/t for steel. In contrast, Borenstein et al. (2019) assumes that there is no price-responsive abatement in manufacturing industries indicating a belief in high MACs. More importantly, my MAC estimates are also significantly above market prices that predominantly stayed in the 0-20€/t range in 2005-2016. Both the level and dispersion of my MAC estimates are likely to suffer from measurement error in emissions. Therefore, I perform several robustness checks in Section 2.5.3. Section 2.5.4 explores potential explanations for why my MAC estimates are consistently above market prices.

2.5.3 Robustness

I examine various alternative specifications to validate the robustness of my elasticity estimates. Here I overview the main findings, while detailed results are reported in Appendix A.5. First, I address the possible bias due to non-classical measurement error in emissions associated with vertical integration. I try several ways of narrowing my sample to firms that are less likely to be vertically integrated and to own plants outside the regulation of the EU ETS. My elasticity estimates do not significantly change when I narrow my sample to firms with the modal NACE4 code in an industry, to firms that are active only in a specific four-digit NACE industry but nothing else, to single plant firms or to small firms. I also run my estimates on a narrower industry specification based on EU ETS activity types and find that the average estimates at the main industry level are very similar. In addition, these also results suggest that any potential industry classification misspecification is unlikely to significantly influence my findings. Finally, I find similar results using fixed effects estimators that can partially control for persistent measurement error at the firm level[60]. Second, I experiment with different control functions. Appendix A.3.1 presents the theoretical foundations for including emission prices in the control function. Since I do not observe individual emission prices, I use various proxies based on the market price, firm level allowance allocation and banking decisions. The elasticity estimates are similar across all specifications. The correlation of the $\hat{\beta}_e$ with the main specification is in between 0.7 and 0.95. Adding emission price controls does increase the standard error estimates. Using second-order polynomials does not change the emission elasticity estimates substantially. Returns to scale is slightly lower and the second-order polynomial specification attributes a bigger part of productivity endogenous. I also experiment with different instruments. Using only instruments from last

[60] As well known in the literature, fixed effects methods can exacerbate the measurement error in the capital coefficient and result in a downward bias for β_K. Likewise, I find that the capital coefficient estimate is lower in these specifications.

year's information set (I_{t-1}) does not change the estimates meaningfully but increases standard errors.

I also consider the robustness of my marginal abatement cost estimates. These are influenced by the elasticity estimates and individual firm level emission intensities in the data. In addition to potentially biased elasticity estimates, emission intensities might be undermeasured for firms that are vertically integrated. Therefore, this issue is likely to be more pronounced for MAC than for elasticities. In particular, measurement error is expected to increase MACs estimates. I analyze the MAC distribution for the different industry specifications where the measurement issue is likely to be the most severe. The results suggest that the impact of measurement error on MAC estimates are likely to be small. Even in the specifications with the lowest MAC estimates, median Λ_{it} across all firms is around 70€/t, around 20% lower, than in the baseline estimate. MACs in other specifications are similar or higher than in the baseline. Finally, the ranking of industries is strikingly similar across all specifications. Altogether, these results suggests that my key findings are unlikely to be driven by measurement error, industry misspecification, control function and instrument choice.

2.5.4 Why Are Marginal Abatement Costs Significantly Above Allowance Prices?

Economic theory suggests that without frictions a market for emission allowances equates marginal abatement costs with the allowance price. As a consequence, marginal abatement costs are equal across firms. The EU ETS market price moved in the range of 5 to 15 €/t in most of my sample period (see Figure 2.1). Most of my individual marginal abatement costs estimates are significantly above this price range. My estimated median across all industries is an order of magnitude higher. This suggests that firms should find it optimal to increase emissions which drives down their marginal abatement costs and increase market prices. I argue that the combination of asymmetric adjustment costs, long-run expectations and the dynamic nature of abatement decisions are the most likely explanation for the discrepancy between economic theory in a frictionless world and my empirical estimates. Start with adjustment costs[61]. A

[61] Adjustment costs are consistent with my model and empirical strategy but are not included in my marginal abatement costs estimates.

majority of abatement options are at least partially irreversible. Energy efficiency improvements and switching from coal to gas are challenging to reverse. Modern, emission efficient plants are difficult to switch to older capital. The irreversibility creates asymmetric adjustment costs. Adjusting to higher levels of abatement expenditures or abatement capital might be costly but it is even more costly to go back to lower levels. Supporting the existence of significant adjustment costs, I find that the dispersion in the marginal product of emissions is comparable to capital and is above material and labor (Appendix A.6). Moving to expectations. The program were created to slowly restrict the supply of allowances after a relatively generous starting pool. Firms knew they had to decrease emissions on the long run. As a response to low prices, the European Commission considered drastically decreasing the amount of allowances or introducing price floors for a long time. Firms rightly expected supply to tighten and prices to rise on the long run. Prices were consistently above 30€/t in 2020 supporting a possible regime change in price dynamics.

In this environment, it can be optimal for firms with high marginal abatement costs not to increase emissions. The sum of adjustment costs and emission prices might be higher than the marginal abatement costs. Then, the costs of increasing emissions overweigh the benefits even in the static sense. Additionally, the firm has an incentive not to raise emissions today if it expects to emit less on the long run. This is realistic in an environment with high future price expectations. In this case, the firm has to pay adjustment costs again when it eventually settles on low emissions. How do firms reach a state with high marginal abatement costs? Alternatively, how do firms become overly emission efficient? Firms might start with relatively modern and emission efficient plants in the beginning of the program. The technology might be inherently emission efficient even without explicit abatement. Finally, in an uncertain dynamic environment with adjustment costs in inputs, dispersion in marginal products can arise naturally as adjustment is not perfect to changing conditions[62]. Since marginal abatement costs are equivalent to the marginal product of emissions in my model, this provides a natural explanation for the dispersion in my MAC estimates.

I consider two alternative explanations. First, firms might find that the transac-

[62]For instance, see Asker, Collard-Wexler and De Loecker (2014) for dispersion in the marginal product of capital.

tion costs of purchasing additional allowances are too high. In 2015, 48% of firms with above market MACs had more allowances that they ended up using. Generally, firms with high MACs tend to be significantly long on allowances. Transaction costs are not a relevant constraint for these firms. Around 75% of underallocated have MACs at least 15€/t higher than the market price. With transaction costs 15€/t, around 90% of all plants would find it optimal to additional allowances and increase emissions assuming the market price did not change. Second, firms might not make abatement and trading choices optimally. Around 50% of firms do no trade in my data. This suggests, that many firms might take the allowance allocation as a hard cap on emissions. As a consequence, firms decrease emissions below the cap even if it would be cheaper to buy allowances in the market. However, 75% and 50% of firms who are overallocated end up emitting at least 15% and 34% less than allowances received. It is difficult to see why firms would go under their limits by this margin. Alternatively, survey evidence shows that most firms think of abatement as a byproduct of energy efficiency improvements. My results are consistent with this hypothesis only if firms do not internalize the impact of emissions costs on energy efficiency projects. In this case, my estimates suggest that firms overinvest in energy efficiency improvements. Altogether, transaction costs and lack of optimization are unlikely to contribute significantly to the gap between MACs and market prices.

2.6 Implications for the EU ETS Market and Environmental Policy

Market prices were generally considered to be too low in the 2005-2016 time period. For most of this time frame, prices were below the social cost of carbon (SCC). An efficient regulation should equate allowance prices with the social cost of carbon. The Obama Administration's SCC estimate was around $30/t and served as a guidance for global social cost of carbon estimates in the early years of the EU ETS. The current scientific consensus is that the SCC is likely above $100/t in 2020 (Carleton and Greenstone (2021)). The European Commission long recognized the need to decrease emission caps or to introduce price floors to help raise prices. My results shed some light on what to expect if market prices get closer to the social cost of carbon. My marginal abatement cost estimates predict where to expect abatement at different levels of market prices. At a carbon price of 30€/t, many cement, chemicals and power firms would find it optimal to decrease emissions[63]. In contrast, most glass and non-ferrous metal producers would rather buy allowances at market prices. When market price reaches the 100€/t range, the majority of firms in most industries are expected to decrease emissions. However, there is significant overlap in the distribution across industries to expect abatement in all industries at reasonable price levels.

My marginal abatement cost estimates are not informative of by how much, how soon and through what mechanism would firms decrease emissions as a response to different market prices. Generally, the key limitation of my approach is the difficulty of predicting non-marginal changes. Any non-marginal response is associated with adjusting the capital, labor and material inputs as well as emissions. Of these, capital is particularly problematic as it is a dynamic input with likely significant adjustment costs. Therefore, I would need to specify, solve and simulate a full dynamic model of investment. Nevertheless, my results shed some light on the key characteristics of the equilibrium at different levels of emission caps. Since marginal abatement cost estimates are relatively continuous, it is unlikely that higher emission caps would lead to dramatic carbon price increases.

Finally, my estimates can provide guidance on which industries to target with environmental policies. It is beneficial for society to incentivize firms with marginal abatement costs below the SCC to decreases emissions. It is unlikely that the United States will introduce a carbon price or carbon market in the foreseeable future due to political constraints. However, subsidizing emission intensity improvements might be feasible. Targeting industries with the lowest marginal abatement cost estimates (cement, chemicals, power) would increase social welfare the most. Since technology is likely similar for the same industry in different locations, these results are relevant for a wide range of countries.

[63]All of my statements in this section assume downward adjustment costs away. Small and homogenous adjustment costs do not change these statement substantially.

2.7 Conclusions

This paper uses production function estimation methods to understand the trade-off of productivity and carbon dioxide emissions. I propose a production and abatement model to help identify the abatement elasticity of emissions without explicitly observing abatement expenditures. I estimate the model using data from the EU ETS for 2005-2016. Then, I use my estimates and the model's implications to calculate marginal abatement costs for participating firms. MACs are generally above prevailing emission allowance prices but below reasonable estimates of the social cost of carbon. I find vast heterogeneity within and across industries. The distribution of MACs is continuous and spans a wide range. As a consequence, it is unlikely that the price increase to tightening emission caps would be radical. My methodology can be applied to other setups where data on financial statements and emissions are available. This is important, because many environmental policy decisions depend on the quantitative trade-off between growth and harm to the environment. For instance, my estimates can help target industries with subsidies to decrease emissions when introducing a carbon tax is not politically feasible.

One caveat of my analysis is that I only observe firm level data on production variables whereas emissions are at the plant level. This potentially introduces non-classical measurement error especially for firms that are vertically integrated. I perform several robustness exercises to understand how this might influence my results. I find that my key results are unlikely to be driven by this data limitation. Nevertheless, future research using plant level data is needed to exclude this possibility. Manufacturing census data is available for several countries in the EU. An additional benefit of using these data sources is the ability to differentiate mechanisms that operate at the plant and firm levels. In general, this data does not allow a detailed empirical analysis of abatement mechanisms. An additional caveat is that my marginal abatement costs results are only indicative of likely market outcomes to different policies. In particular, they are not informative of the extent and timing of emission adjustments. Answering these questions would require a fully specified dynamic model of investment and abatement.

3

Technology Transition and Environmental Regulation in Power Generation

3.1 Introduction

Slowing down climate change and global warming is likely the most important challenge of humanity in the 21st century. Now it is well established, that greenhouse gas (GHG) emissions are the key reason behind climate change and global warming. Power generation is responsible for around one-third of total greenhouse gas emissions in developed countries[1]. As such, any plans to slow down climate change must put a large weight on transitioning away from fossil fuels in power generation. Not surprisingly, leading developed countries made commitments to almost completely eliminate carbon-dioxide emissions from power generation by the middle of the 21st century. For instance, Germany's 2010 plans included decreasing GHG emissions in power generation by 80-90% by 2050. Since then, many countries made similar commitments. However, it is not clear how this transition will take place and what economic policy can and needs to do to support it.

Since GHG emissions are a negative externality, absent of regulation, neither the

[1] For instance, the Energy Information Administration (EIA) reports 33% for the United States in 2018

equilibrium nor the transition is efficient. Transition is likely to be highly welfare relevant in power generation. High entry costs and long-lasting capital result in slow transition. For instance, coal power plants require an initial investment $1-5 billion and have a lifetime of at least 40 years. Early retirements are costly exacerbating the dynamic inefficiency of the negative externality. Lack of other frictions, a carbon tax would restore the efficient equilibrium both in a static and dynamic sense. However, introducing carbon taxes is politically not feasible in the United States[2]. Therefore, second-best policies are likely to be important in the energy transition of the United States. It is both theoretically and empirically uncertain how close these second-best policies can get to the efficient outcome.

In this paper, I study the speed of transition and how it interacts with policy instruments. I focus on PJM, the biggest electricity system in the United States. Recent market trends and geographical characteristics suggest that the key transition in PJM is away from coal and towards natural gas. I develop a non-stationary fully dynamic competitive model of power generation. The problem is essentially non-stationary as it is focused on the path from the current state towards a new equilibrium. I introduce multiple technologies with different cost structures in the model, a key feature of electricity markets. In every period, gas power plants can enter while coal power plants can exit the market. To simplify computation I assume that other technologies are exogenously on the market but can not enter or exit. I model yearly profits as an outcome of competition on hourly spot markets. Demand variation is a key feature of how firms make profits in this market. High marginal cost power plants only produce in high demand hours but usually tend to have lower fixed costs to recover. I calibrate most of the parameters, such as marginal costs, entry and other fixed costs, demand, of the model using publicly available data. I pick the few

[2]The Obama administration tried introducing a cap and trade system that implicitly puts a price on carbon emissions and under certain conditions is equivalent to a carbon tax. However, they could not get any legislation passed due to serious opposition both in the US Senate and in the US population. Most Democratic presidential candidates in 2020 did not talk about a carbon tax explicitly even though they support environmental regulation.

parameters with no clear guidance from the literature to match historical patterns in the data and produce reasonable results. I consider three instruments that are easy to implement in the model, carbon taxes, gas entry subsidies and coal exit subsidies. I pick the levels of subsidies so as to produce a similar long term outcomes as the carbon tax.

My findings highlight the importance of analyzing the full transition path when comparing environmental policy instruments. Policies that lead to similar long-term outcomes induce vastly different transition dynamics. A carbon tax set to $\$30/CO_2$ ton is associated with an almost immediate entry of the long-run gas capacity. The impact of carbon tax is radical because it makes coal immediately redundant for competition. In contrast, gas entry and coal exit subsidies do not change spot market competition fundamentally and result in a slower and smooth transition. The existing coal fleet slows down gas entry. These differences are also welfare relevant. Both gas entry and coal exit subsidies improve on the baseline scenario by inducing capacities that are close to the carbon tax optimum. Neither of these eliminate the static inefficiency due to the emission externality. However, the welfare improvements are relatively minor. Introducing high gas subsidies that overincentivize entry on the long run improve welfare compared to lower levels. The dynamic analysis points to higher optimal gas subsidies compared to two-period model.

This paper contributes to three separate literatures. First, I contribute to the large literature in economics, operation research and engineering about modeling investments in power generation. Most of the economics literature uses two-period models[3]. Recently there is interest in applying fully dynamic models. Bushnell and Ishii (2007) describe a fully dynamic model with market power but acknowledge that it is only possible to be solved for a few number of firms and time periods. Cullen and Reynolds (2017) solves the planner's problem in a fully dynamic model where plants

[3]This is understandable given the complications of spot markets of electricity. Most papers take this seriously (second stage) and focus on a long term equilibrium.

choose in each period whether to be available for production. Long term planning models are widely used in operations research, engineering and policy analysis. For instance, both the Energy Information Administration (EIA) and the Environmental Protection Agency (EPA) uses a long term planning models to simulate the impact of environmental policy[4]. These models minimize the total cost of producing electricity under constraints defined by technological (e.g. capacities, transmission constraints) and economic conditions (marginal costs, startup costs. The methods essentially solve the social planners problem and can be applicable as long as the decentralized and the social planner solution coincide. In contrast, I present an approach to find the decentralized equilibrium in a model that is able to incorporate several of the complicating features of electricity markets. As such, my method can be used in situations when the decentralized and social planner solutions differ.

Second, my model extends non-stationary dynamic competitive models (in the Hopenhayn (1992) style) by adding multiple technologies and an hourly spot market. Similarly to Costantini and Melitz (2008), the non-stationary transition is key to my model and I extend their computational algorithm to solve the model. Adding multiple technologies expands the dimensions of the search space but is otherwise relatively straightforward. My spot market model replicates market outcomes closely while also significantly decreasing computational time. The key observation is that it is enough to focus on 50 representative hours that cover demand variation well. Then, once outages are accounted for, a fully competitive capacity constrained Bertrand model performs relatively well.

Finally, I contribute to the literature on investments in power generation and environmental regulation. In particular, my paper fits in the literature of estimating the welfare costs of second-best policies. Most of the literature uses two-period models (Fischer and Newell (2008)) which limit the range of comparable policy options or long-term planning models (Murray et al. (2014), Fell and Linn (2013)). Linn and

[4] EIA et al. (2016) contains a detailed analysis of the expected impact of the Clean Power Plan

McCormack (2019) finds that recent coal production decline and exit decisions were driven mostly by gas price decrease and not environmental regulation. Bushnell et al. (2017) focuses on the incentives to choose mass-based or rate-based standards under the Clean Power Plan. Abito et al. (2019) uses empirical methods by Bajari, Benkard and Levin (2007) to study gas power plant entry and the Clean Power Plan in PJM. Aldy, Gerarden and Sweeney (2018) and De Groote and Verboven (2019) empirically study investment subsidies for wind and solar respectively. Most of these papers focus on final equilibrium outcomes. My contribution is to focus on the entire transition path and directly compare how different tools influence the speed of transition. My quantitative comparisons are novel as they include differences through the entire transition path.

3.2 Industry Background

I study wholesale electricity markets in the United States. Power generation (wholesale) is the production of electricity. Electricity is physically delivered to end users through transmission and distribution lines. Households are connected to distribution lines and purchase their electricity from retail companies. Historically, the industry was organized as a set of utilities; fully vertically integrated regulated monopolies for distinct geographic regions. Operating electricity infrastructure (transmission and distribution lines) is a natural monopoly. In contrast, wholesale and retail markets are not necessarily natural monopolies. A long deregulatory process started in the 1990s with liberalizing wholesale electricity markets[5]. Power generation markets opened up for competition. Supply and demand of electricity needs to be balanced in every time period (there is very limited storage). Lack of a single producer, a centralized entity is necessary to ensure this balance. In liberalized electricity markets this role is fulfilled by independent system operators (ISOs). In addition to operating

[5] The process stopped mainly due to the California electricity crises in the early 2000s

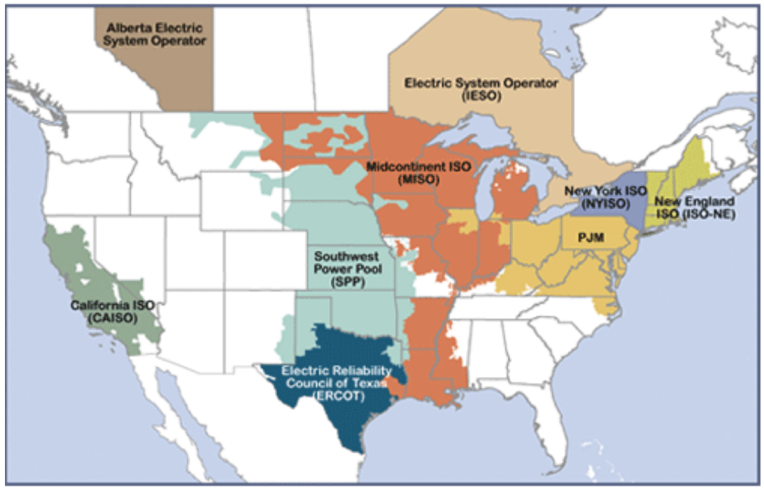

Source: FERC.

Figure 3.1: Liberalized Wholesale Electricity Markets in the US

wholesale markets for electricity, ISOs also operate transmission infrastructure[6]. The distribution network remained with traditional utilities. In most states the retail market remained regulated. Figure 3.1 shows the current state of liberalized wholesale markets. Several states and smaller regions remained under the traditional regulatory framework. Currently, there are six fully competitive wholesale electricity markets in the United States[7]. My analysis uses data from PJM, the biggest ISO in the US.

3.2.1 Supply and Demand for Electricity in PJM

Demand for electricity varies significantly both across and within days. Figure 3.2 presents the daily load profile of PJM and the within day variaton for a representative day. Demand is also very inelastic. Households are usually not very responsive to changes in pries. Households usually pay a fixed price for electricity that does not vary by the hour. Even when it does, it is questionable whether households respond to it. The literature estimates price elasticities for real-time pricing programs in the range of -0.04 to -0.15 (Allcott (2012*a*)). However, Ito (2014) finds that household respond

[6]ISOs are natural monopolies and regulated on a cost plus basis

[7]SPP is a power pool that resembles competitive markets.

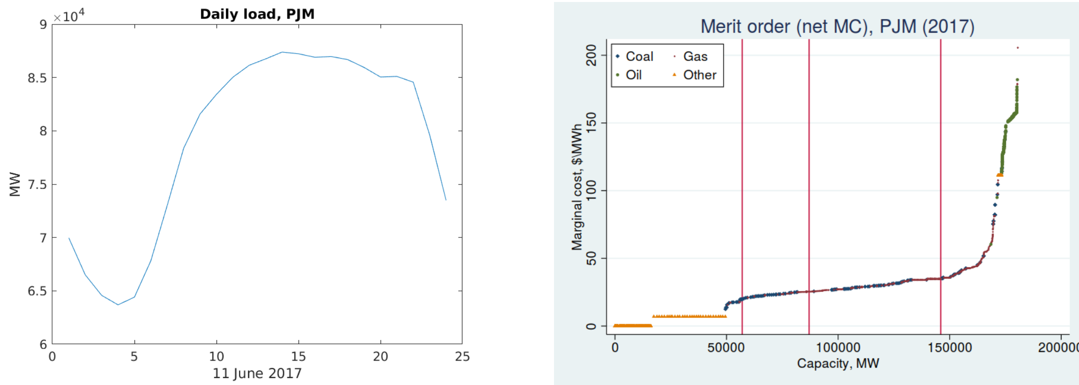

(a) Daily load

(b) Merit order

Figure 3.2: Demand and Supply in Power Generation, PJM (2017)

to the average price of electricity. Even this response is very inelastic[8]. Demand variation combined with inelasticity puts significant burden on the system operator. Although it is mostly predictable, last minute changes do occur. Because demand is inelastic, supply has to be flexible enough to be able to satisfy this variation. A prime example is the morning rampup that happens from around 5 to 10am. Power plants need to increase output or need be started to satisfy this increase in demand. The system needs technologies that are able to start flexibly and quickly.

Electricity supply is characterized by a set of technologies. In this paper I classify technologies to seven distinct groups based on primary fuel input and plant type. The categories are coal, gas, hydro, nuclear, oil, solar and wind[9]. First, technologies can be categorized as dispatchable and non-dispatchable[10]. Dispatchable technologies are able to adjust their production. In contrast, wind or solar power plants are non-

[8] Another way to see this is to estimate the value of lost load (essentially reservation price). Frayer, Keane and Ng (2013) estimates VOLL around 6000$\$/MWh$.

[9] There are additional technologies that I do not include in the discussion: biomass, biogas, geothermal and others. Some plants can use multiple fuels. The two common power plant types when this happens are turbines that can operate with oil and gas or biomass and coal. This specification is just for modeling and descriptive purposes. In the data I have the heat content of the actual fuel used and also the thermal efficiency of the plant. This allows me to precisely estimate the marginal cost of each generator.

[10] Dispatch refers to the system operators instructions to produce.

dispatchable as their production solely depends on weather[11]. Second, technologies can be categorized based on flexibility. Nuclear, coal and combined-cycle gas turbines (CCGT) are slow to adjust production[12]. They are usually considered baseload plant as they historically produced almost always at a lost cost serving baseload. In contrast, conventional combustion turbines (fueled by gas, oil or both) are able to react quickly but are relatively expensive. Traditionally, these plants served the peak load. Hydroelectric power plants vary greatly in flexibility[13].

Figure 3.2b presents the merit order for PJM in 2017. Power plants are ordered by their marginal costs resulting in a supply curve for the industry. Solar, wind and nuclear plants have very low marginal costs and as such, these technologies are on the left of the supply curve. In PJM, coal and gas produce in a similar range of marginal costs. Inside each category, specific power plants differ significantly based on the efficiency and the heat rate of the fuel used. Finally, oil power plants produce at the highest marginal cost levels. Technologies also differ in fixed costs such as construction and yearly operational costs. Generally, there is a negative trade-off between marginal and fixed costs. Additionally, high marginal costs plants need to be flexible to be able to produce.

The system operator has to ensure that supply and demand balances in every instant. Any imbalance can result in drops or raises in system frequency that can result in damaged equipment and in severe cases blackouts. In liberalized US wholesale markets this is achieved by a sequence of two-sided uniform price multiunit auctions. Power plant operators submit quantity-price pairs that describe the quantity they are willing to supply at a given price. Similarly, demand submits demand schedules.

[11] Or at least they can not increase their production.

[12] More precisely, they are capable of very small adjustment occasionally (required for frequency regulation) but can not ramp up their production quickly.

[13] Conventional hydro plants have reservoirs that allows to control water flows to produce when the price of electricity is the highest. Of course, this is limited by earlier water inflows and the capacity of the reservoir. Pumped-storage hydro plants are essentially storage tools. Run of river plants have no reservoirs and hence no control over production.

The system operator aggregates these offers to build industry supply and demand curves. Essentially, the intersection of the two curves determines prices, quantities, production and consumption. In reality, markets clear with a complicated algorithm that also considers constraints. These constraints come from generator characteristics (capacity constraints, ramping rates, etc) and system constraints (transmission). In addition, participants are allowed to trade bilaterally and centralized futures markets also exist. Both are financial markets only, physical delivery is completely determined by the auctions. Most ISOs also use tools to mitigate market power. Every US ISO has price caps that limit the maximum offer price. Moreover, most ISOs apply "offer" capping to mitigate local market power (due to transmission congestion). The ISO can change the offer of a power plant to marginal cost plus a small (e.g. 10%) margin. Although these caps do not necessarily bind they likely deter firms from submitting bids that are far away from true marginal costs. Offer capping makes these markets more competitive by mitigating market power due to congestion.

Finally, the system operator also has a responsibility for the security and stability of the electricity system on the long run. ISOs usually plan to be able to match maximum demand plus a reserve margin. Reserve margins meant to provide an insurance against unexpected power plant shut downs, congestion or changes in demand. The socially optimal level of reserve margins balances the cost of additional capacity to the benefits of less frequent shortages. PJM supplements the energy market with ancillary services markets for short-run security of supply and capacity markets for the long-run. See Appendix B.1 for details.

3.2.2 Relevant Market Trends

There are two large changes in the industry during the time of the data. First, subsidized renewable energy started to enter US power markets. Subsidies were considered necessary to induce innovation and drive down the cost of these technologies. Wind and solar are both high fixed costs zero marginal cost technologies that result in lower

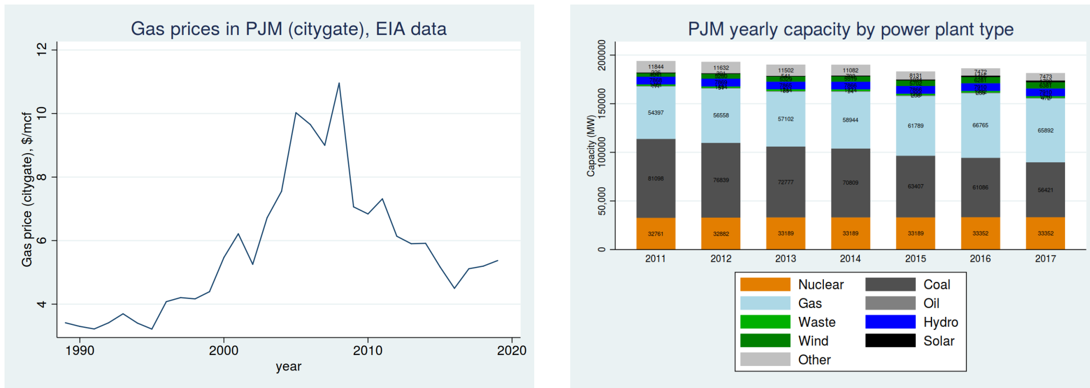

(a) Gas prices (2005-2017)

(b) Generation mix (2011-2017)

Notes: data source is EIA Form 860 using PUDL classifications.

Figure 3.3: Gas Prices and Generation Mix, PJM (2004-2017)

prices on average. Second, the shale gas revolution in the United States led to a drop in current and long-run expected natural gas prices. In the meantime, coal prices were essentially unchanged. The business case of gas-fired plants improved. As their major substitute, this hurt coal the worst. The key change in PJM was a slow entry of gas and exit of coal power plants that started around 2011 and coincided in timing with decreasing gas prices. Renewable generation was less influential due to geographical characteristics and the lack of extensive subsidies. Since gas power plants do set the market clearing price relatively frequently, average prices also decreased across the US and in PJM as well.

3.3 Model

I introduce a non-stationary dynamic competitive entry-exit model that is motivated by the characteristics of the power generation market in PJM. The model is designed to answer questions of the power generation market for which transition is of primary importance. There are two key differences to standard entry-exit models[14]. First, multiple technologies exist. I allow for entry of the newly competitive (gas) and exit of the less competitive (coal) technology because these are the institutional realities of

[14]The entry and exit dynamics in my model are similar to Hopenhayn (1992) but there is no firm level persistent heterogeneity[15]. Instead, firms only draw an i.i.d. scrap value.

PJM today. I assume away entry of any other technology and exit of other technologies currently in the market. Second, power generators compete in hourly organized spot markets. This is essential as hourly demand varies substantially. High marginal cost technologies set prices at high levels of demand. This generates rents for technologies with lower marginal costs that help finance their higher entry and fixed costs. High marginal cost technologies earn rents from scarcity of supply. To solve the non-stationary equilibrium I use a computation method similar to Costantini and Melitz (2008). I extend this method for the case of two distinct technologies.

Both coal and gas power plants emit carbon dioxide and this negative externality is the only friction in the model. Environmental regulation influences the speed of technology transition because technologies differ in their environmental impact. I consider a carbon price and various alternative policy options that are easy to implement in the model. The transition to less polluting technology is slowed down by the existing fleet of capacity. Introducing a carbon price can both result in a different equilibrium and speed up the transition process. For instance, making current coal capacity exit to make way for gas expansion can increase welfare if this results in a faster drop of costs.

3.3.1 Model Setup

Firms make entry and exit decisions in every time period. Firms operate power plants with different technologies some of which might enter and exit. I call coal (c) and gas (g) endogenous as these are the only technologies that can exit and enter. The two technologies have similar marginal costs ($c_c \sim c_g$) but coal has higher entry ($CE_c > CE_g$) and operating costs ($F_b > F_p$)[16]. Coal plants can also expect higher scrap values[17] Both technologies are dispatchable and can change production levels

[16]I do not have to assume only gas enters but since this is clearly the solution to my calibrated model I assume away potential coal entry

[17]I do not consider heterogeneity in costs. Every endogenous coal plant has the same parameters. Similarly, every entering gas plant is the same. This assumes away differences in efficiency.

immediately. Capacities are denoted by K_c and K_g. Firms are infinitesimal. Competition is essentially between MWs of technology. Capacities of other technologies are given exogenously. There is an exogenous state variable z_t with a known path that is observed in the beginning of the time period 0. z_t contains fuel prices, demand and the capacities of exogenous technologies. The endogenous state variables are capacities of each technology $K_t = (K_t^c, K_t^g)$. The timing in each period is the following:

1. Entry decisions are made at the beginning of the period. Firms pay entry costs. Entrants become incumbents.

2. Incumbents observe scrap values F_t^{ex} and decide whether to exit. If exit, receive scrap values immediately.

3. Incumbent firms compete on the hourly spot markets (knowing the sequence of exogenous hourly states).

4. The industry states update and next period starts.

3.3.2 Spot Market

Let θ denote demand distributed with cdf $F(\cdot)$ with support normalized to $[0, \bar{\theta}]$. Demand is perfectly inelastic. In the model, firms behave as price takers on the spot market given the capacities on the market. Knowing endogenous capacities (K_c, K_g), exogenous capacities and marginal costs (c_c, c_g) is sufficient to calculate market outcomes, prices and profits. Market clearing is simplified. First, I calculate the merit order (industry supply curve) given marginal costs and capacities. Second, the intersection of supply and demand determines prices for every hour. $\{P_t^h\}$ denotes the set of hourly prices in time period t. If a firm's marginal cost is below the market clearing price it uses its entire capacity to produce. The firm that sets price produces the satisfy the remaining load. Yearly profits are calculated as the sum of hourly profits. The spot market model assumes away congestion, market power, the interdependency of hours due to starting costs and additional profits in the ancillary services markets.

3.3.3 *Dynamics and Equilibrium*

In each period firms maximize value by deciding whether to continue or exit. Their choices satisfy the following Bellman equation:

$$V_t(K_t, z_t, F_t^{ex}, a) = \max_{e,C}\{\Pi(K_t, z_t, a) + \beta\mathbb{E}[V_{t+1}(K_{t+1}, z_{t+1}, F_{t+1}^{ex}, a)], F_t^{ex}\}, \qquad (3.1)$$

where a denotes the technology of the incumbent and $\Pi(K_t, z_t, a)$ denotes spot market profits. The expectation is over F^{ex}, the scrap values. I let F_{ex} distributed as an exponential random variable with parameter λ and cdf $G(\cdot)$. I use the exponential distribution as it has a full support and therefore guarantees that there is always exit and entry. I introduce the expected value function as the value of the firm before observing its scrap value:

$$EV_t(K_t, z_t, a) = \mathbb{E}(V_t(K_t, z_t, F_t^{ex}, a) = \int_{F^{ex}} V_t(K_t, z_t, F_t^{ex}, a)dG.$$

Exit strategies are characterized by an exit threshold. We can solve for the expected value function analytically:

$$\begin{aligned} EV_t(K_t, z_t, a) &= \Pi(K_t, z_t, a) + \beta EV_{t+1}(K_{t+1}, z_{t+1}, a) \\ &\quad + \lambda[1 - G(\Pi(K_t, z_t, a) + \beta EV_{t+1}(K_{t+1}, z_{t+1}, a)] \end{aligned}$$

This expression allows me to calculate the stationary equilibrium by value function iteration on the expected value function. Similarly, setting EV_T at the long run new stationary level allows to calculate the path of EV_t by backward induction.

Entrants can only enter with technology gas or coal by paying the entry cost. Entry is realized immediately and entrants become incumbents from that period on. Entrants are ex-ante identical. Entry decisions are made before observing scrap values. Entrants enter if the value of entry is positive:

$$V_t^E(K_t, z_t, a) = EV_t(K_t, z_t, a)) - CE_a$$

A dynamic equilibrium is characterized by capacities $\{K_t^c, K_t^g\}$, the price path $\{P_t^h\}$ and exit thresholds. An equilibrium have to follow the following conditions:

1. **Firms maximize value.** Firms exit decisions satisfy equation (3.1) (anticipating entry and decisions of all other incumbents). These choices lead to capacities $\{K_t^c, K_t^g\}$ and price path $\{P_t^h\}$ that are consistent with firms choices.

2. **Free entry.** In equilibrium $V_t^E(K_t, z_t, a)$ must be non-positive for both technologies in all time periods. Entry must be zero if $V_t^E(K_t, z_t, a)$ is negative.

The equilibrium is non-stationary and finite horizon. Firms take capacities as given but their overall behavior determines capacities. I assume that after period T firms receive the stationary payoffs. Due to the path of exogenous states I assume only gas can enter. Solving the model is non-trivial, since the solution space is a vector of capacities. Hourly market clearing complicates calculating yearly profits and increases computational time. I solve the model using a computational algorithm similar to Costantini and Melitz (2008). I modify the algorithm to accomodate multiple technologies. The algorithm is described in detail in Appendix B.2.

3.4 Empirical Approach

In this section, I describe my approach to applying the model to the practical application. First, I discuss the data. Second, I describe my approach to estimating marginal costs and implementing spot market clearing. Then, I explain how I calibrate key parameters that drive dynamic decisions. Finally, I discuss how I model environmental regulation.

3.4.1 Data

I use publicly available data for US power markets from 1990 to 2018. EIA (Energy Information Administration) collects data on power generator characteristics in Form 860 and Form 861. These include the owner, capacity, technology, location, operating status, first and last year of operation. The data allows me to reconstruct the generation portfolio of PJM every year. EIA Form 923 (and previously 906/920) collect detailed monthly generator operations data. This includes fuel receipts and generator output that allows for the estimation of the heat rate (efficiency) and fuel input cost of the generator which are key ingredients to estimating the marginal cost. Unfortunately fuel contract data currently is only available for regulated plants[18]. The data allows me to produce reasonable marginal cost estimates for every power plant in every year. Capital cost estimates are also widely available. EIA published capital cost estimates in 2010 and 2017 for every power plant technology. PJM publishes capital cost estimates in its annual market reports[19]. The data is not as easily available for before 2010. Additionally, PJM publishes net revenue estimates for new gas power plants for every year which is useful to estimate spot market revenues. Finally, PJM publishes hourly system load and prices.

3.4.2 Spot Market

I estimate marginal costs for each power plant in PJM using data from the EIA. I use gross marginal costs that include fuel costs, variable fixed costs and SO_x costs. EIA provides yearly variable fixed costs estimates for every technology[20]. SO_x costs

---

[18] EIA made the entire data available upon request before 2017. The key implication for my work is that I have to extrapolate fuel costs extensively based on similar plants with available data.

[19] Table 7-23 in 2017 Volume 2 contains gas power plant cost estimates, Table 6-15 in 2012.

[20] These include maintenance and costs associated with raw water treatment, waste treatment and catalysts.

come from SO_x emission rates and yearly market prices as reported by the EIA. When I implement carbon prices, I use carbon emission rates and add carbon costs to gross marginal costs. I correct for outages at the plant level. I use outage rates as reported by PJM: 5% for gas, 15% for coal. I assume nuclear, wind and hydro output is based on exogenously given capacity factors. Scheduled maintenance is exogenous for nuclear plants and weather conditions determine output for wind and hydro. In the dynamic model I assume coal is homogenous to simplify exit decisions and solving the model. Therefore, I have to make a choice on the characteristics of the representative coal plant. I calibrate the heat rate of the representative coal plant to match historical hourly prices in PJM in 2011-2017. I select 50 hours by a k-means clustering algorithm to represent all hours of the year. I find that the using these 50 hours results in essentially the same results while reducing the computational burden of the dynamic model.

Figure 3.4 reports the key result of the spot market exercise. Generally, the figure provides confidence in my way of modeling the spot market. On average, replicated prices are relatively similar to market prices in every year from 2011 to 2017. The homogenous coal scenario is picked so as to match average prices so its good performance is not surprising.

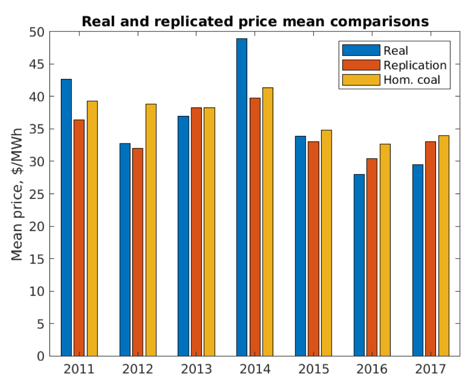

Notes: homogenous coal is the scenario when I replace all coal plants with a representative technology. The representative heat rate is chosen to match average prices.

Figure 3.4: Mean Predicted and Real Prices, PJM (2011-2017)

Although my spot market model matches average prices, it struggles to generate realistic variation in prices. Replicated prices are less dispersed than prices in the data (Appendix B.4). The homogenous coal scenario does not generate price variation in 2017. However, in the long run, there is price variation even in this scenario. This is not surprising. Real-world dispatch takes into consideration ramping, congestion and market power which are absent in my model. As a result, my model yearly profits for high marginal costs plants and overestimates for low marginal cost plants (see Appendix B.4). However, in the relevant price range for coal and gas plants, my estimated profits are not far from data.

3.4.3 Calibration and Estimation

I calibrate most of the parameters of the model using publicly available data from PJM and the EIA. I also make parameter choices to fit the data best. Table 3.1 summarizes the calibrated and estimated parameters. I assume demand is the observed historical quantity in 2017 and that it stays the same after 2017. This is not unreasonable, PJM's forecasts demand not to grow mostly due to energy efficiency improvements balancing the need for more energy. I calibrate entry and variable costs using data from the EIA. I estimate marginal costs for all power plants based on existing heat rates and fuel prices. I assume fuel prices are as of 2017 and do not change. I set the discount rate β to 0.95 which is a common choice. I assume a time period is a year and I set the number of years to 39. I set the price cap to \$400/MWh and the value of lost load to \$1000/MWh[21]. When implementing the carbon price, I use the Obama administration's social cost of carbon estimate of \$30/$tonCO_2$. This is on the lower end of the current academic consensus. I implement the carbon tax as a factor that increases the marginal cost of production. I do not assume a pass-through of carbon taxes to prices, this is determined by hourly spot market competition. I choose the heat rate for the entering new gas power plants based on the best available

[21] These choices are not crucial, my results are not sensitive to these. The price cap is almost irrelevant as it never binds, while the VOLL estimate is if anything on the lower end.

technology in EIA's 2010 estimate. Gas power plants did not improve substantially recently, and the consensus is that they are unlikely to improve further in the future.

Table 3.1: Parameters and Exogenous Variables Used

Notation	**Measure**	**Value**	**Source**
F	Load cdf	K-means(50), 2017	PJM
Nucl, hyd	MWh	Capacity factors*MW	EIA
F_c, F_g	\$/MW/year	50000; 10000	EIA
CE_c, CE_g	m\$/MW/year	3.6, 0.978	EIA
Fossil cap., mc	MW, \$/MWh	Author estimates	EIA
O& M mc	\$/MWh	3.5, 5	EIA
Outage	%	5% gas, 15% coal	PJM
β		0.95	
T	Periods	39	
$p_{cap}, VOLL$	\$/MWh	400, 1000	
scc	\$/t	30	EIA (Obama)
Scrap mean	m\$	0.425, 0.001	Reasonable
Coal heat rate	mbtu/MWh	11.1	Fit
Gas heat rate	mbtu/MWh	7: Best available	EIA (2010)
Fuel prices		After 2017=2017	
Entry subs.	m\$	0.09	Same long-run EQ
Exit subs.	m\$	0.6	Reasonable

Notes: the first value is always for coal and the second is for gas. Scc: social cost of carbon. Parameters are set relatively close to real values but also to illustrate the mechanisms of the model. Fit: picked to fit data. Reasonable: picked for reasonable comparisons for different scenarios.

I start my dynamic simulation in 2011. Since data is available until 2017, I can use the time periods of 2011 to 2017 to compare my model to real world outcomes. I pick the coal scrap value so that the baseline scenario leads to almost all exiting in the 39 years of my modeling horizon. I set the scrap value for gas at 0.001 so that I do not incentive any gas exits for capturing scrap values. As described in Section 3.4.2, I pick the homogenous coal heat rate to match average spot prices. I choose gas entry subsidies to match stationary equilibrium outcomes for the calibrated carbon price. Finally, I choose the coal exit subsidy and the high gas entry subsidy to incentivize faster and relatively comparable transition. This allows me to compare transition dynamics of different environmental policies that result in similar long term outcomes. The subsidies are a transfer from the government to the firms. I assume away any frictions of raising government revenue.

3.4.4 *Welfare*

Welfare in this model is the total cost to society of satisfying consumer demand. I calculate it using the following formula:

$$W_t = -ConsumerC_t + \Pi_t - EC_t + SC_t - VOLL_t - CarbonC_t + Gov_t,$$

where $ConsumerC_t$ is yearly consumer spending on electricity, Π_t denotes yearly spot profits, EC_t is entry costs, SC_t is the sum of scrap values, $VOLL_t$ denotes the value of loss load, $CarbonC_t$ is the cost of carbon emissions for society and Gov_t is government balance from the power generation market. I consider scrap values as increasing welfare as the literal interpretation of these is what firms can recover from their investment. My entry cost and scrap value calculation are simplified. I only consider the total change in capacity[22]. When total generation is below total demand, I use the value of loss load estimate to calculate the harm to consumers. I calculate carbon costs using the social cost of carbon. Finally, government revenues include carbon tax revenues, entry subsidies and exit subsidies[23]. Appendix B.5 describes the details of how I calculate the building blocks of welfare.

3.5 Results

3.5.1 *Stationary Equilibrium*

The model has a stationary (or long-run) equilibrium that the non-stationary equilibrium converges to. Table 3.2 describes the key characteristics of the stationary equilibrium. Column 1 represents the baseline scenario. In the long run, all of coal capacity exits and is replaced by gas power plants. The reason is simple, gas is superior to coal in the model. Gas is cheaper to build and maintain and with current

[22]In the model, there is gas exit in gas even when gas capacities are increasing. This results in higher scrap values and entry costs. I assume this channel away from welfare calculations.

[23]Notice that carbon taxes can generate higher cost for consumers and profits for companies in this market with heterogenous technology.

fuel prices it has lower marginal cost as well. The carbon price scenario (column 2) has a very similar long-run equilibrium. Spot market prices are higher (pass-through is perfect) but the gas capacities are only slightly higher. This is because the profitability of gas power plants in the stationary equilibrium is not effected much by the carbon tax. Gas is most of the time the marginal plant in the stationary equilibrium and hence its carbon cost gets perfectly passed through to prices. An alternative way of looking at it is that since carbon prices are low, they do not influence outcomes too much. The coal exit subsidy (column 3) can not influence the long-run equilibrium level of gas capacities. Since coal exits in any case, exit subsidies do not influence long-run market outcomes. In column 4, I set the gas entry subsidy to induce the same capacities as the carbon price. Long-term welfare is very similar across all scenarios. Capacities are essentially the same across all scenarios. The impact of the carbon price is limited, since on the long-run spot market it barely generates any reorganization of production. Simplified welfare calculations show that welfare differences are small[24]. These results highlight the limitations of looking at only the stationary equilibrium. Differences across policies are minor. However, it is possible that the same policy tools induce vastly different transition paths. This is what I explore in the next section.

3.5.2 Non-Stationary Equilibrium and Welfare

In this section, I discuss the key results of this paper. The solution to my non-stationary model is mainly characterized by the path of coal and gas capacities. Figure 3.5 presents these outcomes for the baseline, carbon tax, coal exit subsidy and two gas entry subsidy scenarios. With the calibrated parameters, coal completely

[24] I assume away entry costs and scrap values on the stationary path. Even though capacity does not change, there is entry and exit in the model. Some firms get high scrap values resulting in exit. Firms enter so that capacity remain unchanged. This would cause only a small difference without differences in capacity, entry and scrap values. However, the difference can be significant with high entry subsidies making comparisons uninformative. Therefore, I do not report the high gas entry subsidy scenario.

Table 3.2: Stationary EQ Under Different Policies

	Baseline	Carbon price	Coal exit subsidy	Gas entry subs. low
Descriptives				
New Gas K (GW)	50.91	51.17	50.91	51.17
Coal K (GW)	0	0	0	0
Mean price ($/MWh)	35.36	46.48	35.36	34.83
Simplified welfare				
Static welfare (b$)		-0.053	0	-0.053
Entry cost (b$)		+0.84	0	+0.84
Total welfare (b$)		-0.21	0	-0.21

Notes: Gas subsidies are set to induce gas capacities as with carbon tax. The equivalent entry subsidy is m$0.09 . Exit costs not included in equilibrium, gas scrap values are low but not 0. Higher entry costs mean lower welfare. Welfare results are simplified. Total welfare is calculated as the discounted future payoff from static welfare minus the entry cost in the initial period. Only the difference to the basic scenario is reported. Stationary exit-entry assumed away.

exits in 40 time periods (years) even in the baseline scenario[25]. Similarly, the transition to gas capacities in the stationary equilibrium is complete in 40 years in the baseline scenario. All other scenarios lead to faster transition towards the stationary equilibrium. Moreover, scenarios lead to vastly different transition paths. A carbon tax of $\$30/tCO_2$ induces immediate gas power plant entry that is close to the long-run equilibrium. Carbon tax immediately increases prices (see Appendix B.6 for details) on the market incentivizing entry[26]. It also creates a big cost advantage for gas power plants and results in coal production close to zero. The gas subsidy that results in the same stationary equilibrium as the carbon tax does not differ substantially from the baseline scenario. Introducing a higher gas entry subsidy does lead to higher gas capacities on the long run and a faster but not immediate transition. Neither gas subsidies changes coal exit speed fundamentally from the baseline. Finally, coal exit

[25]The baseline scenario results in somewhat faster early entry and exit than what I observe in the data. This is mainly a result of high coal scrap values.

[26]The carbon price I consider is on the lower range of current academic consensus. Lower values can lead to less immediate entry. Higher values would produce similar outcomes.

subsidies can also speed up transition by providing incentives for coal plants to exit earlier. This in turn, also leads to faster gas entry. Exits can be faster than in the carbon tax case, but not entry. Coal plants in the market hinder the ability of gas to enter as long as they are relatively competitive. The impact of the carbon tax is radical because it makes coal immediately redundant for competition. In line with results on generation, emissions are decreasing almost immediately in the carbon tax scenario and more smoothly in others. My simulations generate lower emissions than the goals of the Clean Power Plan of the Obama administration (see Appendix B.6 for details). Changing the levels subsidies influence the speed of transition, but not the general characteristics. Subsidies result in smooth transitions since they do not fundamentally alter the spot market merit order[27]. It is clear from these comparisons though, that subsidies that result in similar stationary equilibrium can not generate gas entry that is comparable to a high carbon tax. Welfare comparisons are necessary to examine the attractiveness of each scenario further.

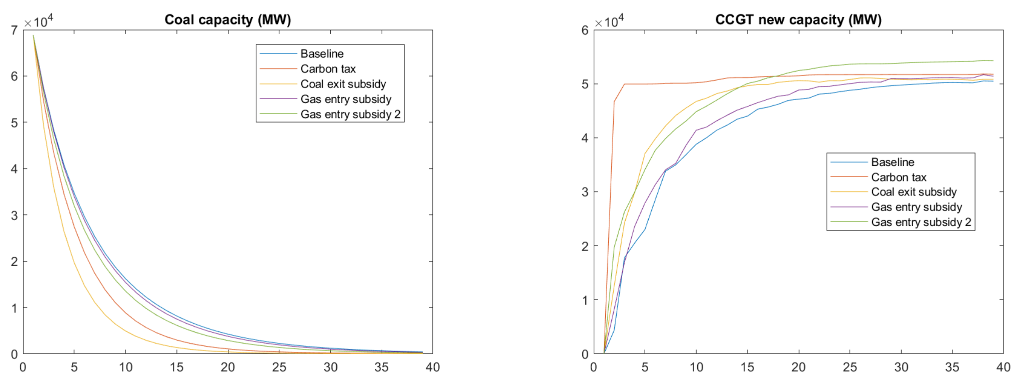

Notes: baseline calibrated parameters as described in Section 3.4.3 and Table 3.1. Carbon tax is $30/tCO_2$. Gas entry subsidy 1 set to induce the same stationary equilibrium as the carbon tax. Coal exit subsidy and gas entry subsidy 2 picked to induce relatively fast coal exit.

Figure 3.5: Coal and Gas Capacities in Different Regulatory Scenarios

Table 3.3 reports welfare results. Confirming theoretical predictions, carbon pricing at the social cost of carbon leads to the highest welfare across all scenarios.

[27]This is partially due to the functional form of coal exit scrap values. However, this is likely a good representation of heterogenous coal in the real world. Coal marginal costs are essentially continuous and they determine exit. The result is a relatively smooth transition.

However, consumers pay higher prices. This can be alleviated by redistributing the government revenue of the program. Notice that profits are higher in the carbon tax scenario because of the cost heterogeneity. Entry and exit subsidies improve on the baseline scenario by getting closer to the carbon price outcome. Since capacities in the small gas entry scenario are very close to baseline, it is not surprising that it increases welfare only moderately. The gains also remain low with substantial coal exit subsidies[28]. Although coal exits fast in this scenario, the remaining coal plants keep producing for long which limits the gains. The key drawback is that coal exit subsidies come at substantial costs for the government. Finally, high gas entry subsidies get closest to carbon price outcomes even though they induce too much capacity on the long run[29]. Both consumer and government costs are low making implementation relatively easy. Altogether these results confirm the welfare importance of transition. Tools that are close to efficient on the long run (low gas subsidies) perform poorly when considering transition dynamics.

[28]This is partially because I count scrap values as increasing welfare. Without this, coal exit subsidy scenario is closer to substantial gas entry subsidies. In general, scrap values do not make much of the difference between scenarios.

[29]It is likely that the optimal gas entry subsidy can not get very close to the carbon price for two reasons. First, subsidies are a dynamic instrument and do not improve on spot market efficiency. Second, it is unlikely that they can incentivize immediate entry at a reasonable long run cost of overcapacity. Finally, my assumption of calculating only capacity changes in welfare might help the high gas entry subsidy scenario. It is unlikely to change other scenarios significantly, since the probability of gas exit is very low in my preferred calibration without entry subsidies.

Table 3.3: Simplified Non-Stationary Welfare Under Different Policies

	Baseline	Carbon price	Coal exit subsidy	Gas entry subs. low	Gas entry subs. high
Consumer costs (b$)	-326	-433	-328	-320	-308
Profits (b$)	62	72	131	61	61
Scrap value (b$)	60	57	54	60	59
Entry costs (b$)	-35	-45	-39	-36	-40
Carbon costs (b$)	-126	-98	-114	-123	-116
Gov. (b$)	0	99	-65	-3	-12
Total welfare (b$)	**-365**	**-349**	**-362**	**-363**	**-357**

Notes: all lines include discounted ($\beta = 0.95$) totals. VOLL is 0 across all scenarios and therefore not reported. Only impact of net capacity changes included in welfare.

4

Do Players Bid Truthfully in VCG Auctions? - Evidence from the New England Frequency Regulation Market

4.1 Introduction

I analyze bidding behavior in the New England (NE) frequency regulation market. My interest in this market comes from the market design change in 2015. As a response to a regulatory intervention, the market design changed from a single dimensional uniform price auction to a multidimensional Vickrey-Clarke-Groves (VCG) auction. This setup is unique as there is virtually no observational evidence on how the VCG mechanism performs for multiunit auctions[1]. The VCG auction is attractive in this environment as it is truthful, efficient and robust to bid skewing, a general concern of scoring auctions. Despite the favorable theoretical predictions, I find indirect evidence that market participants do not bid their true marginal costs. Optimal strategy under VCG is very simple if bidders know their own marginal costs. In contrast, finding optimal bidding strategies in uniform price multi-unit auctions is difficult. Hortaçsu

[1]The topic is relevant today as, although in rankable multiobject environment, both Facebook and Google uses VCG auctions to sell online advertisements. For instance, Varian and Harris (2014) notes for Google about VCG: "From what we can tell, it seems to be working well".

and Puller (2008) propose VCG as an alternative to reduce strategic complexity in the Texas balancing energy market. My findings indicate that players were not able to learn optimal bidding most likely due to a combination of low stake and a complicated market design in which it is difficult to link bids to outcomes. Altogether, my results suggest that using VCG bids as estimates of true marginal costs can be misleading.

Vickrey auctions have a long history in market design. Vickrey (1961) proposed a single unit auction design in which the winning player pays the highest losing bid[2]. This design is attractive for two main reasons. First, players' optimal strategy is to bid truthfully their valuations. Second, the resulting allocation is efficient: the player with the highest valuation receives the good. However, the extension to multiunit auction environments is not obvious. Consider uniform price auctions for multiple identical objects. Players bid quantity-price pairs. With n objects to sell the n highest unit bid receives an object. The price is determined by the highest losing bid. Then, each player pays this price for each unit allocated. Truthful reporting is not optimal when players have any market power (Wilson (1979)) and the uniform price auction is almost never efficient (Ausubel et al. (2014)). In contrast, Vickrey auctions are always efficient and truthful. In VCG auctions for n identical objects allocation is determined by the n highest bids. Then, the mth highest winning bid pays the mth highest losing bid that did not belong to the same bidder. This way bidders are only interested in maximizing surplus which results in truthful reporting and an efficient allocation.

Despite some favorable theoretical properties VCG was rarely applied to multiu-

[2]Lucking-Reiley (2000) notes that economists attributed the invention of the second price sealed bid auction to Vickrey. However, the design was used at least as early as in 19th century for stamp auctions.

nit auctions[3]. There are several theoretical and practical issues for implementation[4]. Two of these question whether players would bid truthfully in practice. First, repeated interactions can create incentives not to reveal true valuations. Second, VCG can be complicated for market participants to fully understand. Therefore, empirical evidence is needed to test the truthfulness property. The experimental auction literature established early that players overbid in second price auctions (Kagel (1995)). Experimental studies generally find that players overbid under multiunit VCG auctions as well (see Kagel and Levin (2016) for a summary). The usual limitations of experiments apply. The stakes are rarely high[5]. Additionally, the limited amount of interactions hinder participants ability to learn optimal behavior. In any case, observational evidence is clearly needed to complement these findings.

To test whether players bid truthfully in a multiunit VCG auction I analyze bidding in the NE frequency regulation market. Frequency regulation is necessary in every electricity system. Small differences between supply and demand can result in drops or rises in system frequency which is costly for consumers. These small differences are balanced out by fast ramping power plants that are able to react (produce more/less energy) in a very short time frame. In 2011, the Federal Energy Regulatory Commission (FERC) issued Order 755 that substantially changed the market design. Order 755 required independent system operators (ISOs) to introduce a new

[3]There is an extensive literature on simultaneous ascending-bid auctions for rankable multiobject and combinatorial auctions. These are practical implementations of the VCG mechanism for multiple, non-identical objects (see Milgrom (2000) and Ausubel (2004) for proving outcome equivalence to VCG.). The combinatorical clock auction (CCA, as proposed by Ausubel (2004)) received substantial attention especially for spectrum auctions. However, bidding behavior is difficult to empirically study due to the low number of observations.

[4]See Ausubel and Milgrom (2004) and Rothkopf (2007) for an exhaustive list. Here I list a few. First, the mechanism might result in low revenue for the auctioneer. Second, a well-known feature of VCG auctions is that players with higher market power receive higher payments. The two features can be exaggerated by non-convexities in energy auctions. Hobbs et al. (2000) use data for the PJM electricity market in an hour with relatively high demand to simulate the outcomes of the VCG auction. They find that this difference can be significant: the player with the strongest market power gets around twice the average payment per MWh than the weakest player. Hobbs et al. (2000) also find that average price suppliers are paid in a VCG auction can be around 50% higher than competitive prices. Fourth, the mechanism is particularly vulnerable to collusion.

[5]A critique somewhat less relevant for field experiments such as List and Lucking-Reiley (2000)

dimension for bidding and compensation that relates to the actual service provided. As a response, the New England ISO implemented a VCG auction. Power plants submit a capacity and a service bid that can be associated with fixed and variable costs respectively. The frequency regulation auction in NE is organized and cleared jointly with the energy auction. The auction takes place every day as 24 separate hourly auctions. I observe all of these bids but not outcomes. Although information is available to replicate outcomes it is unfeasible in practice[6]. Moreover, simplifications are likely to distort the simulated results at the frequency regulation market even if the energy market is replicated to high precision. Therefore, I focus on how power plants change their bids through time.

I find that the fixed cost component of bids varies with indirect measures of market power. First, bids are higher when demand is higher. Second, bids are higher when energy prices are higher. Fuel prices and energy demand determine the relative position of power plants in the energy market that has implications for availability and pricing in the frequency regulation market. I use energy prices as a proxy for outcomes in the energy market. Both of these factors should influence the fixed cost component of bids under a uniform price but should not under a VCG auction. Then, I look at the behavior of individual participants. I categorize individual bidders after visually investigating bid-time diagrams. I find significant heterogeneity across groups in the frequency and likely drivers of bid chances. Group frequency changes bids very frequently following energy prices and demand changes. Group seasonal changes with a seasonal pattern following demand. Group occasional changes bids rarely but mostly when energy prices are high. Finally, more than half of the participants never change their bids. Based on my results, it is likely that most of these players had only a small chance of winning in the auction. These findings suggest heterogeneity in bidding strategies and sophistication. I argue that the most likely explanation is that

[6] Even ISOs only estimate the solutions using supercomputers and highly tailored optimization algorithms.

learning fails to deliver optimal bidding for two reasons. First, learning is difficult as it is difficult to link bids to outcomes due to the complicated market clearing. Second, the stakes are low that limits player's willingness to experiment. I consider alternative explanations such as varying marginal costs, multiple plant ownership and the hypothesis that the market is still under an experimental bidding phase. I provide evidence against each of these explanations.

Several papers in the experimental auctions literature directly test the truthfulness property of VCG in a multiunit auction setting[7]. These studies generally find that players bid above their valuations (Porter and Vragov (2006), Engelmann and Grimm (2009). However, none of these experiments studies procurement auctions or scoring auctions with a VCG design.

My paper is also related to the literature on scoring auctions and in particular unit-price auctions. Asker and Cantillon (2008) shows general conditions under which even though bids are multidimensional only unidimensional pseudotypes influence outcomes. The two most closely related theoretical papers are Ewerhart and Fieseler (2003) and Chao and Wilson (2002). Ewerhart and Fieseler (2003) assumes two dimensions (materials and labor), a fixed scoring rule and heterogenity across players in the amount of labor to finish a project. A key feature of the equilibrium is bid-skewing[8]. Under a similar setup motivated by electricity reserves auctions, Chao and Wilson (2002) proves that the VCG design is truthful and efficient[9]. A key reason why ISO NE decided for the VCG design was its robustness against bid skewing.

There is substantial recent empirical work on multi-unit auctions (see Hortaçsu and McAdams (2018) for a summary). Hortaçsu and Puller (2008) studies the Texas

[7]Again, see Kagel and Levin (2016) for a review. There is also a literature on testing dynamic implementations of VCG for combinatorial auctions in the lab.

[8]Another concern is potential information asymmetries between bidders and the auctioneer, see Athey and Levin (2001) and Bolotnyy and Vasserman (2018) for common and private values respectively.

[9]There is recent interest in the design of electricity reserve auctions. In Europe, the European Commission proposes a multi-dimensional scoring auction with uniform prices. Ocker, Ehrhart and Belica (2018) compares the uniform and pay-as bid formats but not VCG.

balancing energy market. Large players followed optimal bidding closely. Although small players deviated substantially, the cost of optimizing bidding likely overweight potential additional profits. I attribute my findings to similar forces. Hortaçsu and McAdams (2010) recovers valuations for US treasury bonds from bidding under a uniform price auction[10]. They estimate a VCG counterfactual assuming truthful bidding to compare the efficiency of different auction formats. My results question the viability of this assumption but not the general results of the paper.

My paper also contributes to the empirical literature on auctions of electricity reserves and frequency regulation. Doraszelski, Lewis and Pakes (2018), the most closely related empirical paper, study the frequency regulation market in the UK after the introduction of a new market design (2003-2006). Although the product is essentially the same, the market design differs substantially from my setup. Regulation was procured through a pay-as-bid unidimensional auction with less frequent (monthly) interactions. Doraszelski, Lewis and Pakes (2018) observe outcomes which allows them to fit a structural model of bidding behavior. They use this model to study the convergence process to the new equilibrium. They find that after 40-50 periods of interaction market participants converged to an equilibrium best described by a differentiated price setting game. In contrast to their results, my findings indicate that market participants do not converge to an equilibrium predicted by theory. Other relevant recent empirical papers include Knittel and Metaxoglou (2008) (markups in California), Ocker, Ehrhart and Ott (2018) (bidding in reserve markets) and Metaxoglou and Smith (2007) (efficiency of reserve markets in California).

The rest of the chapter is organized as follows. In Section 4.2, I describe the industry, the frequency regulation product and the market design. Section 4.3 introduces theoretical models for the market design before and after the policy intervention.

[10]The literature have long recognized that if valuations are not flat, step-function bidding results in partial identification (such as in Hortaçsu and McAdams (2010)). In my setup, power plants can only bid flat supply functions. If marginal costs are also flat this feature significantly simplifies the analysis of bidding behavior

Section 4.4 describes the data and key descriptive statistics. In Section 4.5, I present and interpret my main empirical results. Section 4.6 concludes.

4.2 Industry and Market Background

This paper studies the frequency regulation market in New England, the region under the oversight of the ISO New England (NE). The ISO NE is the regional transmission organization[11] for the states of Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island and Vermont. The ISO NE is responsible for the reliability of the electricity system in New England. The New England electricity system with an average load of 14 GW in 2017 is among the smallest of the market based electricity systems[12].

4.2.1 The Frequency Regulation Market in New England

Frequency regulation is necessary in every electricity system as supply and demand needs to be equal all the time. Small differences between supply and demand can result in drops or rises in system frequency which is costly for consumers. These small differences are balanced out by fast ramping power plants that are able to produce more/less energy in a very short time frame (from milliseconds to seconds). In deregulated electricity systems, such as New England in the US, the system operator procures this service from power plant operators through a sequence of auctions. The product is an option for the system operator to call the power plant to increase/decrease output. Most dispatchable[13] power plants can offer this service for

[11]ISO: Independent System Operator. ISOs are similar to regional transmission organizations (RTOs) with slightly less responsibilities. In the paper, I will use the term ISO even when talking about RTOs as the two terms are very similar.

[12]It lacks behind PJM (88 GW) but is at the same magnitude as California (26 GW) and New York (18 GW).

[13]Dispatchable power plants can change their output by decision. All conventional power plants are usually dispatchable. Non-dispatchable technologies include wind and solar.

a part of their operating range[14]. For example, resource 34578 in my data has an economic maximum energy capacity of 543 MW and on average it offered its range from 420-540 MW for regulation with significant variation. Only power plants that are not producing at their operating limits can offer the service. As a consequence, power plants need to have leftover capacity to be able to provide regulation. Winners in this auction have to switch their power plants to regulation mode which presents a fixed cost. There are additional variable costs (energy, wear and tear) power plants incur when the option is called.

The frequency regulation market in New England is relatively small. The yearly cost of procuring the service was \$28 million in 2017 and it was in the \$20-30 million range every year from 2013 to 2017. In contrast, the total wholesale cost of electricity in 2017 was \$9.1 billion of which \$4.6 billion were energy costs. Other categories include capacity payments, transmission costs and ancillary services[15]. Frequency regulation accounted for a 0.33% share of total wholesale electricity costs or 0.65% share of energy costs[16]. Frequency regulation in New England is provided mostly by gas and pumped storage power plants. In 2008, the only year with available information on winners, the market shares of gas and pumped storage plants were 90% and 5% respectively

The amount of frequency regulation needed for the system depends mostly on electricity demand and varies by day and hour. The average yearly regulation requirement remain unchanged at around 60 MW in 2013-2019. Figure 4.1a shows the

[14]Power plants have to satisfy a set of requirements to be eligible to sell this product, see for instance ISO New England (2015) for details.

[15]The term ancillary services groups different products the system operator uses the ensure system reliability. It includes frequency regulation and various reserves. Similarly to regulation, reserve products are sold as options. There are two key differences. First, when the option is called power plants have more time to respond (few minutes for spinning reserves to hours for black start). Second, while in frequency regulation the option is almost always called in reserves it is rarely. In NE, the cost of procuring frequency regulation is around the same as procuring reserves.

[16]In comparison to other US system the ISO NE procures less regulation capacity at a higher price resulting in similar total regulation payments for the size of the system (Xu et al. (2016)). The comparison masks a number of differences between electricity systems, regulation products and market rules.

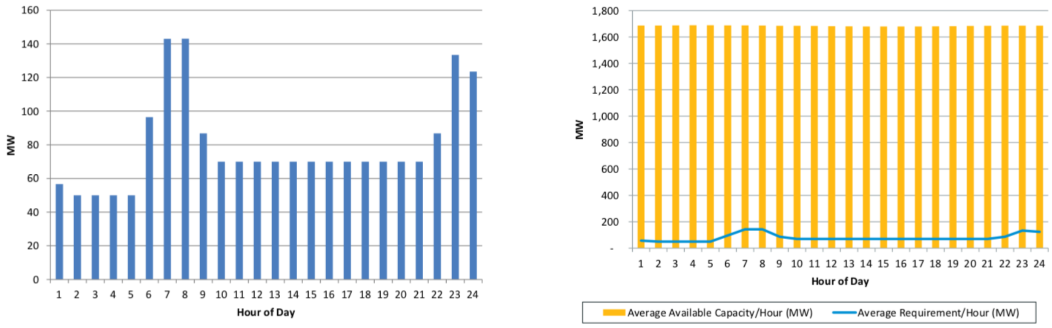

(a) Average hourly requirement, 2017 (b) Average available capacity, 2017

Source: ISO NE Internal Market Monitor (2018)

Figure 4.1: Regulation Requirement and Available Capacity, 2017[17].

average hourly requirements in 2017. The hourly variation is substantial: a multiplier of 5 exists between the lowest (30 MW for hours 1-5 in weekdays) and highest (150 MW for weekdays hours 7-8) hours. Figure 4.1b shows that in all hours of the day there was plenty of available capacity for regulation. I argue in Section 4.2.3 that due to product differentiation some producers can have strong market power. In around 10% of hours in 2017, only one unit covers regulation. In almost 50% of hours 3 or 4 units provide regulation (see Appendix C.1). These numbers suggest, that although the regulation requirement is low, the optimization results in distributing supply across several power plants.

4.2.2 The Electricity Market in New England

Frequency regulation is offered by the same power plants who participate in the energy market. Energy market outcomes and conditions have a strong influence on frequency regulation outcomes. Therefore, understanding the key characteristics of the energy market is essential for understanding the regulation market. Energy systems are characterized by substantial seasonal and daily demand variation. Because supply is mostly the same in all hours the strongest driver of prices is demand variation. Because the same capacity can not be used for energy and regulation simultaneously, high demand periods result in lower available supply for frequency regulation. This in turn has an effect on competition and prices.

The New England electricity system is characterized by a high share of gas and nuclear generation and a lack of coal power plants. Gas and nuclear power plants have a share of 48% and 31% of total electricity production in 2017 respectively[18]. The remaining 21% is distributed between hydro (8%), wind (3%), coal(2%) and other categories (8%)[19]. Yearly variation in energy prices is mostly determined by gas price variation. Gas fired power plants set the price of electricity in 60-70% of hours in real-time markets. The 4 biggest firms own 48% of capacity. In the electricity market, a pivotal supplier existed in 58% of hours in 2017 but there is significant yearly variation.

4.2.3 Frequency Regulation Market Rules in New England

The ISO NE uses a sequence of auctions (day-ahead and real-time) to procure electricity and several ancillary services products. Power plants submit a joint bid for these products and the auctions clear jointly. The ISO NE minimizes the total cost of procuring all products. By co-optimizing clearing the system operator improves efficiency and decreases wholesale electrictity costs. Complementarities between these markets arise from several sources. First, startup costs and minimum operating limits create strong non-convexities in costs[20]. Second, if the optimizing solution results in losses for power plants the ISO needs cover these losses (make-whole payments). For instance, power plant operating at its minimum capacity might not be able to recover its startup costs from market prices alone.

Market clearing is easiest to illustrate through an example. The example is designed to illustrate that in a market with ample regulation capacity, the outcome of

---

[18]All following evidence refers to the market in 2017. Source: ISO NE Internal Market Monitor (2018)

[19]Other includes oil, landfill, gas, methane, solar, steam and wood.

[20]For example, for a conventional gas fired power plant with 100 MW of capacity increasing production from 0 to 1 MW is not possible if the minimum limit is 10 MW. Increasing production from 0 to 10 MW is much more costly than increasing production from 50 to 60 MWs.

the frequency regulation market might be far from an outcome from separate market clearing.

Table 4.1: Example for Joint-Clearing of Energy and Frequency Regulation

Unit	**Min gen**	**Cap**	**Reg cap.**	**Energy bid**	**Reg bid**	**En won**	**Reg won**
1	20	100	10	30	10	95	5
2	20	100	10	35	10	85	10
3	20	400	40	100	0	0	0

Table 4.1 describes the available power plants and their offers. Units 1 and 2 are small combined cycle gas turbines (CCGT) that are efficient but relatively inflexible. They can produce up to 100 MW. Unit 1 is cheaper with an energy bid of \$30/MWh. Unit 3 is a bigger, less efficient conventional gas power plant. All plants are able to provide regulation up to 10% of their energy capacity when that capacity is available. Unit 3 is by far the cheapest source for regulation with a bid of zero. Suppose every plant receives the market clearing price for the product times the quantity won. The energy load is 180 MW and the reserve requirement is 15 MW. There are ample supplies for regulation. Under the joint cost-minimizing solution only Units 1 and 2 operate. By not operating Unit 3 the system operator can save the high energy cost of operating Unit 3. Unit 1 produces 95 MW of energy and provides 5 MW of regulation. Unit 2 produces 85 MW of energy and provides 10 MW for regulation. The market clearing price for energy is \$35/MWh. Unit 1 could have earned an additional \$5/MWh by producing on the energy market were not instructed to provide regulation. System operators have different ways of recovering these lost opportunity costs (LOC). The ISO NE includes this in the final market clearing price since 2013. This means that the market clearing price for regulation is 10+5=\$15/MWh.

This example illustrates three key points that make it very difficult to replicate market outcomes from bidding data only. First, the ISO prefers power plants that are already operating to provide regulation[21]. This can lead to relatively uncompetitive

[21]Doraszelski, Lewis and Pakes (2018) models the decision of the system operator as choosing from differentiated products where differentiation comes from whether the power plant operates.

outcomes even under ample supply of regulation such as in New England. Second, the market clearing price for frequency regulation can be very different from a simple separate clearing of equating demand with an aggregate supply curve. Lost opportunity cost creates a wedge between winning bids and clearing prices. Third, there is a high probability that bids below market price do not to win in the auction.

Until 2015, the frequency regulation product was procured by a relatively simple multi-unit auction. Offers included available capacity (MW) and a capacity bid ($/MWh). Bids were one-step functions and not increasing supply curves as common in multi-unit auctions. Additionally, power plants also submitted bids for the enegy market along with technical parameters such as startup costs and ramping rates. Selection was determined by the joint cost-minimization of electricity and regulation auctions subject to technical parameters such as startup costs and ramping rates. The clearing price equals the shadow price of the constraint on the regulation requirement. This is equivalent to a uniform price auction. Until 2013, the clearing price did not include lost opportunity cost. Plants were refunded separately for their LOCs. Since 2003, resources were compensated for the mileage (service) provided without explicitly bidding for it (FERC (2011)). Service or mileage is defined as the sum of total absolute changes in the output of plant providing regulation. Service is essentially the variable cost component of frequency regulation. The underlying service price was set to one-tenth of the capacity price to make the share of mileage payments around 50% of total regulation compensation. Dispatch was and is optimized based on bids and ramping rates (Cramton (2012)). The ISO sends an Automatic Generation Control (signal) to instruct power plants to increase/decrease output.

In 2011, the Federal Energy Regulatory Commission (FERC)[22] issued Order 755 that substantially changed the market design. FERC wanted to address unfair compensation practices and improve the efficiency of the regulation market. The intention was to introduce a more fair compensation scheme that rewards performance mostly

[22]FERC regulates wholesale electricity markets in New England.

to support storage resources providing frequency regulation. Order 755 introduced three main changes in frequency regulation markets for the electricity systems under FERC's mandate[23]. First, FERC required ISOs to introduce a market based compensation for the actual regulation service provided. Fast ramping resources, such as batteries, usually provide more mileage for the same capacity sold. Second, FERC required ISOs to link compensation to accuracy. Third, FERC required ISOs to include the lost opportunity cost in the market clearing price. Order 755 started a long procedure of updating market rules. All ISOs proposed different solutions based on their market characteristics and existing rules[24]

The New England ISO proposed a VCG auction to procure frequency regulation (ISO New England (2012)). The key theoretical problem of designing an auction is how to combine two bidding dimensions into a selection criteria. As Cramton (2012) summarizes, the ISO NE chose VCG because of its simplicity and efficiency even when only a few resources are selected. Cramton (2012) describes in detail the disadvantages of other possible auction formats[25]. Lost opportunity cost was included in the market clearing price as of July 1, 2013. The final accepted market rule was introduced on March 31, 2015. Since March 31, 2015 participants submit a service bid ($/Mile) along all previously used bidding parameters. Selection is based on minimizing the expected total cost of providing frequency regulation. The service is provided based on an ACG signal just as before. Compensation is based on actual service provided multiplied by market clearing prices along both dimensions[26]. The

---

[23]These include the California ISO, Midwest ISO, the New York ISO, ISO New England, the PJM Interconnection and the Southwest Power Pool. Source: FERC (2011).

[24]ISO NE's initial response was submitted in April 30, 2012. The actual implemented market rule changed significantly from this initial proposal.

[25]Most other ISOs (CAISO, NYISO) implemented a simple scoring auction with a fixed capacity-to-service multiplier. The main problem with this approach is the potential for skewed bidding such as in Ewerhart and Fieseler (2003). The sequential auction format proposed by Chao and Wilson (2002) and implemented in other reserve markets is not a real option in NE as their are only a few resources selected at each hour.

[26]Clearing prices are calculated in a nonstandard way. The original proposal of the ISO NE was rejected by FERC. Standard VCG does not produce prices along each bidding dimension.

mileage payment is reduced in proportion to the precision the power plant was able to follow the AGC signal.

4.3 Theoretical Predictions

This section describes two auction models. First, in Section 4.3.2 I introduce a bidding model with capacity constraints for the market environment before FERC Order 755. The model shows that in equilibrium power plants bid above marginal cost. Moreover, bids can vary based on changes in market power induced by changes either in energy demand or lost opportunity costs. Second, I introduce a multi-unit VCG auction model. The model shows that in equilibrium power plants bid truthfully. As long as the marginal cost of providing regulation stays the same bids should remain the same.

4.3.1 Setup

The system operator (ISO) attempts to procure D_t^c of frequency regulation capacity and D_t^m of mileage. Resource i can offer capacity q_{it} at a marginal cost of c_{it}^c. c_{it}^c represents fixed costs including incremental maintenance for providing regulation and fuel market or other potential risks. When power plants are called by the regulator they incur a marginal cost of c_{it}^m for providing regulation mileage. c_{it}^m represents increased operating and maintenance costs, and incremental fuel costs resulting from the generator operating less efficiently when on regulation (ISO NE Internal Market Monitor (2018)). In addition, since the power plant can not produce for the energy market it loses the potential profit it could have earned (lost opportunity cost or loc_{it}). I assume that marginal costs are constant for the entire range of the power plant's regulation capacity. Demand for each hour (D_t^c, D_t^m) in the frequency regulation market is common knowledge. The capacity requirement is published for every hour

It only produces bundled payments for each participant. In Section 4.3.3 I describe precisely the mathematical formulation of the problem. I also prove that the modified VCG design is still truthful in Appendix C.3.

of the year much before the market starts. Mileage is relatively predictable based on the past[27]. Power plants only know their own marginal costs (c_{it}^c, c_{it}^m) and the distribution of other's marginal costs. Power plants can only offer their entire range of regulation capacity at a uniform price. Since marginal costs are constant and bids can not be increasing, the analysis of bidding behavior is simpler than in general multi-unit auctions.

4.3.2 Before Order 755 - A Uniform Price Auction

Bidders submit a one-dimensional bid b_{it} for all available regulation capacity. Payments are calculated based on a uniform market clearing price which is set by the bid of the last MW unit did not win. Winners earn profits per unit of MW:

$$\pi_{it} = (p_t^c + p_t^m q_{it}^m) - (c_{it}^c + c_{it}^m q_{it}^m) = R_{it} - \tilde{c}_{it},$$

where p_t^m and q_{it}^m denote service price and mileage provided. The service price is calculated by dividing the market price by the mileage ration that is set to 10 by the system operator. I assume that all plants incur the same amount of mileage for a unit of capacity provided[28]. As a consequence, every winning MW earns the same R_{it}. The key to observe is that power plants can be characterized by pseudo-type $\tilde{c}_{it}$ which represents the entire cost of the power plant supplying 1 MW of regulation capacity[29]. $\tilde{c}_{it}$ includes fixed costs and the cost of providing service as well. $\tilde{c}_{it}$ is the only parameter that determines the relative position and one-dimensional bid (b_{it}) of a power plant.

The setup allows me to use results from one-dimensional multi-unit auctions to analyze bidding behavior. Optimal bids increase in $\tilde{c}_{it}$ and the market power of the

[27]I can allow for a known stochastic distribution of mileage. This will introduce an additional complication only when the system operator and participants disagree over this distribution.

[28]This assumes away ex-ante differences in costs incurred by providing regulation service across plants that is not explained by marginal costs.

[29]Pseudo-types play an important role in the scoring auction literature in which many multi-dimensional auctions simplify to single dimensional auctions in pseudo-types. The argument I make here is a simplified version of Asker and Cantillon (2008).

firm[30]. Any factor that influences $\tilde{c}_{it}$ increases b_{it}. High energy price increase bids because they directly raise mileage costs on the energy produced for all plants other than hydro. Energy prices also change the relative market power of bidders but this effect is likely less important. Most bidders are gas power plants and energy prices effect costs similarly for these plants. Altogether, I predict energy prices to increase bids in the uniform price auction. Demand only influences bids through the market power channel. When demand is high, there is less competition, all firms have more market power; and hence, bids are high. Finally, introducing lost opportunity costs are unlikely to change these predictions. Higher demand on the regulation market does not influence LOCs on the energy market. When energy prices are higher, lost opportunity costs are higher since there is more profits to be earned in the energy market for everyone. LOCs shift upward, therefore increasing bids. The impact on relative market power is likely to be second-order in importance.

4.3.3 After Order 755 - Multi-Dimensional Multiunit VCG Auction

This section presents a simplified model of the VCG auction implemented by the ISO NE[31]. The model is based on Chao and Wilson (2002). Power plant bids are a two-dimensional vector b^c_{it}, b^m_{it}. Both b^c_{it} and b^m_{it} are relevant for the entire capacity range. Hence, bids are not increasing supply functions. The system operator minimizes the total cost of procuring regulation given bids and constraints in every time period:

$$\begin{array}{l} J_t(b^c_t, b^m_t) = \min_{q_t, m_t} \sum_{i=1}^n b^c_{it} q_{it} + \sum_{i=1}^n b^m_{it} q^m_{it} \\ \text{(s.t.)}: \sum_{q_{it}} = D^c_t, \quad \sum_{m_{it}} = D^m_t \end{array}$$

Selection is based on cost-minimization. An alternative to think about selection is to think about the added value of each resource to the system. Added value is essentially

[30] For instance, see Hortaçsu and Puller (2008). This is a direct consequence of results by Wilson (1979) applied to procurement auctions.

[31] Here, I present the standard VCG model without the modifications requested by FERC. In Appendix C.3 I present an example to show that the modifications kept the truthfulness property.

a score with lower scores winning in the auction. Compensation is a bundled payment for providing both capacity and mileage:

$$bundle_{it} = J_{-it}(b^c_{-it}, b^m_{-it}) - J_t(b^c_t, b^m_t) + b^c_{it} * q^*_{it} + b^m_{it} * (q^{m*}_{it}).$$

The key feature of the VCG mechanism is that the compensation of a plant is independent of its bid; however, whether a plant gets selected is not. This is because $J_t(b^c_t, b^m_t)$ includes the terms $b^c_{it} * q^*_{it}$ and $b^m_{it} * (q^{m*}_{it})$ which therefore cancel out. The result is truthful reporting. Chao and Wilson (2002) shows that this mechanism results in truthful bidding and an efficient allocation. If marginal costs are reported truthfully the allocation must be ex-ante efficient[32]. The system operator minimizes total procurement costs for the true marginal costs which by definition leads to the efficient allocation. The model does not include LOCs, make-whole payments and startup costs. The beauty of the VCG design is that the truthfulness property is robust to these complications. This was the key reason why the ISO NE suggested the VCG design. Cramton (2012) argues that the VCG design results in truthful reporting even with LOCs. In Appendix C.2, I provide an example to illustrate why players bid truthfully even with LOCs.

As a consequence of truthful bidding, neither energy prices nor demand influences capacity bids. Energy prices increase service marginal costs; and therefore service bids. Energy prices also increase LOCs but these do not change capacity bids under the VCG auction format. Finally, since demand is independent of costs, it does not impact bids either.

[32] Since I assume away uncertainty in mileage, this means the allocation is also ex-post efficient. Bid skewing incentives are only present if the system operator and the power plants disagree about expected mileage or if there is heterogeneity across plants. This is a well documented incentive in the scoring auction literature also noted by Chao and Wilson (2002). See for instance Athey and Levin (2001) for common value or Bolotnyy and Vasserman (2018) for a private value environment

4.3.4 Summary of Testable Predictions

The key difference between the two regimes lies in bidding incentives. Plants optimally report their true marginal costs across all bidding dimensions under VCG. The uniform price auction merges different costs into a one-dimensional bid and introduces markups related to market power. I do not observe marginal costs or auction outcomes that would allow me to infer them. Therefore, I need to test indirect implications of the two model such as the impact of energy prices and demand on bids. In addition, I also analyze variability of bids. The key observation is that under VCG, power plants should rarely change their bids. Underlying marginal costs are unlikely to change[33]. However, under the uniform price auction power plants should change bids more frequently. Energy prices and demand shift marginal costs and market power and influence bidding incentives. Table 4.2 summarizes the predictions of the two theoretical models.

Table 4.2: Theoretical Predictions

#	**Model**	**Prediction**
1	VCG	Capacity bids are independent of energy prices.
2	VCG	Capacity bids are independent of regulation demand.
3	VCG	Energy prices influence mileage bids.
4	Uniform price	Energy prices have an impact on capacity bids.
5	Uniform price	Regulation demand influences bids.
6	Comparison	The impact of energy price is higher for bids before 2015.
7	Comparison	The impact of demand is higher for bids before 2015.
8	Comparison	The variance of capacity bids are higher before 2015.

4.4 Data and Descriptive Statistics

In this section, I describe the data and present key descriptive statistics of the NE frequency regulation market. Prices increase after the introduction of the new market design in 2015. Bid descriptives show that this price increase is likely caused by

[33] Doraszelski, Lewis and Pakes (2018) calculates the marginal costs of providing regulation in the UK. They assume that the marginal cost do not change over time. Their marginal cost estimates are very low (in the range of 1.5-3$/MWh in 2003 prices.)

selection due to new eligibility requirements for participants. I find that both the average and the mean of bids for power plants always bidding decreased after the market design change. The median plant does not change bids through time but is also unlikely to win.

4.4.1 Data

Most of the data comes from publicly available information on the New England ISO's website[34]. I observe all bids for energy and frequency regulation auctions with a 3 months time lag from 1999. Bidder and plant identities are masked consistenly across the dataset. I can follow individual plants through time and I see whether plants are owned by the same company. I observe a set of physical characteristics (capacities, startup costs, ramping rates)[35]. Data is not available for settlements. I do not observe the winners of these auctions or the actual market shares of participants[36]. This is crucial because due to the complications of the market clearing process I can not replicate auction outcomes with reasonable precision. The sample I analyze in this paper goes from 2013 to 2019. There is a day-ahead (DA) and a real-time (RT) auction. I focus on DA as this is the key mechanism to procure regulation. Similarly to most electricity markets, RT is a deviation market from DA and is of secondary importance. Most participants keep their DA bids when submitting for RT.

4.4.2 Prices

Table 4.3 presents the mean and standard deviation of hourly regulation capacity prices (measured in \$/MWh) respectively. The price increase from 2013 from 2015 (Column 1 to 2 and 3) is mostly due to LOCs being included in prices. As I show

---

[34]I download all data using ISO NE Web Services API available at: https://webservices.iso-ne.com/docs/v1.1/

[35]Although I do not know the identify of the plant, some characteristics might be possible to infer from bidding behavior on the energy market.

[36]Communications with the ISO NE made it clear that it is almost impossible to get this data. The ISO NE has a data policy that is signed by market participants. Any additional data is only available if allowed by all market participants which is almost impossible to achive.

later in detail, the price increase after 2015 is mostly due to selection of power plants eligible to offer regulation. The standard deviation of prices does not show a clear pattern but it is considerable in magnitude. Moreover, the distribution shows that high prices are not uncommon which suggests plants can win on the auction even with high bids as long as their LOCs are relatively low. Service prices are a magnitude smaller, again with significant standard deviation (see Appendix C.4 for details). Total service payment only accounted for 14% in 2017 (ISO NE Internal Market Monitor (2018))[37]. Figure 4.2 shows that there are significant differences across daily average capacity prices at least partially caused by outcomes in the energy market. The correlation of hourly regulation capacity and energy prices is 0.40 before and 0.23 after the market design change with strongly yearly and monthly variation. The data is suggestive that energy and regulation prices move less strongly together after the VCG design was introduced.

Table 4.3: NE Regulation Capacity Price Distribution

	2013	2014	2015b	2015a	2016	2017	2018	2019
Mean	11.7	19.0	23.6	23.6	27.3	29.2	28.3	22.0
Standard deviation	18.6	37.4	30.8	30.8	33.3	38.1	44.5	21.7
Distribution								
10th percentile	6.0	6.3	8.3	8.3	9.4	8.3	8.5	8.3
25th percentile	6.7	7.3	11.2	11.2	12.6	10.8	11.0	10.5
Median	7.5	9.6	16.2	16.2	18.6	17.0	16.9	14.7
75th percentile	10.6	16.0	26.8	26.8	31.4	32.5	31.4	24.8
90th percentile	18.7	38.2	42.2	42.2	51.8	60.8	56.3	42.5
Observations	8760	8760	2135	6625	8784	8760	8759	8759

Notes: An observation is an hour. 2015b and 2015a is before and after the rule change respectively.

4.4.3 Bid Descriptives

Table 4.4 presents bid descriptive statistics for plants who bid every year from 2013 to 2019. Since 2015, the set of eligible plants shrank as the ISO New England changed

[37]Unfortunately, detailed information of settlements is not available. The system operator expected service prices to be much higher based on the 50-50% rule before 2015.

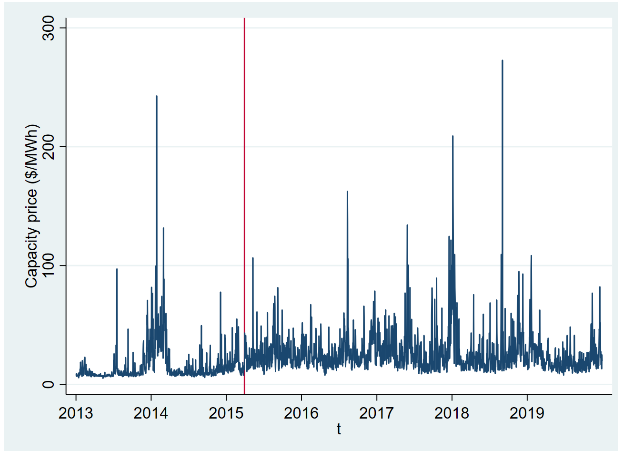

Figure 4.2: Daily Average (Unweighted) Regulation Capacity Prices, 2013-2017

the technical parameters of providing regulation. I observe 81 plants who bid in all years. First, I take the yearly mean and standard deviation of hourly capacity bids for each of these plants. Then, I report the mean and median of the 73 mean and standard deviation. I repeat the exercise for service bids.

Table 4.4: NE Regulation Individual Bid Statistics - Always Bidders

	2013	2014	2015b	2015a	2016	2017	2018	2019
Mean - cap	24.80	28.50	28.98	22.08	22.75	22.14	21.54	21.44
	(15.85)	(19.40)	(15.02)	(8.99)	(8.50)	(8.00)	(7.88)	(9.00)
Sd - cap	5.59	7.70	3.19	1.84	2.01	2.76	2.38	3.44
	(0.00)	(0.00)	(0.00)	(0.56)	(0.00)	(0.00)	(0.00)	(0.42)
Mean - serv				1.81	2.01	1.57	1.27	1.45
				(0.14)	(0.19)	(0.06)	(0.01)	(0.05)
Sd - serv				0.37	0.29	0.42	0.18	0.42
				(0.08)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	73	73	73	72	73	73	73	73

Notes: An observation is a plant-year. 2015b and 2015a is before and after the rule change respectively. Only plants who bid every year included. Statistics aggregated to the plant-year level. Variable "Mean - cap bid" represents the individual plant's mean capacity bid in a given year. Means, medians in parenthesis.

The mean and median of individual mean capacity bids dropped significantly after the market design change (rows 1-2). The mean capacity bid decreased from 30.28 $/MWh in 2014 to 23.79 $/MWh in 2016 while the median decreased from 19.40

\$/MWh to 8.99 \$/MWh. Similarly, the mean standard deviation of capacity bids dropped (rows 3-4). However, the median standard deviation of bids is zero in 2014, 2015, 2016, 2017 and 2018. The median plant submitted the same bid for all hours in these years. Similarly, the median plant did not change its service bid in 2016 and 2017. One potential explanation is that these power plants are either certain winners or never win in these auctions. There is some evidence for experimentation. The median plant changes its bid through 2015 after the design change.

Doraszelski, Lewis and Pakes (2018) estimates a marginal cost of 1.5-3 \$/MWh for gas plants for providing regulation including the mileage costs as well. Bids are way above this estimate in 2016 indicating bidding above marginal costs.

4.5 Empirical Results

In this section, I present the main empirical results of this paper. First, I report regression outcomes on capacity bids as dependent variable. Players bid high when demand is high and when energy prices are high. Both findings are against the theoretical predictions of Section 4.3. Then, I investigate individual bidding behavior in more depth. I categorize bidders into four groups and find significant differences across these groups. I attribute these findings to the failure of learning the optimal bidding strategy. Finally, I discuss alternative explanations.

4.5.1 Bid Regressions

In order to investigate how power plants modify their bids through time I estimate regression models for capacity bids. I focus only on power plants who bid in all years from 2013 to 2019 and who did vary their bids in a given year. This leaves me with 45 and 47 bidders before and after 2015 respectively. Changer plants bid lower on average confirming the hypothesis that they have a higher chance of frequently winning in these auctions (Appendix C.4 for details). I estimate regression models for predicting capacity bids of power plants before and after the market design change.

I choose to include all years in my data in the preferred specification:[38]

$$b_{it}^{c} = \alpha + \beta_1 * p_t^{energy} + \beta_2 * D_t^c + \gamma controls_{it} + \lambda_i + \epsilon_{it}, \tag{4.1}$$

where D_{it}^c is regulation requirement, λ_i are plant specific fixed effects, while controls include regulation status (available or not), regulation capacity[39] and regulation response rate. Theory models predict (see Table 4.2):

$$1: \beta_1^{after} = 0, \quad 2: \beta_2^{after} = 0, \quad \text{or at least} \quad 6: \beta_1^{after} < \beta_1^{before}, \quad 7: \beta_2^{after} < \beta_2^{before}.$$

Table 4.5 presents the results of this model for the two specifications before (columns 1 and 2) and after the policy change (columns 3 and 4). I estimate the two models separately which allows me to estimate separate behavior in the two regimes including the fixed effects. In order to allow for autocorrelation in the bids submitted by the same power plant, I cluster standard errors at the power plant level[40]. My preferred specifications include power plant fixed effects (columns 2 and 4). I report OLS estimates (columns 1 and 3) for comparison. The estimates reject predictions 1, 2 and 7 but confirm prediction 6. The coefficient on energy prices (β_1) is statistically different from 0 after 2015. The impact of energy prices on bids is lower after the policy change. On average, a \$1 increase in the energy price led to a \$0.145 and \$0.0352 increase in capacity bids before and after the policy change respectively. Both estimates are significant at the 5% level[41]. In addition, the two coefficient estimates are statistically different at the 1% level. Demand has a sta-

[38] I experiment with alternative specifications of leaving some years out. Winters in New England can experience gas shortages which were especially severe in 2014. NE does not have sufficient gas pipeline capacity to cover winter months when residential gas demand has priority. Gas power plants might not be able to operate leading to unusual market environment. Unavailable gas production is substituted by oil. In addition, there is slightly more experimentation in 2015. Dropping these unusual years do not change the results substantially.

[39] With regulation requirement (demand) included, there is no need to control for month, day and hour.

[40] Because I only have around 50 clusters my standard errors increase significantly.

[41] Using robust but not clustered standard errors would make the coefficients statistically significant even at the 0.1% level.

Table 4.5: Capacity Bid Regressions

	Before OLS	Before FE	After OLS	After FE
Energy price ($/MWh)	0.149***	0.145***	0.0129	0.0352**
	(0.0242)	(0.0233)	(0.0192)	(0.0146)
Reg. requirement (MW)	-0.000251	0.0000781	0.00248*	0.00273***
	(0.00139)	(0.00121)	(0.00134)	(0.000882)
Status	11.98***	1.104	-2.802	-1.572*
	(4.008)	(2.615)	(2.916)	(0.936)
Regulation capacity (MW)	0.0970	-0.0436	0.0255	0.0103
	(0.0685)	(0.0273)	(0.0726)	(0.00782)
Response rate (MW/min)	0.252		0.487	
	(0.331)		(0.325)	
Constant	0.756	17.49***	6.899	11.90***
	(4.365)	(2.497)	(4.888)	(1.158)
Observations	664687	672412	1325959	1325959
Adjusted R^2	0.194	0.562	0.101	0.658
Clusters	45	45	47	47

Notes: an observation is a power plant in a given hour. Dependent variable is the capacity bid. Before and after are relative to FERC Order 755 and include years 2013-2015Q1 and 2015Q3-2019 respectively. FE: power plant fixed effects. Only plants who bid every year and change their bids in a given year included.
Standard errors are clustered at the power plant level and are reported in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

tistically significant impact on capacity bids ($\beta_2 \neq 0$) after the policy change. In contrast, before the policy change demand does not seem to influence bids for the average plant.

4.5.2 Individual Bidding Patterns

Although the regression results in Section 4.5.1 are informative of the aggregate pattern they mask individual differences. As I show in this section, power plants follow different bidding strategies. I classify the bidding behavior of power plants to 5 groups[42]. Plants in the first group (frequency) change their bids in high frequency seemingly following some external signal. The second group (occasional) is characterized by occasional bid changes. Plants in the third group (hourly) seem to follow a recurring bidding pattern. In particular, most of these plants bid the same depending on the hour of the day. This might happen if plants follow regulation requirements in their bids. The fourth group (changer) is characterized by plants who changed their bidding strategy. Finally, some plants never change their bids. Figure 4.3 presents examples for bid timelines for the first 4 groups.

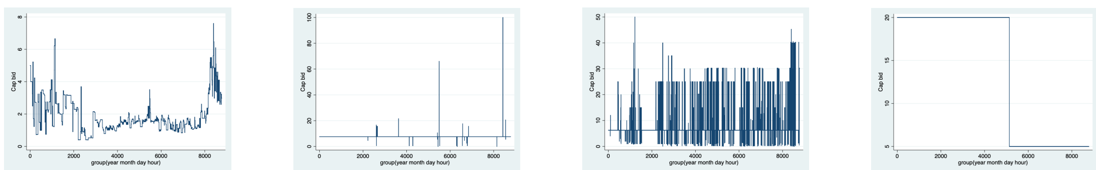

(a) Frequency (b) Occasional (c) Hourly (d) Changer

Figure 4.3: Bidder Timelines 2016. Group - Representative Bidder

I estimate fixed effects models similar to Equation (4.1) for the four groups separately. Table 4.6 presents the results of these estimates. Groups follow quantitatively different bidding strategies. Group 1, seems to follow the energy price closely. Most likely, they actually modify their capacity bids as the underlying fuel (gas) price changes. The impact of energy prices is economically significant for this group. A 1\$/MWh price energy price increase is associated with a 0.0962\$/MWh increase in individual bids. Although, the coefficients of regulation requirement (demand) and regulation capacity are statistically significant their magnitude is very small. The model explains around 53.6% of the variation in bids which is high considering the frequency of bid changes. Occasional bidders seem to modify their bids occasionally without a clear pattern. Plants in the third group seem to modify bids mostly based on regulation demand. The impact is both statistically and economically significant.

[42] The classification is based on visually investigating how bids of an individual plant changes over time. See Appendix C.5 for these graphs for each group.

On average, bids are \$0.96 higher for a 100 MW increase in demand (the difference between demand at 4am and 8am). Finally, changers group plants who appear to have changed bidding strategies.

Table 4.6: Capacity Bid Regressions by Group (After 2015)

	(1) Frequency	(2) Occasional	(3) Hourly	(4) Changer
Energy price (\$/MWh)	0.0962***	0.00450	0.00608	-0.00511
	(0.0278)	(0.00767)	(0.00336)	(0.00446)
Regulation requirement (MW)	0.000497	-0.000331	0.00958**	0.00121
	(0.000715)	(0.000210)	(0.00247)	(0.00119)
Status	-2.515	0.138	-1.736***	-0.352
	(1.537)	(0.319)	(0.106)	(0.262)
Regulation capacity (MW)	-0.00324	0.00641	0.0565***	0.0331***
	(0.00414)	(0.00550)	(0.0121)	(0.000545)
Constant	13.54***	8.954***	5.309***	6.101***
	(1.057)	(0.617)	(0.788)	(0.406)
Observations	398117	224468	210048	105024
Adjusted R^2	0.536	0.638	0.057	0.232
Clusters	12	9	6	3

Notes: an observation is a power plant in a given hour. Dependent variable is the capacity bid. Only after FERC Order 755 included, 2016-2019. All specifications include power plant fixed effects. Only plants who bid every year and change their bids in a given year included.
Standard errors are clustered at the power plant level and are reported in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.5.3 Summary of Findings and Interpretation

Table 4.7 summarizes the results of empirical tests of theory predictions. The results question the theoretical prediction of truthful bidding under the VCG design. In particular, they suggest that bidders followed similar bidding strategies before and after the market design change. The results are surprising because bidding under VCG is simple if bidders know their own marginal costs. For example, Hortaçsu and

Puller (2008) proposes VCG as an option over more strategically complex mechanism such as the uniform price auction. My results suggest that VCG does not resolve the problems of strategic complexity. Power plant operators still follow different strategies possibly due to differences in strategic sophistication[43].

Frequently repeated interactions should allow power plant operators to converge to best response behavior such as in Doraszelski, Lewis and Pakes (2018). I provide four reasons why the learning channel might fail to deliver optimal bidding. First, power plant operators might not be informed of market rules to sufficient detail. Bidders might have started with beliefs that previously optimal bidding is still optimal. Second, although optimal bidding is simple, proving that truthful bidding is optimal is hard. Third, the stakes are not too high. Any optimization is likely to be costly for power plant operators. Simply, the cost of optimizing bidding might overweigh the benefits[44]. Fourth, learning is particularly difficult in this environment. Every hour, only a few power plant operators win in the auction often due to energy auction outcomes. As a result, linking regulation bids to compensation might be difficult even with significant effort[45]. Altogether, my findings warn against using bids from a VCG mechanism to identify marginal costs.

4.5.4 Alternative Explanations

First, I consider the possibility that demand and energy prices are correlated with marginal costs. Gas fired power plants provide the bulk of regulation and gas prices are highly correlated with energy prices in New England. Higher gas prices might increase the opportunity or the actual cost of providing regulation. Lost opportunity cost should not influence bidding under VCG (see Appendix C.2.). If providing reg-

[43]See for example, Hortaçsu et al. (2017)

[44]Hortaçsu and Puller (2008) finds that large power plant operators are closer to best response bidding in the Texas energy balancing market. Small operators follow rule of thumb strategies. They also find that the benefits of optimization are likely to be too small for small operators to invest in optimization.

[45]Doraszelski, Lewis and Pakes (2018) finds that bidders converge to an equilibrium best described by Nash best response bidding in a similar environment.

Table 4.7: Theoretical Predictions

#	Model	Prediction	Finding
1	VCG	Capacity bids are independent of energy prices.	×
2	VCG	Capacity bids are independent of regulation demand.	×
3	VCG	Energy prices influence mileage bids.	×
4	Uni. price	Energy prices have an impact on capacity bids.	✓
5	Uni. price	Regulation demand does influence bids.	×
6	Compare	The impact of energy price is higher for bids before 2015.	✓
7	Compare	The impact of demand is higher for bids before 2015.	×
8	Compare	The variance of capacity bids are higher before 2015.	✓

ulation is not energy neutral power plants fuel prices directly influence costs. This effect is likely to be small and likely to alter service bids rather than capacity bids. However, I find that service bids are responsive to energy prices (see Appendix C.4 for details). A 10 \$/MWh is associated with a 0.02 \$/Mile service bid increase which is around 10% of the market clearing price. However, the impact is not statistically significant using clustered standard errors. Given the economically significant response of service bids there is no reason why power plants should increase capacity bids as well. Finally, fuel prices do not explain why power plants respond to demand changes.

Second, owners of multiple plants might co-optimize. Competitive auction models assume that bidders do not coordinate bidding. My sample consists of companies owning multiple power plants. The VCG design does not take ownership into account. Resources are selected and compensated as if they are owned by different companies. Compensation of a winning bids is based on bids of losing bids. Hence, a bidder might bid higher in its losing resources to increase the compensation of its winning resources. As Cramton (2012) also highlights, this creates an incentive to overbid only for rejected marginal resources. Cramton (2012) argues that it is difficult to know in advance which resource will be optimal. Moreover, any overbidding is limited by the offers of other bidders.

Finally, it is unlikely that market might is still under a phase of experimentation. Doraszelski, Lewis and Pakes (2018) finds that the UK frequency regulation market needed 3.5-4 years with monthly interactions to converge to a new equilibrium. In New England, participants interact daily. Therefore, it is likely that in 2018 or 2019 (after more than 1000 interactions) the market would be in the new equilibrium. In addition, the data shows that firms experimented mostly until the end of the 2015 calendar year.

4.6 Conclusions

I provide evidence that power plants do not bid optimally in the ISO NE frequency regulation auction after the 2015 change to a VCG design. Power plants change their bids when demand changes. High demand is associated with weaker competition and higher market power of participants. However, under the modified VCG auction bids should be independent of market power. My preferred interpretation is that finding optimal bidding by experimentation is hard. Bidders kept strategies used before the market design change without major updating. My results suggest that VCG algorithms might not be superior in complex auction environments. This paper also questions the validity of using VCG bids to identify marginal costs. A key limitation is that I do not observe market outcomes therefore it is difficult to confidently say that power plants "leave money on the table".

5

Conclusions

In Chapter 2, I estimate carbon dioxide marginal abatement costs for various manufacturing industries using observational data for the first time. The empirical strategy I propose is new. The key advantage of my methodology is that it can be used with widely available data on production variables and emissions. I find that marginal abatement costs are the lowest in the cement, chemicals and power industries. When introducing a carbon tax is not politically feasible, targeting these industries with subsidies to decrease emissions would increase social welfare the most. My results have implications for the market design of carbon markets around the world. Cement, chemicals and power firms are the most likely to decrease emissions as the carbon price rises to levels close to the social cost of carbon. My estimates show vast heterogeneity both within and across industries but no clear gap in the distribution of marginal abatement costs. This suggests that it is unlikely to observe drastic price increases when the allowance caps is lowered.

In Chapter 3, I study the transition towards a cleaner power generation fleet in PJM. I extend non-stationary dynamic competitive models by adding multiple technologies and an hourly spot market. Both of these are key characteristics of power generation markets. My findings highlight the importance of analyzing the full transition path when comparing environmental policy instruments. Policies that lead to similar long-term outcomes induce vastly different transition dynamics. Neither entry nor exit subsidies come close to the efficient outcome which is achievable by a carbon tax. My results show that these instruments are a poor substitute for carbon taxes when welfare along the transition path is considered.

In Chapter 4, I analyze bidding in a multi-unit VCG auction. My setup is unique as there is virtually no observational evidence on how the VCG mechanism performs for multiunit auctions. Bidding under a VCG design is simple since truthful bidding is optimal. However, I find that participants bid higher when relative market power increases. Additionally, my results suggest heterogeneity in bidding strategies and sophistication. My preferred interpretation is that players were not able to learn optimal bidding most likely due to a combination of low stake and a complicated market design in which it is difficult to link bids to outcomes. Altogether, my results suggest that using VCG bids as estimates of true marginal costs can be misleading in complex environments.

Appendix A

Appendix to Chapter 2

A.1 Data Construction and Quality

A.1.1 EU ETS Registry Data

I access data through the European Union Transaction Log (EUTL)[1]. The unit of regulation is an installation (plant). However, trading is done at the firm level. The registry administers several types of accounts used for compliance and trade in the program. The two most important are Operating Holding Accounts (OHA) and Person Holding Accounts (PHA). OHAs belong to firms that directly own installations. A typical OHA is a firm at the country level who operates installations. PHAs do not directly own installations but might own them indirectly. A typical PHA is the headquarters or the trading arm of a multinational firm. Compliance with the program is completely through OHAs and I focus my analysis on the activities of them. PHA accounts might be used to reshuffle allowances in between firms that belong to the same parent company.

I get information on installations and OHAs by scraping the "Operating Holding Accounts" tab on the EUTL website. The data contains identifying information on

[1] Data is available through The European Union Transaction Log (EUTL) website at `https://ec.europa.eu/clima/ets/`.

both the installation and the OHA it belongs to. I observe the main activity of the installation, its emissions and surrendered allowances. This is the only source of data at the installation level that I use. I observe a national company identification number for OHAs. Since this data contains mistakes, I use Jaraitė et al. (2013)'s cleaned national ID data[2].

I scrape the "Transactions" tab on the EUTL website to get information on allowance trading. In contrast to the OHA data this requires more cleaning. First, the timing of the program does not perfectly correspond to calendar years. Firms have to surrender allowances by the end of April the following year to cover their yearly obligations. Therefore, I create the relevant year from May to May for transactions trading. Second, the data contains transactions of all types including the initial allocation and auctioning of allowances, trading inside and outside firms, surrenders and administrative transactions. I categorize and clean these based on the description of codes and the entities involved. Eventually, I need to be able to build variables describing banking and trading behavior of firms. I use the initial allocation of allowances and surrenders as observed in the OHA data[3]. I include inside and outside firm transaction and auctioned allowances in net trade yearly balances for OHAs[4]. Then, I generate the full yearly balance by adding net trades to the initial allowances minus surrenders. Finally, the current bank of allowances is derived by adding up the yearly balance to last year's bank. Transferring allowances is not allowed from Phase I to Phase II, so the bank starts at 0 in 2008.

Even though the two EU ETS registry datasets come from the same source, merging them is not trivial. There is no numerical id linking the two datasets. Therefore, I have to merge on account names. The accounts dataset contains OHAs, whereas

[2]Most of the cleaning they do is about company identifiers in Germany that are often misspecified in the original data. Since most accounts were established early, even though done in 2013 the cleaning takes cares of most of the issues. There is also less mistakes in later years.

[3]Surrenders almost always equal to emissions.

[4]The auctioned allowances are hard to identify from the data.

transactions have PHAs as well. An OHA might not show up in the transaction database and a PHA might have the same name as a missing OHA. My scraping algorithm does not always produce consistent results for strings which I also have to correct. Therefore, I treat the accounts dataset as the basis and add transaction information to it using OHA accounts only. This results in an almost perfect merge but possible loses some OHA trades. I match 90% of total transactions done by OHAs to OHA accounts in my data. I treat firms without observed transactions as if they did not trade at all.

A.1.2 Orbis Dataset

Orbis is a global database that includes information on public and private firms' high level balance sheet and income statements that is provided by Bureau van Dijk[5]. Orbis's European module is called Amadeus which covers all countries under the EU ETS. Amadeus also includes information on company ownership. Data quality varies among countries but is generally considered to be better for the EU. In all EU countries, even small firms must file financial statement information to a national registry for administrative and tax purposes. BvD collects information from these national registries[6]. Kalemli-Ozcan et al. (2015) provides a detailed guide on how to use Orbis data for economics research. They are able to cover around 70% of total economic activity in the EU with Orbis while also being able to reproduce firm distributions produced by Eurostat. Furthermore, Orbis data is widely used in the literature on finance, trade and macroeconomics and more recently in the literature on productivity and markups[7]. I access historical Orbis data through the license of the Ford Library of the Fuqua School of Business of Duke University[8]. I follow

[5]Orbis is an umbrella term to describe the entire portfolio of BvDs offerings.

[6]The requirements for filing vary across the EU. See Kalemli-Ozcan et al. (2015) for a review.

[7]Recent papers using Orbis data to measure productivity and markups include productivity and markups include Asker, Collard-Wexler and De Loecker (2014) and Autor et al. (2020).

[8]Depending on the source of access, the data quality can vary. Historical Orbis data on disks is generally considered to be the best quality (Kalemli-Ozcan et al. (2015)). The data I have access to

Kalemli-Ozcan et al. (2015) to process the Orbis dataset.

Because Orbis is based on financial statements it is not as detailed on the operation of firms as manufacturing censuses. Furthermore, it is at the firm and not the plant level. The data includes production variables in monetary terms such as revenues (output measure) and materials costs (intermediate inputs). I observe the wage bill and the number of employees separately. Since I do not observe investments, the capital measure I can use is balance sheet based (fixed assets). In addition, Orbis provides several company identifiers and industry classification codes (SIC, NACE, NAICS). Data availability varies by country and variable. For the EU ETS firms that I am able to match, balance sheet data is available for around 95% of firm-year observations. That is, for 95% of the firm-year observations when Orbis reports a database entry I also observe balance sheet variables. Key lines of income statements such as revenues and EBIT are available for around 90%. The wage bill, number of employees and materials costs are available for around 80%, 80% and 70% respectively. Orbis does not track firms through time consistently. Firms might start reporting late, stop reporting, altogether or miss reporting in some years. Balanced panels of Orbis do not actually represent real entry and exit. Furthermore, the number of firms without missing observations for a ten year time period is very small.

A.1.3 Linking EU ETS Data with Orbis

I link the EU ETS data with Orbis using national company identifiers. Orbis reports several national identifiers for each firm. For most countries, this includes at least a tax identifier, likely one from the trade registry and possibly others. The source of the identifiers used by the EU ETS is not defined and varies by the country. Therefore, I collect all BvD ids that belong to any national ID that I observe in the EU ETS dataset. To arrive at a unique BvD id for each national identifier I make the following choices. First, I keep only consolidated entries when multiple available.

is essentially a replication of the historical disks.

Second, I keep only end of the year reports for firms who report multiple times a year. Finally, I keep only one source if there are multiple sources of reporting for the same company-year. In some cases, there is several BvD id-s belonging to the same national ID. These are usually cases when the national ID belongs to the mother company whereas BvD has more detailed information on subsidiaries. I make the general rule of picking the mother companies as this seems to be the optimal choice for the observations I manually check. After these steps, there are still multiple BvDids for some accounts that I clean manually. The above procedure assigns Orbis BvDid matches to around 80% of EU ETS OHA accounts in my data[9]. The success of the merge varies substantially across countries: it is above 90% in Italy, but below 70% in the Netherlands.

A.1.4 Data Quality

Table A.1 reports statistics describing the quality of the data in more detail. The data covers 12 years (2005-2016) and 8707 EU ETS accounts. Column (1) reports statistics for the full potential sample, notably, around 31% of observations have zero or missing emissions. Column (2) contains the 80% of accounts for which I am able to find an Orbis match irrespective of whether the data is missing for a given year. Column (3) reports the full sample I am able to match to Orbis irrespective of whether Orbis have all the variables available. I lose around 14 thousand (or 17%) of observations due to missing yearly reporting or firm exits in Orbis. Column (3) covers around 75% percent of all emissions. Column (1)-(3) are very similar in variables of the EU ETS accounts data. I tend to lose observations with zero emissions in a given year. Column (4) describes the sample for which I observe the full set of variables required for production function estimation. Compared to column (3) I lose the 25% of observations with zero emissions and the rest to variables missing in Orbis. The biggest drop is due to missing materials costs. The production function sample covers

[9] I experimented with fuzzy matching based on account holder names but did not manage to improve the rate much.

Table A.1: Data Availability and Selection

	(1) All	(2) Orbis ever	(3) Orbis	(4) Prod fn	(5) Balanced	(6) Transactions
zero_emis	0.31	0.28	0.25	0.00	0.00	0.05
log_emissions	9.88	9.92	9.93	9.91	10.37	9.92
log_allocation	10.22	10.27	10.29	10.30	10.66	10.30
log_sales			17.41	17.74	18.31	17.57
log_materialcosts			16.50	16.88	17.52	16.66
log_laborcost			15.57	15.65	16.21	15.67
log_employees			5.10	5.18	5.47	5.18
log_fixedassets			16.98	17.24	17.82	17.17
log_ebit			15.06	15.08	15.52	15.14
log_balance_year						8.79
anytrade						0.53
Observations	104484	86320	72013	31767	9828	51577

Notes. An observation is a firm-year. Variable means reported. Columns indicate different samples.

around 50% of all emissions and has firms that are generally bigger. The difference is around 36% (or 36 log points) in sales, 26% in capital and 7% in the number of employees. As expected, we miss variables and observations for smaller firms more often. However, the difference in emissions is very small. Moving to a balanced sample with all production variables available (Column (5)) comes with a loss of 70% of observations. The balanced sample is also very different in size (57 log points bigger in sales) and emissions (36 log points). Column (6) contains the full sample for which I have transactions, Orbis and EU accounts data for (Column (3)+transactions). Compared to Column (3) I lose 20% of observations with no emissions in a given year[10]. The remaining firms are modestly bigger (in sales, capital and labor) but do not emit more. Again, this is expected, as it is more likely that bigger firms ever show up in the transactions dataset. Around 53% of firm-year observations does participate in trading in the EU ETS with lower shares trading in earlier years.

[10]This is somewhat surprising since most OHAs receive free allocations; therefore should show up in the transactions dataset even if they are not trading. This is the main reason I use initial allocation of allowances from the accounts data.

A.2 Industry Classification

A.2.1 Detailed Description of Industries

This section describes the industry categorization in detail. EU ETS defines a main activity for each plant under regulation. Whether a firm is included in the EU ETS is roughly based on these activity types[11]. I assign activity categories to the firm based on the most common activity among its plants. Around 87% of the firms have plants with only one activity type. A common pattern across the remaining firms is to have a plant classified as power which I read as the plant has on-site power generation. I categorize the firm in this case by the other activity type it has[12]. Eventually I am interested in assigning the main activity type to a firm. Therefore, all activities leading to a final product are grouped together. An example is coke production. Although the production of coke might be closer to refining of oil than the production of steel I include it under the category "steel". Steel firms likely own both a coke and a steel plant and this way the firm can clearly be classified as a "steel" firm. I very briefly describe the manufacturing process and the main sources of carbon dioxide emissions. Finally, I also provide rationale for some of the less obvious groupings.

The category cement contains plants across the whole chain of cement production[13]. Cement production starts with mining limestone. For most cement plants, the mine is either on-site or very close due to high transporting costs. Limestone is heated with supplements to produce clinker, the main ingredient of cement. The data contains both integrated and unintegrated cement plants but also standalone limestone plants. Around one-third of plants seem to be limestone only, while the

[11]The EU ETS has a finer product-based categorization used for benchmarking which I do not have access to.

[12]The most common is an other and power combination, but almost all other categories have a combination with power as well.

[13]EU ETS activities: "Production of cement clinker", "Installations for the production of cement clinker in rotary kilns or lime in rotary kilns or in other furnaces", and "Production of lime, or calcination of dolomite/magnesite".

share of unintegrated plants is not clear. Carbon emissions come from the chemical process of separating $CaCO_3$ to CaO (clinker) and CO_2 and the production of heat and electricity.

Ceramics include production of bricks, tiles, porcelain and gypsum or plasterboard, and other ceramics. Ceramics are hard, brittle and resistant materials made by firing nonmetallic minerals (usually clay) at a high temperature[14]. Plasterboard is included here since it is mostly produced from clay and the key product is used in construction. Gypsum or plasterboard is around 2% of total firms in the ceramics category and are unlikely to significantly influence the estimates. The required heat is usually generated by burning fossil fuels (mostly carbon) which is the key source of CO_2 emissions.

The category chemicals include petrochemicals (e.g. plastic), carbon black, ammonia, hydrogen, nitric acid and sodium carbonate[15]. The biggest category is the production of bulk chemicals (75% of observations) that is likely petrochemicals. The most common process of manufacturing petrochemicals is steam cracking which requires high temperatures. Similarly to petrochemicals, carbon black is manufactured from heavy petroleum products and is most often used in tires. Ammonia is produced by reacting nitrogen in a high temperature high pressure environment. Nitric acid production requires ammonia and heat. All chemicals production is highly energy intensive. The bulk of carbon emissions is coming from creating a high heat environment by burning fossil fuels.

Glass[16] includes industrial and commercial glass manufacturers and mineral wool.

[14]EU ETS activities: "Manufacture of ceramics" ,"Installations for the manufacture of ceramic products by firing, in particular roofing tiles, bricks, refractory bricks, tiles, stoneware or porcelain", and "Production or processing of gypsum or plasterboard".

[15]EU ETS categories: "Production of bulk chemicals", "Production of carbon black","Production of adipic acid","Production of ammonia", "Production of hydrogen and synthesis gas","Production of nitric acid" ,"Production of soda ash and sodium bicarbonate"

[16]EU ETS activity types: "Installations for the manufacture of glass including glass fiber", "Manufacture of glass" and "Manufacture of mineral wool". I decided not to include sodium bicarbonate production. It is an important raw material but its production process puts it closer to chemicals than glass.

Glass is most often produced by melting raw materials (silica, sodium carbonate) in a furnace at high temperatures. Mineral wool production is relatively similar and is often categorized as glass[17]. The source of carbon emissions is again generating heat by burning fossil fuels.

Non-ferrous metals is the production of several metals that do not contain iron such as aluminium, copper and zinc[18]. Non-ferrous metals production is highly electricity intensive as it requires electrolysis to convert raw materials (ore) to the final metal. Most carbon emissions come from the onsite generation of power.

The category paper covers both industrial and commercial production of various types of paper[19]. Paper production starts with manufacturing paper pulp most often from wood. The process is relatively energy intensive and often achieved by on-side generation of heat and electricity from burning fossil fuels.

Power covers electricity generation[20]. Standalone power plants generating electricity for the whole market are straightforward and are included in this category. Power plants onsite of industrial facilities are more difficult to classify and it is not clear that the EU ETS treats these consistently[21]. Therefore, I am restricting the power category to plants that belong to firms whose main NACE industry classification is power generation. I assign other power plants in category other.

The category refineries is fairly straightforward[22]. Refining uses fractional distillation which is highly energy intensive. Crude oil is heated and refined products are

---

[17]EU ETS categorizes a Hungarian firm focusing on mineral wool as a glass producer.

[18]EU ETS activities: "Production of primary aluminium", "Production of secondary aluminium", and "Production or processing of non-ferrous metals".

[19]EU ETS activity types: "Production of paper or cardboard","Industrial plants for the production of (a) pulp from timber or other fibrous materials (b) paper and board","Production of pulp"

[20]Includes EU ETS activity types: "Power", "Combustion of fuels", "Combustion installations with a rated thermal input exceeding 20 MW".

[21]I am aware of mistakes such as categorizing the only Hungarian refinery as a "Combustion of fuels" plant.

[22]EU ETS categories:"Refineries","Mineral oil refineries" and "Refining of mineral oil"

differentiated by their boiling points. Most refineries produce the required heat by burning oil or coal on-site which is the main source of carbon emissions.

Steel includes the production of coke, steel and other ferrous metals (alloys of iron) such as stainless steel[23]. Coke is a key ingredient to steel production and is often produced at the same plant as steel. Coke is produced by heating coal to a very high temperature in the absence of air. Ferrous metals are produced in furnaces at very high temperatures. Most carbon emissions come from generating the heat required to melt the raw materials.

Finally, other covers the category type other, all missing activity types and onsite power generation. Aviation is not included in the EU ETS accounts data and therefore neither in my dataset.

A.2.2 Activities and NACE Classifications

In this section, I provide details on how my activity classification corresponds to standard industry classifications. I pick NACE Rev. 2 as this is the primary classification system Eurostat uses[24]. The key takeaway of the analysis is that my activity classification resemble well 3-digit NACE codes. NACE codes are assigned to a firm "... according to its principal economic activity. The principal activity is the activity which contributes most to the value added of the unit." (Eurostat, 2008, p.29). Table A.2 shows the most common primary NACE3 industry classification code and its share. Activity categories are fairly homogenous: for most activities, the most common NACE3 code represents around or above 70% of firms. The standard complication of industry classifications is of vertically integrated firms. Vertical integration is likely responsible for most of the remaining codes. A cement producer might mine

[23]EU ETS activity types:"Production of pig iron or steel","Installations for the production of pig iron or steel (primary or secondary fusion) including continuous casting","Production or processing of ferrous metals","Metal ore roasting or sintering","Metal ore (including sulphide ore) roasting or sintering installations","Coke ovens" and "Production of coke". Coke could be included with refineries as the distillation process is not unlike oil refining.

[24]NACE: Nomenclature des Activités Économiques dans la Communauté Européenne. The results using NAICS (North American Industry Classification System) are similar.

limestone and manufacture ready-to-mix concrete. In the data, the second and third most common NACE3 codes for cement are "Quarrying of stone, sand and clay" and "Manufacture of concrete products for construction purposes" with shares of 7% and 5% respectively. For steel, the remaining codes are several categories of processing steel (casting, manufacturing of pipes, etc) which similarly fits the vertical integration logic. That is the key reason I am using activity categories that can span a vertically integrated industry. The only outlier is power, which represents any power generating unit that is likely located at manufacturing sites in my data. Power NAICS codes represent a wide array of industries in my dataset.

Table A.2: Most Common NACE3 Primary Codes by Activity

Activity	Description	NACE3	Share
Cement	Cement, lime and plaster	235	79
Ceramics	Clay building materials	233	74
Chemicals	Basic chemicals, fertilisers and nitrogen compounds...	201	77
Glass	Glass and glass products	231	78
Non-ferrous metals	Basic precious and other non-ferrous metals	244	71
Other		353	23
Paper	Pulp, paper and paperboard	171	71
Power	Electric power generation, transmission and distribution	351	100
Refineries	Refined petroleum products	192	75
Steel	Basic iron and steel and of ferro-alloys	241	57

Table A.3 shows the most common primary NACE4 industry classification code and its share. My activity categories are less homogenous in this case but the logic of vertical integration being mostly responsible for the differences holds similarly.

Table A.3: Most Common NACE4 Primary Codes by Activity

Activity	Description	NACE4	Share
Cement	Cement	2351	40
Ceramics	Bricks, tiles and construction products, in baked clay	2332	55
Chemicals	Other organic basic chemicals	2014	24
Glass	Hollow glass	2313	37
Non-ferrous metals	Aluminium	2442	38
Other	Steam and air conditioning supply	3530	23
Paper	Paper and paperboard	1712	60
Power		3511	100
Refineries	Refined petroleum products	1920	75
Steel	Basic iron and steel and of ferro-alloys	2410	57

A.3 Model Details

A.3.1 Intermediate Input Demand and Control Functions

First, I derive a formula for the intermediate input demand function from the static profit maximization of the materials choice. I assume that productivity is factor neutral, there is no uncertainty, materials choice is static, capital, labor and abatement are already chosen. Not all of these are necessary but they make the exposition simple. For deriving the control function it is easier to start from the original system of two equations: the production function and the abatement technology separately. Under the assumptions of the model with perfect competition, material choice satisfies:

$$\max_{M_{it}} \Pi_{it} = P_{it} f(K_{it}, L_{it}, M_{it}) - \sum P_{it}^X X_{it} - \tau_{it} E_{it},$$

where X refers to inputs {K,L,M}, P_{it}^X is input price. Emissions are determined as a function of abatement and all other input choices: $E_{it} = g(a_{it}) f(K_{it}, L_{it}, M_{it})$. Profit maximization leads to the following first-order condition:

$$(P_{it} - \tau_{it} g(a_{it})) \frac{\partial f(\cdot)}{\partial M} = P_{it}^M.$$

The expression shows that materials demand depends on emission prices and abatement decisions (or emissions). The general expression is complicated but it is possible to arrive at an expression for the transformed productivity term $\hat{\Omega}_{it}$ as a function of K, L, M, E, P^M, P and τ. The functional form depends on the production function $f(\cdot)$ and the abatement technology $g(\cdot)$.

Next, I derive the expression for the Cobb-Douglas case with a constant elasticity abatement technology.

$$\Pi_{it} = P_{it}(1 - a_{it}) \Omega_{it} K_{it}^{\alpha_K} L_{it}^{\alpha_L} M_{it}^{\alpha_M} - \sum P_{it}^X X_{it} - \tau_{it} E_{it}$$

where $E_{it} = B_{it}(1 - a_{it})^{\frac{1}{\alpha_E}} \Omega_{it} K_{it}^{\alpha_K} L_{it}^{\alpha_L} M_{it}^{\alpha_M}$. The first order condition of materials choice is:

$$P_{it} \alpha_m (1 - a_{it}) \frac{Y_{it}}{M_{it}} - \tau_{it} B_{it} (1 - a_{it})^{\frac{1}{\alpha_E}} \frac{Y_{it}}{M_{it}} = P_{it}^M$$

Then, I have to rewrite this to the transformed production function that has emissions as input. After taking logs and reorganizing:

$$\tilde{\omega}_{it} = log\left(P_{it}^M + \tau_{it}\frac{E_{it}}{M_{it}}\right) - log(P_{it}) - log(\alpha_m) - \alpha_k k_{it} - \alpha_l l_{it} - \alpha_e e_{it} - (\alpha_m - 1)m_{it}$$

where $\tilde{\omega}_{it} = (1-\alpha_E)\omega_{it} - \alpha_E log(B_{it})$. This it the familiar Cobb-Douglas materials control function with additional terms to correct for materials choice's impact on emission costs. What is key to observe here is that emission prices τ_{it} do not enter linearly and complicate the otherwise linear functional form. A first-order Taylor approximation of the first term and reorganizing the m_{it} terms leads to:

$$\tilde{\omega}_{it} \approx \tau_{it}\frac{E_{it}}{P_{it}^M M_{it}} - log(P_{it}) - log(\alpha_m) - \alpha_k k_{it} - \alpha_l l_{it} - \alpha_e e_{it} - (\alpha_m)m_{it}$$

I do not observe τ_{it} and only have proxies for it. I observe the market price of emission allowances, allowance allocations and the stock of allowances each firm holds in the beginning of the year. The functional form of how effective emission prices depend on these is not clear. Suppose $\tau_{it} = \tau_t + \delta\hat{\tau}_{it}(Allocation, Emission, Bank)$ Then, I arrive at a functional form that I can use for estimation:

$$\tilde{\omega}_{it} \approx \tau_t\frac{E_{it}}{P_{it}^M M_{it}} + \delta\hat{\tau}_{it}\frac{E_{it}}{P_{it}^M M_{it}} - log(P_{it}) - log(\alpha_m) - \alpha_k k_{it} - \alpha_l l_{it} - \alpha_e e_{it} - (\alpha_m)m_{it} \quad \text{(A.1)}$$

A.3.2 Imperfect Competition Case

This section discusses conditions under which my empirical strategy is valid with imperfect competition and unobserved prices. I also detail how the interpretation of some of the estimates are different in this case. I follow De Loecker (2011) and Asker, Collard-Wexler and De Loecker (2014) and introduce a CES demand system with monopolistic competition. Suppose the log production function is

$$Q_{it} = exp(\omega_{it} + u_{it})K_{it}^{\alpha_k} L_{it}^{\alpha_l} M_{it}^{\alpha_m} E_{it}^{\alpha_e}$$

and the demand for a firm's product has constant elasticity:

$$Q_{it} = P_{it}^{-\rho} * exp(\xi_{it})$$

where ρ is a parameter of demand elasticity and ξ_t are idiosyncratic demand shocks. Note that I could also introduce aggregate demand shifters or allow demand to depend on prices of competitors such as in De Loecker (2011). We can use the demand system to obtain an expression for price and arrive at revenue production functions of the form:

$$r_{it} = \omega_{it}^* + \alpha_k^* k_{it} + \alpha_l^* l_{it} + \alpha_m^* m_{it} + \alpha_e^* e_{it} + \xi_{it}^* + u_{it}.$$

The new input elasticity parameters measure the combined impact of production and demand: $\alpha_h^* = \alpha_h[1-\rho^{-1}]$ for $h = k, l, m, e$. Similarly, the new productivity term picks up both demand and production effects $\omega_{it}^* = [1-\rho^{-1}]\omega_{it}$ and demand shocks need to be scaled according to demand elasticities $\xi_{it}^* = \rho^{-1}\xi_{it}$. Alternatively, we can combine ω_{it}^* and ξ_{it}^* into a joint term $\tilde{\omega}_{it}$. In this case, $\tilde{\omega}_{it}$ picks up a weighted sum of productivity and demand shocks where the weights are determined by the demand elasticity.

Next, I derive material demand assuming materials are decided last and static optimization in materials hold. Similar equations hold with slightly different timing assumptions as well. The firms solves:

$$\max_{M_{it}} \Pi_{it} = e^{\omega_{it}^* + \xi_{it}^*} K_{it}^{\alpha_k^*} L_{it}^{\alpha_l^*} M_{it}^{\alpha_m^*} E_{it}^{\alpha_e^*} - \sum P_{it}^X X_{it}.$$

The resulting intermediate input demand equation takes the following form:

$$m_{it} = (1-\alpha_m^*)^{-1}\left[\omega_{it}^* + \xi_{it}^* + \alpha_k^* k_{it} + \alpha_l^* l_{it} + \alpha_e^* e_{it} + log(\alpha_m^*) - log(P_{it}^m)\right].$$

This shows that intermediate input demand is monotonic in the new productivity term ω_{it}^* as well as its components. The control function based on the inverse of this equation can only control for ω_{it}^* even the best case when there is no variation in demand shocks and input prices. When output is measured in revenues is controlling for output prices (or its determinants when unobserved is correct), since output prices are also endogenous to input choice.

A.3.3 Abatement Choice, Emissions and Productivity

This section provides an illustrative proof for why emissions are correlated with unobserved productivity. I assume that productivity is factor neutral, there is no uncertainty, abatement choice is static and other inputs are already chosen. These are not necessary but make the illustrative example simple. With these assumptions, abatement expenditures are chosen to maximize per period profits:

$$\Pi_{it} = P_{it}(1 - a_{it})f(K_{it}, L_{it}, M_{it}) - \sum P_{it}^{X} X_{it} - \tau_{it} E_{it}$$

where $E_{it} = g(a, \cdot)f(K_{it}, L_{it}, M_{it})$. The first order condition is

$$\frac{\partial g(a, \cdot)}{\partial a} = \frac{P_{it}}{\tau_{it}}$$

In this simple example, the abatement choice only depends on output and emission prices and the abatement technology. By assumption, abatement does not influence other input choices other than emissions. Therefore, neither productivity, nor the functional form of the production function influences the abatement decisions. However, emissions are still positively correlated with productivity (even conditional on other inputs choices) since they are explicitly a function of it. Notice that this would be the case even if emissions would not depend directly on Ω, since in this case the abatement choice would. When $g(a, \cdot) = B_{it}(1 - a)^{\frac{1}{\alpha_E}}$ emissions are:

$$E_{it} = B_{it}^{2\alpha_E} \tau_{it}^{-\frac{1}{1-2\alpha_E}} (P_{it}\alpha_E)^{\frac{1}{1-2\alpha_E}} f(K_{it}, L_{it}, M_{it})$$

This shows that emissions are increasing in emission intensity (B_{it}) even though abatement also increases in B_{it}. Emissions also increase in the emission input elasticity, output prices, productivity and other inputs.

A.3.4 Alternative Models

In this section, I show how my model can accommodate three alternative modeling assumptions. First, the interpretation of a_{it} as abatement stock is consistent with the model. Suppose $Y_{it} = (1 - a_{it})f(\cdot)$, $E_{it} = (1 - a_{it})^{\frac{1}{\alpha_E}} f(\cdot)$, but a_{it} is not interpreted as abatement stock. Total abatement expenditures in time period t are unchanged: $a_{it}\bar{Y}_{it}$. Therefore, the marginal abatement costs (the derivative of the former in emissions) is unchanged as well.

I also show an alternative model, when my empirical method fails. Suppose that the production function is Cobb-Douglas and that $Y_{it} = (1 - a_{it})f(\cdot)$, $E_{it} = (1 - \bar{a}_{it})^{\frac{1}{\alpha_E}} f(\cdot)$. $\bar{a}_{it}$ is abatement expenditure stock which evolves to the following similar specification: $\bar{a}_{it} = a_{it} + \bar{a}_{it-1}$. After the inversion output depends on last year's abatement stock ($\bar{\alpha}_{it-1}$):

$$Y_{it} = f(\cdot)^{1-\alpha_E} E^{\alpha_E} + \bar{\alpha}_{it-1} f(\cdot).$$

Empirically, omitting $\bar{\alpha}_{it-1}$ from the production function results in biased elasticity estimates. Input choices and in particular E_{it} is likely negatively correlated with abatement stock. Therefore, $\bar{\alpha}_E$ is likely to be underestimated (the direction of the bias is not that clear since the functional form is not additive in logs.). Notice however, that the interpretation of α_E is different in this model since it measures the impact of abatement stock and not just this year's abatement expenditure. For the inversion to work and for emissions to proxy for abatement expenditures there must be a one-to-one mapping between a_{it} and E_{it}.

Second, I assume that the abatement technology only requires capital. I show that in this case, my empirical approach is valid and my model estimates can be interpreted as a function of the parameters of this new model. The key to these results is that capital is an observed stock measure and hence its measurement is compatible with a stock interpretation of abatement. Consider the following model with capital only abatement technology: $Y_{it} = f((1 - a_{it})K_{it}, L_{it}, M_{it})$ and $E_{it} = (1 - \bar{a}_{it})^{\frac{1}{\alpha_E}} f(K_{it}, L_{it}, M_{it})$. Then, $a_{it}K_{it}$ is capital allocated for abatement. As such, abatement capital is indistinguishable from ordinary capital in the beginning of the period. Although, this might be a strong assumption, the idea that a_{it} can be thought of as the share of abatement capital is useful. The key here is that capital is a stock measure with dynamic implications and therefore the abatement technology inherits these characteristics. When a_{it} increases the change additional abatement capital decreases emissions and not the change in abatement expenditures investment to abatement capital. Notice that the one to one mapping in this case is between a_{it} (the share of abatement capital) which is a stock measure. As such, it can capture the persistent impact of abatement expenditures. This model leads to the following production function:

$$Y_{it} = (\Omega_{it}^{1-\alpha_E\alpha_K} B_{it}^{-\alpha_E\alpha_K}) K_{it}^{\alpha_K(1-\alpha_E\alpha_K)} L_{it}^{\alpha_L(1-\alpha_E\alpha_K)} M_{it}^{\alpha_M(1-\alpha_E\alpha_K)} E_{it}^{\alpha_E\alpha_K}$$

The functional form is the same and my empirical method produces consistent estimates of the production function. The transformed estimates can yield drastically different estimates of α_E. Note that the interpretation of the estimates is also different. The elasticity of emissions intensity to abatement capital and to abatement expenditure are different. Especially, since abatement capital measures the effects of a stock measure, the estimates are likely higher. Allocating an a_{it} share of the capital input gives up less output (how much less depends on α_k). However, to allocate a similar amount in monetary terms, the firm would need to allocate a much higher share of capital than output. Therefore, the two models measure the same underlying phenomenon. The interpretation of the capital only model might be more reasonable for some circumstances.

Third, assume emissions do not increase linearly with potential output: $Y_{it} = (1 - a_{it})f(\cdot)$ and $E_{it} = (1 - a_{it})^{\frac{1}{\alpha_E}} f(\cdot)^{\alpha_Y}$. When $\alpha_Y = 1$, we get back the linear case. With α_Y above 1, emissions are increasing faster than linear in output. The reason it is important is because in this case, there is an additional channel through which technology effects emissions. Any methodology that is trying to separate the impact of abatement expenditures from technology have to account for this. The transformed production function takes the same Cobb-Douglas form: $Y = \Omega^{1-\alpha_E\alpha_Y} B^{-\alpha_E} E^{\alpha_E} (K^{\alpha_K} L^{\alpha_L} M^{\alpha_M})^{1-\alpha_E\alpha_Y}$. The input elasticities are slightly different. When With α_Y is above one, relative importance of emissions is higher in the final production function. The overall implications are similar.

A.3.5 Marginal Product of Emission is Marginal Abatement Costs

In this section, I show that the marginal product of emissions in the transformed production function is equivalent to the marginal cost of abatement for the Cobb-Douglas case. First, notice that abatement expenditures in this model are measured by $a\Omega f(\cdot)$. From the abatement technology (equation 2.2) we can derive the following expression:

$$af(\cdot) = f(\cdot) - B^{-\alpha_E} E^{\alpha_E} (f(\cdot))^{1-\alpha_E} = f(\cdot) - B^{-\alpha_E} E^{\alpha_E} (f(\cdot))^{1-\alpha_E}$$

Notice that emission influence abatement expenditures only through the second term which is the transformed production function $(f(\cdot) - \bar{\Omega} K^{\bar{\alpha}_K} L^{\bar{\alpha}_L} M^{\bar{\alpha}_M} E^{\alpha_E})$. Taking the partial derivative of $af(\cdot)$ in E measures the marginal abatement costs. The above shows that this is equivalent to -1* the marginal product of emissions in the transformed production function.

A.4 Vertical Integration and EU ETS Coverage

The EU ETS regulates emissions at the plant level for certain industries and above a certain level of activity. Firms that own these plants likely have other activities that make the firm level measurement inaccurate. The firm variables also include other activities such as an unregulated plant in the same industry or plants in other industries. It is likely that output and inputs are generally higher than the aggregates of EU ETS plants the firms own. For instance, a vertically integrated cement-concrete firm has higher revenues and other inputs for the same emissions as a similarly productive non-vertically integrated firm. Alternatively, emissions are underestimated for vertically integrated firms and for those active in other industries. The key to understand the potential bias this might cause in the coefficient estimates is to realize that essentially this is an issue of measurement error in emissions. Suppose we observe $\bar{e}_{it} = e_{it} + \nu_{it}$, where e_{it} is log emissions and ν_{it} is measurement error. At best, when the measurement error is classical it causes to coefficient estimates of the emission input elasticity attenuated towards zero. However, ν_{it} is likely to be negatively correlated with e_{it}. Since vertically integrated firms they tend to be bigger (have additional activities), they would have had higher emissions if they were a similar sized non-vertically integrated firm[25]. Therefore, ν_{it} is likely a high negative value for these high emissions firms (again compared to the hypothetical if they were the same size but non-integrated firm). In the followings, I show that the sign of the bias of $\hat{\beta}_e$ is not clear even in the single-variable OLS regression of output on emissions. The bias in this case is:

$$plim\hat{\beta}_e = \left(1 - \frac{\sigma_\nu^2 + \sigma_{e\nu}}{\sigma_e^2 + \sigma_\nu^2 + 2\sigma_{e\nu}}\right)\beta_e$$

In contrast to the classical measurement error case ($\sigma_{e\nu} = 0$), a negative $\sigma_{e\nu}$ can result in an overestimate of a positive β_e. $|\sigma_{e\nu}| > \sigma_\nu^2$ when the correlation of e and ν is high and the standard deviation of e is much higher than of ν. This is not unlikely, and in the extreme case, can result in a strong upward bias of the coefficient estimate. When the measurement error is correlated with emissions is when the potential bias for β_e is likely to be the most severe. However, ν might also be correlated with other inputs or unobserved productivity. In this case, other coefficient estimates are likely to have a stronger bias. Again, the directions of this bias are unclear. This exercise shows that the bias due to mismeasurement in my case is unclear even in an simple OLS setup. Showing the sign of the bias in my production function estimation method is much more difficult. In any case, it is likely that the sign or magnitude of the bias is not possible to determine ex-ante. that instrumenting for e_{it} with e_{it-1} does not resolve this issue, since e_{it-1} is likely to be similarly mismeasured for vertically integrated. Collard-Wexler and De Loecker (2016) proposes a method to deal with classical measurement error in inputs in production function estimation which requires a second, uncorrelated measure. Unfortunately, in my setup this is hard to come by since the real measurement is a hypothetical (using allocations has the same issue as emissions).

A.5 Robustness

This section reports the results of several robustness exercises. The key results are discussed in detail in Section 2.5.3. First, I address the robustness of my elasticity

[25] Alternatively, we observe firms that should not be in the data that have emissions measured too low.

and MAC estimates to measurement error due to vertical integration. I restrict the sample to firms for whom the vertical integration issue is less likely to be serious. Table shows my elasticity estimates are similar across different specifications. Table A.5 reports the distribution of marginal abatement costs across the restricted samples. Table A.6 shows that the industry rank of MACs is essentially unchanged across the specifications.

Table A.4: Robustness of Elasticity Estimates to Measurement Error

	Standard	Mode	Mode+	ETS	Narrow	Single	Small
Elasticity							
β_k	0.146	-0.015	-0.022	-0.008	-0.017	0.002	-0.000
β_l	0.315	0.039	0.037	0.029	0.031	0.012	-0.002
β_m	0.455	-0.018	0.004	-0.018	-0.001	-0.018	-0.005
β_e	0.068	0.007	0.004	0.005	0.004	0.012	0.007
Elasticity - absolute							
$\lvert\beta_k\rvert$	0.146	0.018	0.024	0.019	0.036	0.014	0.003
$\lvert\beta_l\rvert$	0.315	0.042	0.050	0.035	0.069	0.018	0.009
$\lvert\beta_m\rvert$	0.455	0.038	0.045	0.035	0.042	0.019	0.008
$\lvert\beta_e\rvert$	0.068	0.011	0.018	0.011	0.019	0.012	0.008
Standard error							
se β_k	0.025	0.008	0.008	0.004	0.007	0.005	0.004
se β_l	0.042	0.008	0.010	0.009	0.016	0.007	0.001
se β_m	0.036	0.007	0.015	0.009	0.013	0.008	0.003
se β_e	0.016	0.004	0.010	0.003	0.005	0.003	0.001
Scale and persistence							
scale	0.984	0.013	0.023	0.007	0.018	0.008	-0.000
ω_{it} persistence	0.907	0.003	-0.010	0.014	-0.817	0.010	-0.004
Correlation							
β_e	1.000	0.894	0.880	0.954	0.874	0.972	0.995
Observations	18338	10861	6711	11210	17761	13370	18338

Notes. An observation is a firm-year. Firms from activity other are excluded. Years: 2005-2016 Standard: Cobb-Douglas with linear materials control function, high level activity aggregation, means reported. All other specifications show the mean of the individual differences to standard estimates.
Mode: Firms with modal NACE4 code only. Mode+: Mode and no other industry codes. Mode ETS: Firms with modal EU ETS mainactivity. Narrow: estimated at the EU ETS mainactivity level. Single: single plant firms only.

Second, I turn to the robustness of my elasticity estimates to different control functions. I estimate several specifications when I use different proxies and functional forms to control for emission prices (τ_{it}) in equation A.1. I experiment with different ways to proxy the first and second terms of equation A.1. I do not directly observe these terms and it is not ex-ante clear what the best proxies are. "ME1" I uses variables $\tau_1 = \tau * e/m$ and $\tau_2 = log(A/E)E/M$ as proxies for the first and second

Table A.5: Robustness of Marginal Abatement Cost Estimates

	Standard	Mode	Mode+	ETS	Narrow	Single	Small	Poly2
Λ_{it} ***distribution***								
25th percentile	37	32	42	40	42	51	45	77
Median	90	68	70	108	156	115	104	173
75th percentile	239	189	152	277	587	314	271	468
90th percentile	1045	1198	628	1037	2178	1535	1177	1632
Observations	18338	10861	6618	11210	35109	13370	18338	18338

Notes. An observation is a firm-year. Firms from activity other are excluded. Years: 2005-2016. Λ_{it} is measured in €/ton CO_2 equivalent. Λ_{it} represents the expected marginal abatement cost for firm i in time t at the observed level of output and emissions. Standard: Cobb-Douglas with linear materials control function, means reported. Mode: Firms with modal NACE4 code only. Mode+: Mode and no other industry codes. ETS: Firms with modal EU ETS mainactivity. Narrow: estimated at the EU ETS mainactivity level. Single: single plant firms only. Poly2: standard but control function is second-order polynomial.

Table A.6: Median MAC industry Ranking in Different Specifications

method	low1	low2	high2	high1
cd_m	Cement	Power	Refineries	Non-ferrous metals
cd_m_ind_ets_m	Cement	Power	Refineries	Non-ferrous metals
cd_m_ind_mod	Cement	Power	Chemicals	Non-ferrous metals
cd_m_ind_mod_r	Power	Cement	Non-ferrous metals	Chemicals
cd_m_mainact	Refineries	Cement	Glass	Ceramics
cd_m_poly2	Cement	Power	Refineries	Non-ferrous metals
cd_m_small	Cement	Power	Refineries	Non-ferrous metals
cd_m_sp	Cement	Chemicals	Non-ferrous metals	Refineries

Methods are Standard, ETS, Mode, Mode+, Narrow, Poly2, Small and Single plant. Industries are ranked by median marginal abatement costs (Λ_{it}) Low1: lowest, low2: 2nd lowest, high2: 2nd highest, high1: highest.

term respectively. "ME2" includes the emission price τ_t and allocation level a_{it} to simplify the functional form. "ME3" includes the emission price τ_t and the size of the allowance bank b_{it}. Again, b_{it} proxies for emission prices. "ME4" replaces b_{it} with $log(B_{it}/E_{it})$. Finally, "ME5" includes τ_1 and $log(B_{it}/E_{it} * E/M$ as proxies. Table A.7 reports the difference of the different emission control proxy cases to the main specification. Generally, the estimates are very similar. The two specifications that are the most different are likely to be misspecified (ω_{it} persistence is too low).

Finally, Table A.8 show the robustness of my elasticity estimates to using a second-order polynomial control function. MAC estimates for this specification are reported in Table A.5.

Table A.7: Robustness to Emission Prices in the Control Function

	CD M proxy	ME1	ME2	ME3	ME4	ME5
Elasticity						
β_k	0.146	-0.010	-0.002	0.000	0.004	-0.016
β_l	0.315	-0.025	-0.010	-0.025	0.010	0.013
β_m	0.455	0.033	0.020	0.010	-0.025	-0.008
β_e	0.068	-0.022	-0.002	-0.016	0.015	0.003
Elasticity - absolute						
$\lvert\beta_k\rvert$	0.146	0.012	0.014	0.014	0.009	0.018
$\lvert\beta_l\rvert$	0.315	0.027	0.040	0.031	0.043	0.034
$\lvert\beta_m\rvert$	0.455	0.035	0.026	0.026	0.039	0.052
$\lvert\beta_e\rvert$	0.068	0.025	0.021	0.021	0.032	0.021
Standard error						
se β_k	0.025	0.011	0.006	0.001	0.004	0.030
se β_l	0.042	0.046	0.010	0.014	0.013	0.053
se β_m	0.036	0.053	0.013	0.011	0.010	0.063
se β_e	0.016	0.040	0.018	0.015	0.016	0.055
Scale and persistence						
scale	0.984	-0.024	0.006	-0.030	0.005	-0.008
ω_{it} persistence	0.907	-0.252	-0.082	0.026	-0.060	-0.194
Correlation						
β_e	1.000	0.538	0.925	0.924	0.824	0.798
Observations	18338	18338	18338	18338	18338	18338

Notes. An observation is a firm-year. Firms from activity other are excluded. Years: 2005-2016.

CD M proxy: Cobb-Douglas with linear materials control function, means reported. All other specifications show the mean of the individual differences to CD M proxy estimates. ME: different specifications with materials proxy controlling for emission prices.

Table A.8: Production Function Estimates - Second Order Polynomial Control

	(1) Cem	(2) Cer	(3) Chem	(4) Glass	(5) Non-ferr	(6) Oth	(7) Paper	(8) Power	(9) Ref	(10) Steel
Other input elasticities										
β_l	0.493	0.411	0.150	0.148	-0.057	0.343	0.357	0.051	0.299	0.303
	(0.167)	(0.060)	(0.177)	(0.117)	(0.150)	(0.067)	(0.067)	(0.088)	(0.238)	(0.066)
β_m	0.277	0.311	0.496	0.608	0.721	0.405	0.362	0.396	0.199	0.617
	(0.136)	(0.040)	(0.188)	(0.100)	(0.092)	(0.028)	(0.074)	(0.053)	(0.171)	(0.048)
β_k	0.079	0.092	0.179	0.116	0.113	0.146	0.171	0.284	0.070	0.051
	(0.032)	(0.020)	(0.072)	(0.024)	(0.060)	(0.022)	(0.030)	(0.053)	(0.092)	(0.021)
Emission elasticity										
β_e	0.090	0.166	0.027	0.037	0.092	0.028	0.142	0.038	0.280	0.011
	(0.088)	(0.047)	(0.092)	(0.086)	(0.088)	(0.033)	(0.040)	(0.040)	(0.095)	(0.040)
Returns to scale										
$\beta_k + \beta_l + \beta_m + \beta_e$	0.938	0.981	0.853	0.909	0.869	0.922	1.032	0.769	0.848	0.982
	(0.072)	(0.041)	(0.204)	(0.067)	(0.096)	(0.073)	(0.035)	(0.073)	(0.144)	(0.037)
[1em] ***Productivity***										
Persistence	0.889	0.828	0.836	0.838	0.895	0.892	0.793	0.800	0.940	0.707
Median	2.905	2.478	3.989	2.816	3.721	2.816	1.729	4.964	6.452	1.508
Inter-quartile range	0.429	0.509	0.453	0.383	0.440	0.559	0.320	0.919	1.785	0.237
Observations	1713	4534	901	1953	274	17674	3988	2684	506	1785

Notes. An observation is a firm-year, all observations (unbalanced panel) included, 2005-2016. Block bootstrapped (firm level) standard errors reported in parenthesis, 100 iterations. Control function is second-order polynomial in log inputs, emission price controls not included. Instruments: $\{k_{it}, k_{it-1}, k_{it-2}, l_{it-1}, l_{it-1}, m_{it-1}, m_{it-2}, e_{it-1}, e_{it-2}\}$. Estimated productivity $\hat{\omega}_{it}$ calculated as $\hat{\phi}(\cdot) - \hat{f}(\cdot)$, does not contain measurement error ϵ_{it}. Persistence is calculated as the regression coefficient on $\hat{\omega}_{it} = \rho\omega_{it-1} + \nu_{it}$.

A.6 Marginal Revenue Product Comparison

I calculate the dispersion of marginal revenue product (MRP) of all inputs in my main specification. Table A.9 reports two measures of dispersion across firms within a year. MRPE shows the most heterogeneity. Table A.10 shows within firm through time dispersion. Emissions vary more than any other input within firms through time. This supports an explanation of possible high adjustment costs combined with long-term high price expectations hindering firms increasing emissions.

Table A.9: Dispersion in Marginal Products - 2015

	Cem	Cer	Chem	Glass	Non-ferr	Paper	Power	Ref	Steel	All
St. dev										
K	0.75	0.91	0.92	0.75	1.15	0.93	1.63	1.48	1.28	1.06
L	0.34	0.41	1.10	0.50	1.56	0.61	1.55	1.35	0.71	0.77
M	0.50	0.51	0.86	0.36	0.48	0.36	1.00	1.13	0.37	0.56
E	1.00	1.06	1.61	0.96	1.69	1.58	2.44	1.98	1.25	1.47
Iqr										
K	0.91	1.14	1.00	0.94	1.12	1.00	1.70	1.52	1.26	1.17
L	0.46	0.44	1.01	0.68	1.48	0.55	1.55	1.69	0.86	0.81
M	0.52	0.52	0.71	0.46	0.75	0.28	1.00	1.89	0.46	0.59
E	0.81	1.38	2.41	0.97	2.85	1.65	3.01	2.96	1.14	1.72
Observations	137	326	107	166	42	319	237	44	162	1540

Notes. Category other excluded. Main specification. Iqr: interquartile range.

Table A.10: Dispersion in Marginal Products Through Time

	Cem	Cer	Chem	Glass	Non-ferr	Paper	Power	Ref	Steel	All
St. dev										
K	0.31	0.38	0.35	0.38	0.30	0.37	0.46	0.54	0.39	0.39
L	0.16	0.21	0.57	0.18	0.22	0.18	0.48	0.37	0.24	0.26
M	0.25	0.31	0.32	0.15	0.11	0.14	0.37	0.31	0.13	0.24
E	0.37	0.55	0.65	0.38	0.57	0.52	0.72	0.47	0.44	0.52
Iqr										
K	0.20	0.29	0.17	0.22	0.21	0.19	0.23	0.20	0.26	0.23
L	0.09	0.14	0.16	0.10	0.19	0.11	0.23	0.18	0.14	0.14
M	0.13	0.19	0.15	0.10	0.09	0.09	0.19	0.12	0.08	0.14
E	0.17	0.28	0.32	0.14	0.42	0.17	0.34	0.19	0.21	0.23
Observations	1586	4093	855	1805	267	3668	2527	472	1670	16943

Notes. Category other excluded. Main specification. Iqr: interquartile range.

Appendix B

Appendix to Chapter 3

B.1 Background - Capacity and Ancillary Services Markets

The system operator also has a responsibility for the security and stability of the electricity system on the long run. Not only is it required to match supply and demand today, ISOs also have to set up incentives such that demand can be satisfied years ahead. ISOs usually plan to be able to match maximum demand plus a reserve margin. Reserve margins meant to provide an insurance against unexpected power plant shut downs, congestion or changes in demand. The socially optimal level of reserve margins balances the cost of additional capacity to the benefits of less frequent shortages.

Energy-only markets require power plants to recover their entry costs from spot market profits alone. The key idea is that when capacity is low profits are high resulting in scarcity rents. With free entry scarcity rents drop to a level of recovering entry costs. Lack of distortions, energy-only markets are capable of incentivizing the socially optimal investment (Hogan et al. (2005)). In real-world energy markets there are several distortions that result in a "missing money" problem. Lack of demand elasticity, inefficient scarcity pricing, market power mitigation (offer and price caps) can all result in socially unoptimal scarcity rents and entry[1].

The idea of Installed Capacity market designs (ICAP) is to complement scarcity rents such that the new equilibrium capacity is optimal. The implementation is market based. Retailers are required to procure capacity up to their peak demand plus the reserve margin. Thus, power plant operators receives payments for holding capacity. Implementation varies across ISOs in whether the market is centralized, what exactly is required from supply receiving payments, whether there is a forward market and whether only new entry is eligible to set prices. The key is that capacity payments are determined in equilibrium. Hence, the most efficient technologies and operators receive capacity payments. Of course, there is a question of how competitive real world capacity markets are. Schwenen (2015) shows that capacity markets in New York were best described by the tacit collusive equilibrium of Fabra, von der Fehr and Harbord (2006). In contrast, PJM mitigates market power in the capacity market strongly and as a result it is fairly competitive Allcott (2012*a*).

There is an inherent trade-off between mitigating market power and incentivizing entry. This trade-off is long observed in the economics literature[2]. Therefore there is a choice between an energy only market without much market power mitigation and market power mitigation with capacity markets. System operators around the world have different answers to this question. While Australia seems to be satisfied with its energy only market without price caps (Moran and Skinner (2008)), the ISO New England strongly believes in capacity markets. The answer to this question likely varies due to differences in fuel access and prices and initial system conditions.

B.2 Model - Computation Algorithm

This section provides a high level review of my computational algorithm.

[1]There are additional reasons, see Cramton and Stoft (2005). Allcott (2012*a*) describes scarcity pricing in a particularly clear way. Fabra (2018) presents a model with inelastic demand. She shows that energy-only markets are not able to achieve socially optimal investments.

[2]For instance, for incentivizing innovation by awarding patents among others.

1. Set up parameters and exogenous fuel path $\{z_t\}$.

2. Guess K1:

   (a) Guess a sequence of $\{K_2, K_3, ..., K_T\}$

   (b) For t=T calculate Π_T (long run spot profits). This determines V_T.

   (c) For t=T-1 calculate Π_{T-1}. V_T and Π_{T-1} determines exit thresholds Φ_{T-1}. Then calculate V_{T-1}.

   (d) Iterate backwards until t=1.

   (e) Predict capacity sequence $\{\hat{K}_1, \hat{K}_2, ..., \hat{K}_T\}$. Check whether it is close enough to $\{K_2, K_3, ..., K_T\}$.

   (f) Iterate until convergence: not a contraction mapping. If $\{\hat{K} < K\}$, the model predicts too much exit given the guess. Means guess capacity sequence did not decrease quickly enough, so need to modify toward the prediction.

3. Check if entry condition holds for the first period. Iterate over until it does.

B.3 Empirical Approach - Spot Market Modeling

There is a substantial literature in modeling spot market competition in energy. The appropriate equilibrium concept is supply function equilibrium (Wilson (1979)). The uniform price auction is known to have multiple equilibria in many settings (Hortaçsu and McAdams (2018)). Klemperer and Meyer (1989) shows that when demand is uncertain and bidders have private values the equilibrium is unique. They also show that this supply function equilibrium lies in between the Bertrand and Cournot equilibrium[3]. The Klemperer and Meyer (1989) model fits power generation well as val-

[3]The literature explored the properties of supply function equilibrium under different conditions. Regardless of the multiplicity empirical papers use first-order conditions that can identify valuations under the assumption that the same equilibrium is played in different markets observed. See Hortaçsu and McAdams (2018) for a good summary.

uations (marginal costs) are relatively well known[4]. There is two general approaches empirical papers in energy take to model spot market competition. The first group of papers approximates the supply function equilibrium by estimating a capacity-constrained Cournot game. This is possible when bids are not observed and marginal costs are assumed to be known. The model is flexible enough to incorporate vertical contracts (Bushnell, Mansur and Saravia (2008)), sequential markets (Ito and Reguant (2016)) and it is easy enough to compute to allow for entry-exit decisions (Myatt (2017)). The second group of papers uses bid data and applies a supply-function equilibrium framework. Both marginal costs and vertical arrangements are identified using first-order conditions (Reguant (2014))[5]. Allcott (2012*b*) uses a supply function equilibria framework for a two-stage entry model. He uses only a few hours to reduce computation time. Finally, there is an additional family of models that assumes no uncertainty. Fabra, von der Fehr and Harbord (2006) describes a model that is essentially a capacity-constrained Bertrand game. They show that if overall capacity is large enough a competitive equilibrium exists. In this equilibrium price is limited by the marginal cost of the most efficient non-producing player. In the end, choosing an approach depends on data availability and the underlying competitiveness of the market. The overall impact of vertical contracts (increase), congestion (decrease) and offer capping (increase) on competitiveness is uncertain. Willems, Rumiantseva and Weigt (2009) finds that the Cournot model approximates the supply function equilibrium very well for the German electricity market. Myatt (2017) estimates a capacity constrained Cournot model with a fringe (assumes away congestion and contracts) for all US ISOs. His estimates show that the model is able to match data well for some but not all markets. In particular, it does not fit well for ERCOT (considered competitive) and New England (no congestion). Allcott (2012*a*)

[4]In contrast to treasury auctions, the other important multi-unit auction environment

[5]Hortaçsu and Puller (2008) shows vertical arrangements are identified if marginal cost are known. Wolak (2000) shows the reverse.

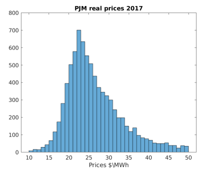

(a) Data, 2017

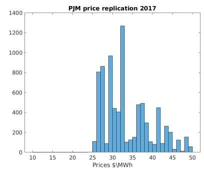

(b) Het. coal, 2017

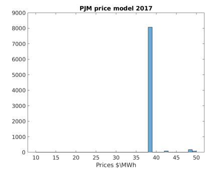

(c) Hom. coal, 2017

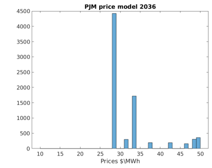

(d) Hom. coal, 2036

Figure B.1: Price Histograms, PJM

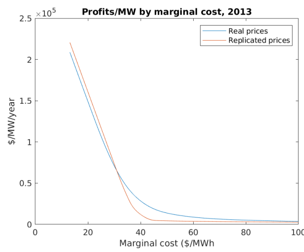

(a) 2013

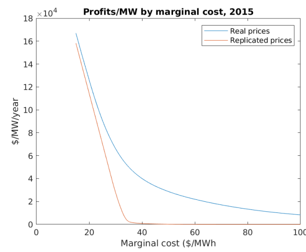

(b) 2015

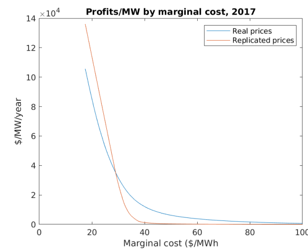

(c) 2017

Figure B.2: Real and Replicated Profits by Marginal Cost, PJM

assumes PJM markets are competitive (plants bid MC) as market power does not seem to influence his estimates significantly.

B.4 Empirical Approach - Spot Market Fit

Figure B.1 presents price variation from my spot market replication exercise. Figure B.1a presents the price distribution in the data in 2017. Figure B.1b shows the key result of my replication exercise when I use the full range of power plants in the market. I replace all coal plants with a representative plant and present replication results in Figure B.1c and d.

I also calculate profits under my replication exercise to real profits using prices in the data. Figure B.2 presents the results for different years.

B.5 Welfare

This section contains the details of the welfare calculations. I calculate yearly consumer costs as

$$ConsumerC_t = \sum_{h=1}^{8760} p_{th} * max(load_{th}, K_{th}),$$

where K_{th} denotes total system capacity of all power plants and p_{ht} is the wholesale price of electricity. I assume that distribution and transmission costs are independent of wholesale prices and output. Therefore, they are irrelevant to welfare calculations.

The total value of loss load (VOLL) is calculated as the difference between production and demand multiplied by the value of loss load:

$$VOLL_t = \sum_{h=1}^{8760} 1000 * -min(0, K_{th} - load_{th})$$

.

I calculate spot profits by the following formula:

$$\Pi_t = \sum_{i=1}^{n} \sum_{h=1}^{8760} [p_{th} - (mc_{it} + p_t^c * emi_{it} - \mathbb{1}_{gas} * ptc_y] * q_{ith} + EC_t - ES_t,$$

where i denotes units, p^c is carbon price, emi_{it} is emission intensity $(CO_2 t/MWh)$. q_{ith} is 0, max capacity or in between when the unit sets the price. ES_t is total exit subsidies[6].

Entry costs are

$$EC_y = max(0, K_y^{wind} - K_{y-1}^{wind}) * (CE_y^{wind} ens_y^{wind}),$$

where ens_y^{wind} is the per MW entry subsidy. This only calculates net entry and I assume away entry and exit that does not change capacities. This is relevant for

[6] Currently, implemented as a shift to the mean of the distribution of scrap values.

both the stationary and non-stationary equilibrium. There is exit in the model even though capacities are growing which results in higher entry.

Carbon costs are calculated as:

$$Carboncost_y = \sum_{i=1}^{n} \sum_{h=1}^{8760} scc_y * q_i yh * emi_{iy}$$

Finally, the government balance equals carbon tax minus entry and exit subsidies:

$$G_y = \sum_{i=1}^{n} \sum_{h=1}^{8760} p_y^c * q_{iyh} * emi_{iy} - max(0, K_y^c - K_{y-1}^c) * exs_y - max(0, K_y^w - K_{y-1}^w) * ens_y^w,$$

where exs_y is the per MW exit subsidy.

B.6 Results - Dynamics

Figure B.3 presents detailed market outcomes for capacities, prices, emissions, generation and profits.

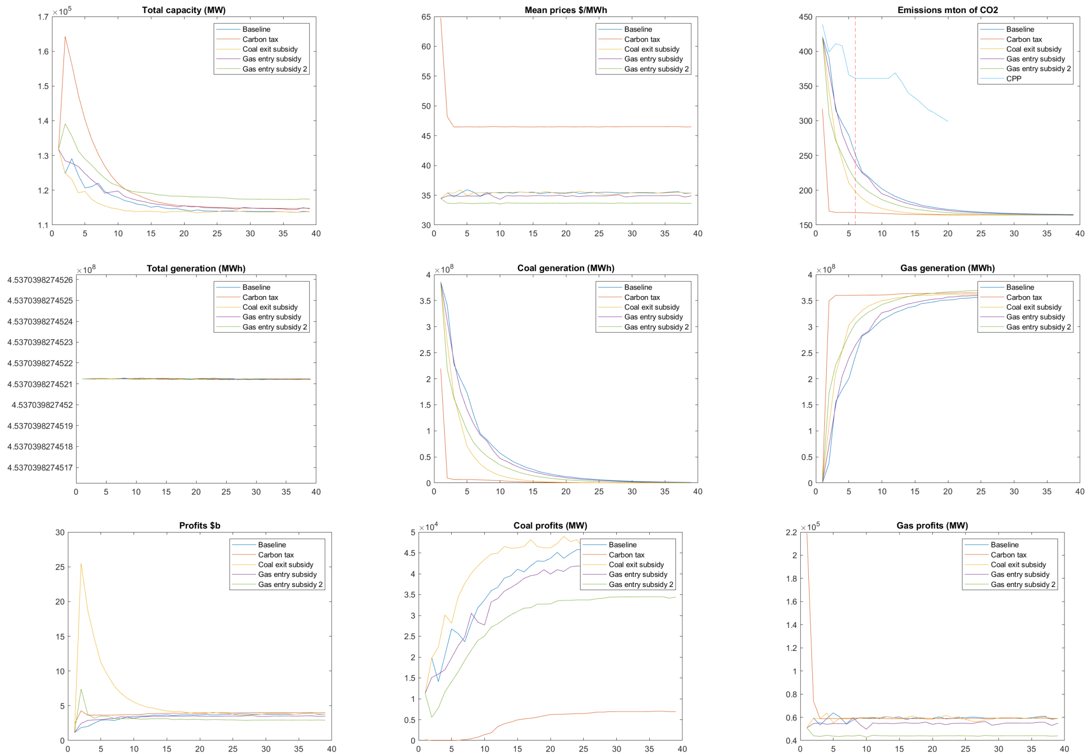

Figure B.3: Detailed Market Outcomes

Notes: emissions CPP scenario with Obama administration plans for PJM states. Total generation is the same across all years as demand is not elastic and the value of loss load is higher than the lifetime cost of entry.

Appendix C

Appendix to Chapter 4

C.1 The Frequency Regulation Market

Table C.1: Number of Units Providing Regulation

# **of units**	1	2	3	4	5	6+
% of hours	11	16	24	23	14	12

Source: Cramton (2012).

C.2 Example - Lost Opportunity Cost

This example is designed to illustrate the role of the lost opportunity cost (LOC) for bidding in the frequency regulation auction. First, I show that the change in calculating LOCs in 2013 did not alter bidding incentives but changed compensation significantly. Second, I show that changes in LOC can effect market power; and therefore the bidding incentives of participants significantly under uniform price auction formats. Finally, I show that changes in LOC should not lead to different bidding behavior under VCG.

The system operator minimizes the cost of procuring frequency regulation. This includes compensation for opportunity cost of power plants. For simplicity, the decision to procure regulation from a certain power plant does not change the energy price significantly. Therefore, the energy price, the energy bid and the LOC are exogenous to the decision of participants and the system operator. Participants only choose a regulation bid that has to be the same for all capacity offered.

Table C.2 describes the available power plants. There are two gas power plants and a hydro power plant available. The energy bids are reflective of their underlying marginal costs. Gas 1 is more efficient than Gas 2. All plants can offer 10 MW of capacity for regulation and they have 0 marginal costs. LOC is calculated by the system operator and is not part of the bid. The capacity requirement for regulation is 15 MW.

Table C.2: Example for Lost Opportunity Cost

Unit	**Energy bid**	**LOC**	**Reg cost**	**Reg cap.**	**Bid uniform**	**Bid VCG**
Gas 1	35	5	0	10	0	0
Gas 2	30	10	0	10	20	0
Hydro	10	30	0	10	0	0

With uniform price, with the LOC not included in the price this is essentially an asymmetric Bertrand game with capacity constraints. In one equilibrium, Gas 2 sets the price at 20: the LOC difference of Gas 2 and Hydro [1]. The market price is 20 and all producing plants get compensated for their LOCs. Gas 1 and 2 receive \$25/MWh and \$30/MWh respectively. The allocation is efficient but plants bid above their marginal cost significantly.

With uniform price, with the LOC included in the price the bids and the outcome is the same. However, the market price is now 30. Both gas plants receive the same \$30/MWh. Under VCG the compensation of plants is determined solely by the bids of other plants. However, selection is based on their own bids. This design ensures that plants bid truthfully and bids are independent of LOCs.There are two ways to

[1] The bids of plants that are not price setting are not unique without introducing uncertainty.

interpret LOC: part of marginal cost or product differentiation. The system operator prefers plants with lower LOCs for regulation. As the example illustrates, lower LOCs give plants market power and lead to bids significantly higher than marginal cost.

C.3 The Modified VCG Auction in New England

FERC required the ISO NE to modify its initial proposal to produce prices for both capacity and price dimensions. These are based on the bundled payment for each resource:

$$bundle_i = J_{-i}(b^c_{-i}, b^m_{-i}) - J(b^c, b^m) + b^c_i * q^*_i + b^m_i * q^{m*}_i$$

The service clearing price is defined as $p_s = \max_i(c_i)$ where i represents all resources that are cleared. Then, the capacity clearing price is set by the maximum of capacity break-even costs (and scaled to represent unit MW terms):

$$p_c = \max_i[(J_{-i}(b^c_{-i}, b^m_{-i}) - J(b^c, b^m) + b^c_i * q^{m*}_i)/q^{m*}_i]$$

Finally, the expected payment for i is:

$$Rev_i = \max_i[(J_{-i}(b^c_{-i}, b^m_{-i}) - J(b^c, b^m) + b^c_i * q^*_i)/q^*_i] * q_j + (\max_i b^m_i) * q^{m*}_i$$

The mechanism remains truthful as the compensation is still independent of the bid of the plant. A power plant's bid still only influences its selection. The mechanism remains efficient as well. However, the expected payments are higher under this modified design. The expected payment equals the bundle payment for a resource that sets both the service and capacity clearing price and it is weakly higher for all other resources.

C.4 Additional Empirical Results

Table C.3: NE Regulation Service Price Statistics

	2015 after	2016	2017	2018	2019
Service price	0.30	0.43	0.34	0.25	0.28
	(0.64)	(0.81)	(0.99)	(0.76)	(0.81)
Observations	6625	8784	8760	8759	8759

Note: An observation is an hour. Means, standard deviations in paranthesis.

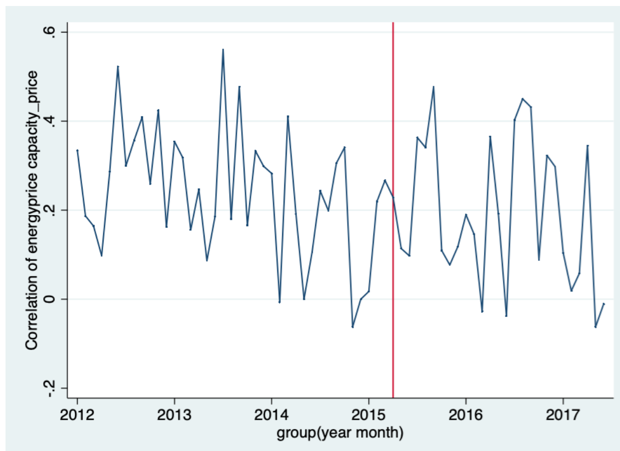

Figure C.1: Monthly Correlation Energy and Regulation Prices, 2013-2017

Table C.4: NE Regulation Individual Bid Statistics - All Bidders

	2013	2014	2015b	2015a	2016	2017	2018	2019
Mean - cap	7.71	8.94	32.03	25.89	30.85	23.33	18.50	16.53
	(0.00)	(0.00)	(15.55)	(8.99)	(12.93)	(10.00)	(6.71)	(6.65)
Sd - cap	1.60	1.74	3.06	2.04	1.63	5.94	1.96	3.03
	(0.00)	(0.00)	(0.00)	(0.69)	(0.00)	(0.47)	(0.00)	(0.00)
Mean - serv				1.73	1.62	1.33	1.12	1.12
				(0.14)	(0.00)	(0.02)	(0.01)	(0.01)
Sd - serv				0.34	0.24	0.37	0.19	0.35
				(0.08)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	352	340	83	82	93	94	90	99

Notes: An observation is a plant-year. 2015b and 2015a is before and after the rule change respectively. All plants included. Statistics aggregated to the plant-year level. Variable "Mean - cap bid" represents the individual plant's mean capacity bid in a given year. Means, medians in parenthesis.

Table C.5: NE Regulation Individual Bid Statistics - Bid Changers

	2013	2014	2015b	2015a	2016	2017	2018	2019
Mean - cap	17.07	27.31	31.00	16.20	8.75	12.26	11.64	13.87
	(12.40)	(26.39)	(19.25)	(7.74)	(8.05)	(7.29)	(6.81)	(7.83)
Sd - cap	11.34	15.61	11.64	2.82	5.39	5.78	5.70	6.78
	(11.13)	(16.52)	(11.59)	(1.09)	(4.38)	(2.95)	(3.09)	(5.36)
Mean - serv				2.46	1.38	0.97	0.45	1.24
				(1.17)	(0.19)	(0.20)	(0.05)	(0.11)
Sd - serv				0.54	0.66	0.79	0.38	0.77
				(0.30)	(0.02)	(0.05)	(0.00)	(0.05)
Observations	36	36	20	49	31	38	34	40

Notes: An observation is a plant-year. 2015b and 2015a is before and after the rule change respectively. Only plants who bid every year and have a positive standard deviation of bids included. Statistics aggregated to the plant-year level. Variable "Mean - cap bid" represents the individual plant's mean capacity bid in a given year. Means, medians in parenthesis.

Table C.6: Service Bid Regressions After 2016

	(1) OLS	(2) FE
Energy price ($/MWh)	0.000464	0.00208
	(0.00286)	(0.00275)
Regulation requirement (MW)	-0.00000580	0.0000391
	(0.0000576)	(0.0000422)
Status	0.378	-0.163**
	(0.373)	(0.0802)
Regulation capacity (MW)	0.00796	0.00249
	(0.00538)	(0.00216)
Response rate (MW/min)	0.0583**	
	(0.0220)	
Constant	-0.654***	0.854***
	(0.219)	(0.192)
Observations	1251218	1251218
Adjusted R^2	0.126	0.572
Clusters	47	47

Notes: an observation is a power plant in a given hour. Dependent variable is the service bid. Only after FERC Order 755 included, 2016-2019. All specifications include power plant fixed effects. Only plants who bid every year and change their bids in a given year included.

Standard errors are clustered at the power plant level and are reported in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

C.5 Individual Bid Graphs

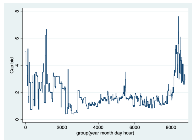
(a) 13547

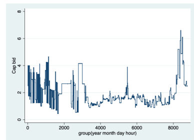
(b) 31662

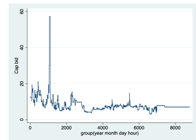
(c) 34578

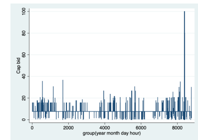
(d) 41406

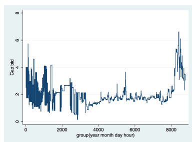
(e) 44171

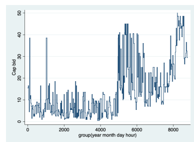
(f) 53298

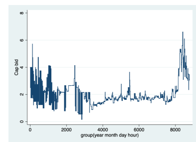
(g) 55388

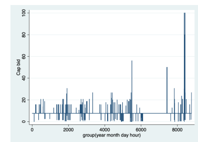
(h) 56535

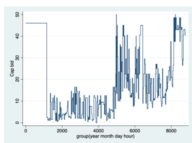
(i) 60188

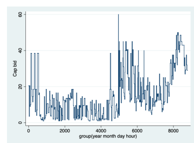
(j) 66602

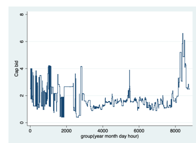
(k) 78679

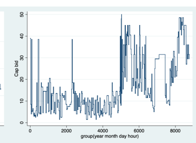
(l) 83099

Figure C.2: Bidder Timelines 2016 - High Frequency Changers

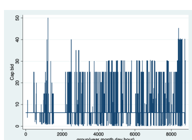
(a) 12297

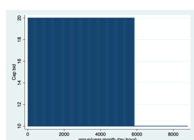
(b) 17063

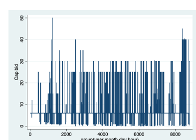
(c) 24780

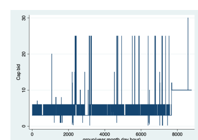
(d) 56483

Figure C.3: Bidder Timelines 2016 - Hour Changers

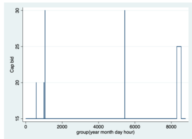

(a) 23789

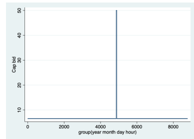

(b) 36992

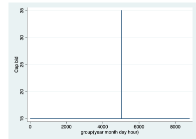

(c) 43414

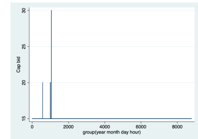

(d) 43551

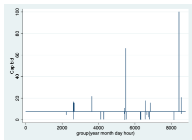

(e) 45714

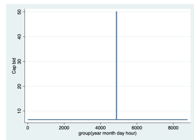

(f) 55408

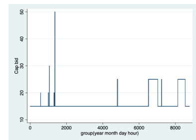

(g) 59958

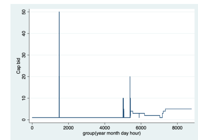

(h) 64231

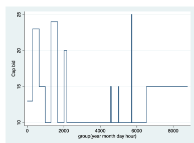

(i) 64814

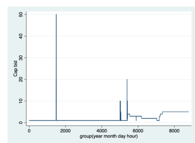

(j) 64231

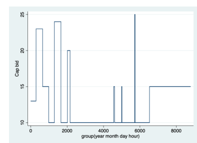

(k) 64814

Figure C.4: Bidder Timelines 2016 - Occasional High Bidders

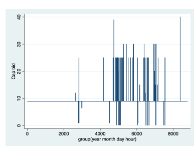

(a) 20846

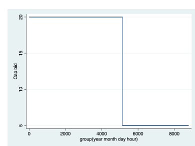

(b) 58855

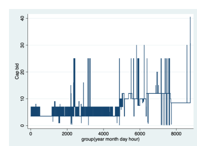

(c) 64475

Figure C.5: Bidder Timelines 2016 - Strategy Changers

Bibliography

Abito, Jose Miguel, Christopher R. Knittel, Konstantinos Metaxoglou, and Andre Trindade. 2019. "Output Reallocation and Investment in Incomplete Emission Regulation."

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica*, 83(6): 2411–2451.

Aldy, Joseph E, Todd D Gerarden, and Richard L Sweeney. 2018. "Investment Versus Output Subsidies: Implications of Alternative Incentives for Wind Energy." National Bureau of Economic Research.

Allcott, Hunt. 2012*a*. "Real-Time Pricing and Electricity Market Design." *NBER Working paper*.

Allcott, Hunt. 2012*b*. "The Smart Grid, Entry, and Imperfect Competition in Electricity Markets." National Bureau of Economic Research.

Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. "Dynamic Inputs and Resource (Mis)Allocation." *Journal of Political Economy*, 122(5): 1013–1063.

Asker, John, and Estelle Cantillon. 2008. "Properties of Scoring Auctions." *The RAND Journal of Economics*, 39(1): 69–85.

Athey, Susan, and Jonathan Levin. 2001. "Information and Competition in U.S. Forest Service Timber Auctions." *Journal of Political Economy*, 109(2): 375–417.

Ausubel, Lawrence M. 2004. "An Efficient Ascending-Bid Auction for Multiple Objects." *American Economic Review*, 94(5): 1452–1475.

Ausubel, Lawrence M., and Paul Milgrom. 2004. "The Lovely but Lonely Vickrey Auction." In *Combinatorial Auctions.* , ed. Peter Cramton, Yoav Shoham and Richard Steinberg. MIT Press.

Ausubel, Lawrence M., Peter Cramton, Marek Pycia, Marzena Rostek, and Marek Weretka. 2014. "Demand Reduction and Inefficiency in Multi-Unit Auctions." *The Review of Economic Studies*, 81(4 (289)): 1366–1400.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *The Quarterly Journal of Economics*, 135(2): 645–709.

Backus, Matthew. 2020. "Why is Productivity Correlated with Competition?" *Econometrica*, 88(6): 2415–2444.

Bajari, Patrick, C Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica*, 75(5): 1331–1370.

Blundell, Richard, and Stephen Bond. 2000. "GMM Estimation with Persistent Panel Data: An Application to Production Functions." *Econometric Reviews*, 19(3): 321–340.

Bolotnyy, Valentin, and Shoshana Vasserman. 2018. "Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement."

Borenstein, Severin, James Bushnell, Frank A Wolak, and Matthew Zaragoza-Watkins. 2019. "Expecting the Unexpected: Emissions Uncertainty and Environmental Market Design." *American Economic Review*, 109(11): 3953–77.

Branger, Frédéric, Jean-Pierre Ponssard, Oliver Sartor, and Misato Sato. 2015. "EU ETS, Free Allocations, and Activity Level Thresholds: The Devil Lies in the Details." *Journal of the Association of Environmental and Resource Economists*, 2(3): 401–437.

Bushnell, James B, Erin T Mansur, and Celeste Saravia. 2008. "Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets." *American Economic Review*, 98(1): 237–66.

Bushnell, James B, Stephen P Holland, Jonathan E Hughes, and Christopher R Knittel. 2017. "Strategic Policy Choice in State-Level Regulation: The EPA's Clean Power Plan." *American Economic Journal: Economic Policy*, 9(2): 57–90.

Bushnell, Jim B, and Jun Ishii. 2007. "An Equilibrium Model of Investment in Restructured Electricity Markets."

Carleton, Tamma, and Michael Greenstone. 2021. "Updating the United States Government's Social Cost of Carbon." University of Chicago, Becker Friedman Institute for Economics Working Paper (2021-04).

Chao, Hung-Po, and Robert Wilson. 2002. "Multi-Dimensional Procurement Auctions for Power Reserves: Robust Incentive-Compatible Scoring and Settlement Rules." *Journal of Regulatory Economics*, 22(2): 161–183.

Collard-Wexler, Allan, and Jan De Loecker. 2015. "Reallocation and Technology: Evidence from the US Steel Industry." *American Economic Review*, 105(1): 131–71.

Collard-Wexler, Allan, and Jan De Loecker. 2016. "Production Function Estimation and Capital Measurement Error." National Bureau of Economic Research.

Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, Ulrich J Wagner, et al. 2020. "Does Pricing Carbon Mitigate Climate Change? Firm-level Evidence from the European Union Emissions Trading Scheme." *Firm-Level Evidence from the European Union Emissions Trading Scheme (November 5, 2020).*

Commission, EU, et al. 2015. "EU ETS Handbook." *Brussels, Belgium. URL: https://ec. europa. eu/clima/sites/clima/files/docs/ets_handbook_en. pdf.*

Copeland, Brian R., and M. Scott Taylor. 2003. *Trade and the Environment: Theory and Evidence.* . STU - Student edition ed., Princeton University Press.

Costantini, James, and Marc Melitz. 2008. "The Dynamics of Firm-Level Adjustment to Trade Liberalization." *The Organization of Firms in a Global Economy*, 4: 107–141.

Cramton, Peter. 2012. "Direct Testimony of Professor Peter Cramton." *New England ISO Filing as a Response to FERC Order No. 755.*

Cramton, Peter, and Steven Stoft. 2005. "A Capacity Market That Makes Sense." *The Electricity Journal*, 18(7): 43–54.

Cullen, Joseph A, and Stanley S Reynolds. 2017. "Market Dynamics and Investment in the Electricity Sector."

Dechezleprêtre, Antoine, and Misato Sato. 2017. "The Impacts of Environmental Regulations on Competitiveness." *Review of Environmental Economics and Policy*, 11(2): 183–206.

De Groote, Olivier, and Frank Verboven. 2019. "Subsidies and Time Discounting in New Technology Adoption: Evidence from Solar Photovoltaic Systems." *American Economic Review*, 109(6): 2137–72.

De Loecker, Jan. 2011. "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity." *Econometrica*, 79(5): 1407–1451.

De Loecker, Jan. 2013. "Detecting Learning by Exporting." *American Economic Journal: Microeconomics*, 5(3): 1–21.

De Loecker, Jan, and Pinelopi Koujianou Goldberg. 2014. "Firm Performance in a Global Market." *Annual Review of Economics*, 6(1): 201–227.

Demirer, Mert. 2020. "Production Function Estimation with Factor-Augmenting Technology: An Application to Markups."

Doraszelski, Ulrich, and Jordi Jaumandreu. 2013. "R&D and Productivity: Estimating Endogenous Productivity." *Review of Economic Studies*, 80(4): 1338–1383.

Doraszelski, Ulrich, and Jordi Jaumandreu. 2018. "Measuring the Bias of Technological Change." *Journal of Political Economy*, 126(3): 1027–1084.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes. 2018. "Just Starting Out: Learning and Equilibrium in a New Market." *American Economic Review*, 108(3): 565–615.

EIA, et al. 2016. "Annual Energy Outlook 2016." Energy Information Administration.

Ellerman, A Denny, and Meghan McGuinness. 2008. "CO2 Abatement in the UK Power Sector: Evidence from the EU ETS Trial Period."

Engelmann, Dirk, and Veronika Grimm. 2009. "Bidding Behaviour in Multi-Unit Auctions–An Experimental Investigation." *The Economic Journal*, 119(537): 855–882.

Eurostat, NACE. 2008. "Rev. 2–Statistical Classification of Economic Activities in the European Community." *Office for Official Publications of the European Communities, Luxemburg.*

Ewerhart, Christian, and Karsten Fieseler. 2003. "Procurement Auctions and Unit-Price Contracts." *The RAND Journal of Economics*, 34(3): 569–581.

Fabra, Natalia. 2018. "A Primer on Capacity Mechanisms." *Energy Economics*, 75: 323–335.

Fabra, Natalia, Nils-Henrik von der Fehr, and David Harbord. 2006. "Designing Electricity Auctions." *The RAND Journal of Economics*, 37(1): 23–46.

Fell, Harrison, and Joshua Linn. 2013. "Renewable Electricity Policies, Heterogeneity, and Cost Effectiveness." *Journal of Environmental Economics and Management*, 66(3): 688–707.

FERC. 2011. "Order 755 - Frequency Regulation Compensation in the Organized Wholesale Power Markets."

Fischer, Carolyn, and Richard G Newell. 2008. "Environmental and Technology Policies for Climate Mitigation." *Journal of Environmental Economics and Management*, 55(2): 142–162.

Frayer, J, S Keane, and J Ng. 2013. "Estimating the Value of Lost Load." *London Economics International LLC.*

Gandhi, Amit, Salvador Navarro, and David A Rivers. 2020. "On the Identification of Gross Output Production Functions." *Journal of Political Economy*, 128(8): 2973–3016.

Gillingham, Kenneth, and James H Stock. 2018. "The Cost of Reducing Greenhouse Gas Emissions." *Journal of Economic Perspectives*, 32(4): 53–72.

Greenstone, Michael. 2002. "The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures." *Journal of Political Economy*, 110(6): 1175–1219.

Greenstone, Michael, John A List, and Chad Syverson. 2012. "The Effects of Environmental Regulation on the Competitiveness of US Manufacturing." National Bureau of Economic Research.

Hall, Robert E. 1988. "The Relation Between Price and Marginal Cost in US Industry." *Journal of Political Economy*, 96(5): 921–947.

He, Guojun, Shaoda Wang, and Bing Zhang. 2019. "Leveraging Political Incentives for Environmental Regulation: Evidence from Chinese Manufacturing Firms."

Hobbs, Benjamin F., Michael H. Rothkopf, Laurel C. Hyde, and Richard P. O'Neill. 2000. "Evaluation of a Truthful Revelation Auction in the Context of Energy Markets with Nonconcave Benefits." *Journal of Regulatory Economics*, 18(1): 5–32.

Hogan, William W, et al. 2005. "On an "Energy Only" Electricity Market Design for Resource Adequacy."

Hopenhayn, Hugo A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica: Journal of the Econometric Society*, 1127–1150.

Hortaçsu, Ali, and David McAdams. 2010. "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market." *Journal of Political Economy*, 118(5): 833 – 865.

Hortaçsu, Ali, and David McAdams. 2018. "Empirical Work on Auctions of Multiple Objects." *Journal of Economic Literature*, 56(1): 157–84.

Hortaçsu, Ali, and Steven L. Puller. 2008. "Understanding Strategic Bidding in Multi-unit Auctions: A Case Study of the Texas Electricity Spot Múarket." *The RAND Journal of Economics*, 39(1): 86–111.

Hortaçsu, Ali, Fernando Luco, Steven L Puller, and Dongni Zhu. 2017. "Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets." National Bureau of Economic Research Working Paper 23526.

ISO NE Internal Market Monitor. 2018. "2017 Annual Markets Report." ISO New England.

ISO New England. 2012. "Filing to FERC, Docket No. ER12-000, Order No. 755 Regulation Market Changes."

ISO New England. 2015. "Regulation Market Manual M-Reg."

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*, 104(2): 537–63.

Ito, Koichiro, and Mar Reguant. 2016. "Sequential Markets, Market Power, and Arbitrage." *American Economic Review*, 106(7): 1921–57.

Jaraitė, Jūratė, Thijs Jong, Andrius Kažukauskas, Aleksandar Zaklan, and Alexander Zeitlberger. 2013. "Matching EU ETS Accounts to Historical Parent Companies A Technical Note." *European University Institute, Florence. Retrieved June*, 18: 2014.

Jones, DL. 2012. "Available and Emerging Technologies for Reducing Greenhouse Gas Emissions from the Iron and Steel Industry." *US EPA, Office of Air Quality Planning and Standards, Sector Policies and Programs Division.*

Kagel, John H. 1995. "Auctions: A Survey of Experimental Research." In *The Handbook of Experimental Economics*. Vol. 1, , ed. John H. Kagel and Dan Levin, Chapter 7. Princeton University Press.

Kagel, John H., and Dan Levin. 2016. "Auctions: A Survey of Experimental Research." In *The Handbook of Experimental Economics*. Vol. 2, , ed. John H. Kagel and Alvin E. Roth, Chapter 9. Princeton University Press.

Kalemli-Ozcan, Sebnem, Bent Sorensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yesiltas. 2015. "How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts and Aggregate Implications." National Bureau of Economic Research.

Klemperer, Paul D., and Margaret A. Meyer. 1989. "Supply Function Equilibria in Oligopoly under Uncertainty." *Econometrica*, 57(6): 1243–1277.

Knittel, Christopher R., and Konstantinos Metaxoglou. 2008. "Diagnosing Unilateral Market Power in Electricity Reserve Market." *Journal of Energy Markets*, 1(1): 65–95.

Levinsohn, James, and Amil Petrin. 2003. "Estimating Production Functions using Inputs to Control for Unobservables." *The Review of Economic Studies*, 70(2): 317–341.

Linn, Joshua, and Kristen McCormack. 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.

List, John A., and David Lucking-Reiley. 2000. "Demand Reduction in Multiunit Auctions: Evidence from a Sportscard Field Experiment." *The American Economic Review*, 90(4): 961–972.

Lucking-Reiley, David. 2000. "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce." *Journal of Economic Perspectives*, 14(3): 183–192.

Mahasenan, N, RT Dahowski, and CL Davidson. 2005. "The Role of Carbon Dioxide Capture and Storage in Reducing Emissions from Cement Plants in North America." In *Greenhouse Gas Control Technologies 7*. 901–909. Elsevier.

Marin, Giovanni, Marianna Marino, and Claudia Pellegrin. 2018. "The Impact of the European Emission Trading Scheme on Multiple Measures of Economic Performance." *Environmental and Resource Economics*, 71(2): 551–582.

Marschak, Jacob, and William H Andrews. 1944. "Random Simultaneous Equations and the Theory of Production." *Econometrica, Journal of the Econometric Society*, 143–205.

Martin, Ralf, Mirabelle Muûls, and Ulrich J Wagner. 2014. "Trading Behavior in the EU Emissions Trading Scheme."

Martin, Ralf, Mirabelle Muûls, and Ulrich J Wagner. 2016. "The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What Is the Evidence After Ten Years?" *Review of Environmental Economics and Policy*, 10(1): 129–148.

Metaxoglou, Konstantinos, and Aaron Smith. 2007. "Efficiency of the California electricity reserves market." *Journal of Applied Econometrics*, 22(6): 1127–1144.

Milgrom, Paul. 2000. "Putting Auction Theory to Work: The Simultaneous Ascending Auction." *Journal of Political Economy*, 108(2): 245–272.

Moran, Alan, and Ben Skinner. 2008. "Resource Adequacy and Efficient Infrastructure Investment." In *Competitive Electricity Markets*. 387–416. Elsevier.

Murray, Brian C, Maureen L Cropper, Francisco C de la Chesnaye, and John M Reilly. 2014. "How Effective are US Renewable Energy Subsidies in Cutting Greenhouse Gases?" *American Economic Review*, 104(5): 569–74.

Myatt, James. 2017. "Market Power and Long-Run Technology Choice in the U.S. Electricity Industry."

Nauclér, Tomas, and Per-Anders Enkvist. 2009. "Pathways to a Low-carbon Economy: Version 2 of the Global Greenhouse Gas Abatement Cost Curve." *McKinsey & Company*, 192(3).

Ocker, Fabian, Karl-Martin Ehrhart, and Marion Ott. 2018. "Bidding Strategies in Austrian and German Balancing Power Auctions." *Wiley Interdisciplinary Reviews: Energy and Environment*, 7(6): e303.

Ocker, Fabian, Karl-Martin Ehrhart, and Matej Belica. 2018. "Harmonization of the European Balancing Power Auction: A Game-Theoretical and Empirical Investigation." *Energy Economics*, 73: 194–211.

Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6): 1263–1297.

Petrick, Sebastian, and Ulrich J Wagner. 2014. "The Impact of Carbon Trading on Industry: Evidence from German Manufacturing Firms."

Porter, David, and Roumen Vragov. 2006. "An Experimental Examination of Demand Reduction in Multi-unit Versions of the Uniform-price, Vickrey, and English Auctions." *Managerial and Decision Economics*, 27(6): 445–458.

Porter, Michael E, and Claas Van der Linde. 1995. "Toward a New Conception of the Environment-Competitiveness Relationship." *Journal of economic perspectives*, 9(4): 97–118.

Rafey, Will. 2020. "Droughts, Deluges, and (River) Diversions: Valuing Market-Based Water Reallocation."

Raval, Devesh R. 2019. "The Micro Elasticity of Substitution and Non-neutral Technology." *The RAND Journal of Economics*, 50(1): 147–167.

Reguant, Mar. 2014. "Complementary Bidding Mechanisms and Startup Costs in Electricity Markets." *The Review of Economic Studies*, 81(4): 1708–1742.

Rothkopf, Michael H. 2007. "Thirteen Reasons Why the Vickrey-Clarke-Groves Process Is Not Practical." *Operations Research*, 55(2): 191–197.

Schwenen, Sebastian. 2015. "Strategic Bidding in Multi-Unit Auctions with Capacity Constrained Bidders: The New York Capacity Market." *The RAND Journal of Economics*, 46(4): 730–750.

Shapiro, Joseph S, and Reed Walker. 2018. "Why is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade." *American Economic Review*, 108(12): 3814–54.

Syverson, Chad. 2004. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy*, 112(6): 1181–1222.

Toyama, Yuta. 2018. "Incentives to Invest, Storable Permits, and Transaction Costs in Market-Based Environmental Regulation."

Varian, Hal R., and Christopher Harris. 2014. "The VCG Auction in Theory and Practice." *American Economic Review*, 104(5): 442–45.

Vickrey, William. 1961. "Counterspeculation, Auctions and Competitive Sealed Tenders." *The Journal of Finance*, 16(1): 8–37.

Willems, Bert, Ina Rumiantseva, and Hannes Weigt. 2009. "Cournot Versus Supply Functions: What Does the Data Tell Us?" *Energy Economics*, 31(1): 38–47.

Wilson, Robert. 1979. "Auctions of Shares." *The Quarterly Journal of Economics*, 93(4): 675–689.

Wolak, Frank A. 2000. "An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market." *International Economic Journal*, 14(2): 1–39.

Xu, Bolun, Y. Dvorkin, D. S. Kirschen, C. A. Silva-Monroy, and J. Watson. 2016. "A Comparison of Policies on the Participation of Storage in U.S. Frequency Regulation Markets." *2016 IEEE Power and Energy Society General Meeting (PESGM)*, 1–5.