NORTHWESTERN UNIVERSITY

Essays on Empirical Microeconomics

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Yuta Toyama

EVANSTON, ILLINOIS

August 2018

ABSTRACT

Essays on Empirical Microeconomics

Yuta Toyama

My dissertation consists of two chapters that empirically study policy-related questions in applied microeconomics by using structural econometric modeling developed in industrial organization. In the first chapter, I study the welfare effects of a cap-and-trade program. I develop an equilibrium framework that incorporates forward-looking behavior and transaction costs. In the presence of transaction costs in permit trading, investment patterns may depart from the first best outcome. Storable emissions permits allow firms to smooth costs over time, and abatement investment introduces dynamic incentives into compliance decisions. I apply the framework to study the first nine years (1995-2003) of the US Acid Rain Program. Using data on permit transactions and electricity production, I estimate the model and show that variable transaction costs are substantial. I use the estimated model to quantify the effect of a cap-and-trade program in comparison to a uniform standard, given a fixed level of aggregate emissions. I find that the total costs of reducing emissions under cap-and-trade are 16.6% lower. Although health and environmental damages from SO_2 emissions increase due to the change in the geographic distribution of emissions, the net benefit of the cap-and-trade is positive. I also examine the potential gains from trade in the

absence of transaction costs. I find more dispersed patterns of investment and less banking of permits, both of which result in cost savings.

In the second chapter, a joint work with Kei Kawai and Yasutora Watanabe, we study how voter turnout affects the aggregation of preferences in elections. Under voluntary voting, election outcomes disproportionately aggregate the preferences of voters with low voting cost and high preference intensity. We show identification of the correlation structure among preferences, costs, and perceptions of voting efficacy, and explore how the correlation affects preference aggregation. Using 2004 U.S. presidential election data, we find that young, low-income, less-educated, and minority voters are underrepresented. All of these groups tend to prefer Democrats, except for the less-educated. Democrats would have won the majority of the electoral votes if all eligible voters had turned out.

Acknowledgments

I am deeply indebted to my advisers, Rob Porter and Mar Reguant, for their encouragement and advises through my doctoral studies. Rob always patiently listened to my ideas since the beginning of the dissertation project and directed me to the right path. Mar gave me strict guidance on my analysis in the project, yet has been always positive and encouraging, which was a huge help especially when I felt less confident with my project. They both are a great role model as an IO economist and a mentor to my future students.

I am extremely grateful to Kei Kawai and Yasutora Watanabe for collaboration in the research project on which the second chapter of my Ph.D. thesis is based. It has been a great joy to work with them. I have learned a lot from them, especially about how to develop analysis and results into a whole research paper.

In the process of working on my Ph.D. dissertation (especially the first chapter), I have benefitted from comments and discussions made by many people, including Vivek Bhattacharya, Igal Hendel, Gaston Illanes, Koichiro Ito, Vishal Kamat, Lola Segura, Max Tabord-Meehan, Yuta Takahashi, and Nicky Vreugdenhil.

Finally, I would like to thank my family for their understanding and continuing supports since when I decided to pursue doctoral studies in Economics.

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Investment Incentives, Storable Permits, and Transaction Costs in the Intertemporal Cap-and-Trade Program

1.1. Introduction

How to achieve environmental sustainability without compromising economic efficiency is a central question in policy debates and the economics literature. A traditional approach to pollution regulation is imposing a uniform standard on emissions intensity, although this approach may not be a cost-effective solution. Pollution abatement often requires costly investments in clean technology, and firms are heterogeneous in their costs of reducing pollution. Instead, economists have advocated market-based solutions to give firms an incentive to internalize negative externalities. One example is a cap-and-trade program where firms trade emissions permits to achieve a target level of aggregate emissions. When assessing the welfare effects of a cap-and-trade system, two critical elements exist: transaction costs and dynamic regulatory environment. Transaction costs associated with permit trading may prevent the first best reallocation of emissions permits, violating the Coase theorem. Dynamics influence how firms make investment decisions in the costly compliance technology. In this paper, I develop an empirical framework for a cap-and-trade program that incorporates firms' forward-looking behavior and transaction costs. I apply this framework to evaluate the welfare effects of the US Acid Rain Program, a federal cap-and-trade program designed to reduce sulfur dioxide emissions.

In a cap-and-trade program, firms face two key dynamic decisions: whether to invest in clean technology, and whether to store (bank) emissions permits. Investments in clean, but costly, technology are an important margin for reducing emissions. A cap-and-trade program typically spans a long time horizon, and the allocation of emissions permits changes over time. Firms, therefore, must take into account the change in the regulatory environment in their investment decisions. Moreover, firms can store (bank) emissions permits across periods. The storability of emissions permits under a permit banking system allows firms to smooth their costs over time. Although these dynamic factors are key elements of cap-and-trade programs, modeling these factors and estimating such a dynamic model is a challenging task. To deal with these complications, previous studies focused on a static decision problem on the steady state (See, e.g., Carlson et al. 2000, Fowlie 2010b, and Chan 2015). An innovation of my paper is to model explicitly dynamic aspects of a cap-and-trade program and bring it to the data, which allows me to evaluate comprehensively the evolution of a cap-and-trade program.

Another important focus of my paper is the role of transaction costs in the permit market. Transaction costs discourage firms from relying on the permit trading as a compliance strategy. Previous studies documented that many firms tend not to trade emissions permits, and instead comply with the regulation using their allocated permits (see, e.g., Jaraitė-Kažukauskė and Kažukauskas, 2015, for the EU Emissions Trading Scheme). In the absence of transaction costs, the Coase (1960) theorem implies that investment in abatement (i.e., reduction of emissions) should be efficient under cap-and-trade programs. In practice, transaction costs distort investment incentives, leading to the efficiency loss. My paper quantifies transaction costs in the permit market and their impacts on investment patterns and abatement costs.

I construct a dynamic equilibrium model of investment and a cap-and-trade program that captures these complex interactions in an equilibrium framework. The novelty of my model is that it includes dynamic investment decisions, permit banking, and transaction costs, each of which was studied separately in the previous literature, in a unified framework.[1] In my model, firms are price takers for emissions permits. They face various tradeoffs in their compliance decisions. Firms can comply with the regulation either by reducing emissions, which can be achieved by decreasing production or investing in clean technology, or by buying emissions permits. However, firms incur two type of transaction costs in the permit market: (1) a sunk cost associated with participation in the permit market and (2) variable costs that depend on the trading volume. These costs discourage firms from permit trading and affect investment patterns. The sequence of permit prices is determined by market clearing conditions, resulting in a dynamic competitive equilibrium.

I apply this framework to study the first nine years (1995-2003) of the US Acid Rain Program, a cap-and-trade program for sulfur dioxide (SO_2) emissions that targets the US electricity industry.[2] The goal of the Acid Rain Program is to reduce aggregate SO_2 emissions from generation facilities to half of their 1980 levels. The regulator distributed emissions permits to the existing generation facilities, and these facilities were required to hold sufficient permits to offset their emissions each year.[3] Regulated sources could choose how to comply

[1]The previous literature that studies various models of a cap-and-trade program includes a static trading model with transaction costs (Stavins, 1995), a theoretical model of permit banking (Rubin, 1996; Schennach, 2000), and a model of long-run abatement investment (Fowlie, 2010b).

[2]I choose 2003 as a terminal period of my analysis because of the announcement of the Clean Air Interstate Rule in December 2003, which had a major impact on the regulatory environment regarding SO_2 emissions. See section 1.2.2 for details.

[3]Emissions permits are called emissions "allowances" in the Acid Rain Program because the term "permit" has another meaning in US environmental law. Because "permit" is the standard terminology in the economics literature, I use the term "permit" in this paper.

with the regulation. For example, they could switch to a cleaner fuel, invest in abatement equipment, or obtain additional permits from the market.

An appealing feature of the Acid Rain Program is the availability of data. Various information is publicly available, including production data for power plants and trading data for emissions permits by electric utilities. Combining these two data sets allows me to identify the key parameters of my model, including transaction costs of permit trading and investment costs. My focus on the Acid Rain Program is also motivated by previous findings that the gains from permit trading might not have been fully realized (Carlson et al. 2000 and Keohane 2006 on Phase I, and Chan 2015 on Phase II). I revisit their findings by taking into account transaction costs of permit trading as a source of inefficiency.

I combined data on permit transactions and production information between 1990 and 2003. Identification of my model relies on optimality conditions regarding firms' decisions and detailed information on production and permit transactions. I use production data before and after the introduction of the cap-and-trade program to estimate the firm-level profit function from electricity production. The profit function implies the marginal profit from emissions across firms. Absent variable transaction costs, marginal profits should be equalized across firms, and equal to the permit price. Variable transaction costs are identified from how marginal profits vary with the trading volume. Firm-level participation in permit trading is employed to identify sunk costs of participation. Finally, I use the first-order-condition for investment to identify the marginal costs of investment.

I estimate the model parameters by simulated nonlinear least squares. The estimates imply that sunk participation costs are quite small, but variable transaction costs from permit trading are substantial. The median marginal transaction cost is estimated to be 98 USD, while permit prices range between 100 and 200 USD in my sample period.

I conduct a series of counterfactual simulations using the model estimates. First, I examine the impact of a cap-and-trade program in comparison to a uniform standard regulation under which all firms are required to have the same emissions rate. I find that the investment pattern under cap-and-trade differs significantly from the pattern under a uniform standard regulation, for a given aggregate level of emissions. Rather than following the uniform emissions rate imposed by the regulator, firms can optimally choose their levels of investment based on the costs and returns under a cap-and-trade program. The aggregate costs of abatement are 16.6% lower than under a uniform standard regulation.

I also discuss the implications for health and environmental damages. A potential concern of a cap-and-trade program is that it could lead to higher health and environmental damages in comparison to uniform standard regulations. Even though the aggregate level of emissions is fixed, the geographic distribution of emissions might differ from that under a uniform standard. Damages from SO_2 emissions depend on the location of emissions sources. health and environmental damages, therefore, can differ across regulatory regimes (see, e.g., Muller and Mendelsohn 2009, Fowlie et al. 2012, Fowlie and Muller 2013, and Chan et al. 2015). I use the data from Muller and Mendelsohn (2009) to calculate the damages under both a cap-and-trade program and a uniform standard. I find that the health and environmental damage does increase under a cap-and-trade program. But the net benefit of the cap-and-trade, calculated by the sum of abatement costs and health and environmental damages, is positive.

In a second counterfactual simulation, I examine potential gains from trade. I find that shutting down transaction costs, as might happen if there was a centralized trading platform, would lead to a more dispersed distribution of emissions and investment levels, reflecting more active trading of emissions permits. There is less permit banking in the absence of transaction costs. Transaction costs discourage firms from selling emissions permits, and firms prefer to accumulate these permits, lowering the allocative efficiency of emissions permits. In the absence of transaction costs, the total abatement costs decrease by around 37%, enough to offset increases in health and environmental damages. Thus, "unrealized" gains from trade are significant in my sample period.

My empirical framework can be applied to other market-based environmental policies, including water trading systems, the current Corporate Average Fuel Economy (CAFE) regulation, and the Renewable Portfolio Standard. A key feature of these regulations is the interaction between investment in clean technology and trading of environmental credits. For example, the recent CAFE standard regulation allows firms to trade CAFE credits with other firms for their compliance. This trading scheme is an alternative to improving fuel efficiencies for a manufacturers' fleet by making a costly investment. Under the Renewable Portfolio Standard, electricity utilities can either invest in renewable technologies or obtain the Renewable Energy Certificates from the market to comply with the regulation.

1.1.1. Related Literature

My paper is related to three strands of literature: (i) the empirical literature on dynamic investment behavior, (ii) the empirical literature on cap-and-trade regulations, and (iii) the evaluation of the Acid Rain Program.

Understanding firms' investment behavior and its welfare consequences is a central theme in the industrial organization literature. Previous work includes Ericson and Pakes (1995), Bajari et al. (2007), Ryan (2012), Collard-Wexler (2013), and Kalouptsidi (2014) in an oligopolistic setting, and Rust (1987), Aguirregabiria and Mira (2002), and Kellogg (2014) in a competitive setting.

A novel feature of my paper is that investment in technology is substitutable with trading of emissions permits. To comply with the cap-and-trade regulation, a firm can either make an investment and reduce emissions, or purchase emissions permits from the permit market, where a firm faces transaction costs. Investment decisions also interact with storability of inputs, namely banking of emissions permits in my model. For example, if tighter future regulatory intensity is announced at the inception of the regulation, permit banking induces investment in early periods. A similar market structure can be found in other settings, including the CAFE credit trading program and the green certificate trading in the Renewable Portfolio Standard.

My paper also contributes to the empirical literature on cap-and-trade programs. Much of the literature test qualitative predictions of models of permit trading. A few recent papers take a structural approach to measure the welfare implications of permit trading.[4] In the context of NOx regulation, Fowlie (2010b) constructs a model of abatement choice to study the effect of rate-of-return regulation on permit trading. Fowlie et al. (2014) construct and estimate a model of dynamic investment and entry/exit game to discuss the implications of hypothetical market-based environmental policies in the US cement industry.[5] Abito (2014) quantifies the impact of rate-of-return regulation on the efficiency of SO_2 emissions regulation by estimating a multi-product cost function, though he does not explicitly consider permit trading.

[4]The literature has examined the independence of outcomes from the initial allocation (Reguant and Ellerman, 2008 and Fowlie and Perloff, 2013) and the internalization of emissions costs (Kolstad and Wolak, 2008, Fowlie, 2010a, and Fabra and Reguant, 2014).

[5]Dardati (2014) also studies how an allocation scheme for closing plants affects entry/exit decisions, using the calibrated model of industry dynamics in the context of the Acid Rain Program.

A distinctive feature of my paper is to model trading behavior in the permit market and banking of emissions permits.[6] The previous papers all assume frictionless permit markets in which cap-and-trade is equivalent to imposing a Pigouvian tax. My model captures emissions permit trading when permits may be banked and transaction costs exist. It can quantify transaction costs and their impact. My framework can also be used to study how the regulatory design of permit trading, such as the availability of permit banking and alternative allocation rules for emissions permits, affects firms' abatement decisions.

Finally, my paper provides new insights for the evaluation of the Acid Rain Program. One approach in the literature is to calculate cost saving due to permit trading by estimating a cost function and a discrete choice model for abatement choice (see, e.g., Ellerman et al., 2000, Carlson et al., 2000, Keohane, 2006, and Chan, 2015). Researchers found that adopting a permit trading program led to significant cost savings compared to traditional command-and-control approaches, though the actual cost did not reach the least-cost solution. Another approach is to focus on aggregate variables to discuss the efficiency of the permit market (Joskow et al. 1998, Ellerman and Montero 2007, and Helfand et al. 2006).[7]

My paper complements this literature by providing an empirical model that incorporates firms' decisions on abatement, permit trading, and permit banking in an equilibrium framework. My model allows me to evaluate the role of permit banking and transaction costs. I decompose the effects of the Acid Rain Program into the effects of permit trading and permit banking. Previous studies often note the importance of permit banking as a

[6]Cantillon and Slechten (2015) might be the closest to my paper. They study participation decisions and price formation for CO_2 emissions permits using trading data in the EU-ETS scheme.

[7]Joskow et al. (1998) finds that prices in the spot market and the EPA auction are close and concludes that "a relatively efficient private market" had developed by mid-1994. Ellerman and Montero (2007) argues for the efficient market of permits by comparing the actual and theoretically predicted volume of aggregate banking. Helfand et al. (2006) uses the monthly permit prices during 1994 to 2003 to test whether the price path follows the Hotelling r-percent rule for intertemporal arbitrage. They reject the Hotelling rule, which is a suggestive evidence of inefficiency of the market.

source of cost efficiency. My paper is the first to quantify the gains. Moreover, I quantify the potential gains from trade which could be achieved in the absence of transaction costs. These simulation analyses require an equilibrium model of the cap-and-trade program. I also use transaction data for emissions permits to estimate the model, in contrast to most of the literature.

1.2. Empirical Setting and Descriptive Analysis

1.2.1. The Acid Rain Program

Fossil-fuel electricity plants, especially coal plants, produce sulfur dioxide (SO_2) emissions as a byproduct of electricity generation. SO_2 is known to have detrimental effects on human health and the environment. Although the federal government introduced command-and-control-type regulations with the Clean Air Act Amendments of 1970, such regulations have not been effective in reducing SO_2 emissions.[8] The failure of the previous regulations led to the introduction of the Acid Rain Program (ARP), a cap-and-trade program, in 1995.

The target of the regulation is electricity generating units (EGUs) that use fossil fuels and have an output capacity greater than 25 megawatts. The regulation was implemented in two phases. In Phase I (1995-1999), a subset of eligible EGUs are under the regulation. These units include 263 EGUs named the "Table 1" group, which were especially dirty and old before the regulation, and an additional 182 EGUs from the Non-Table 1 group as substitution or compensating units. In Phase II (begun in 2000), all eligible EGUs are mandated to comply with the regulation.

The ARP aims to reduce SO_2 emissions from generation facilities to half of their 1980 levels, which determines the total number of emissions permits for each year. Most of the

[8]Ellerman et al. (2000) provide a brief history of the regulation on SO_2 emissions.

emissions permits are allocated for free to incumbent units. The EPA adopts the rule that determines the unit-level allocation of emissions permits based on the characteristics of a unit.[9] The allocation is primarily determined by the product of average heat inputs during 1985-1987 and the target emissions rate for each Phase (2.5 pounds per 1 million British thermal unit (lb/MMbtu) in Phase I and 1.2 lb/MMBtu in Phase II). Some units also obtain additional allocation of permits based on technical and political considerations (Joskow and Schmalensee, 1998).

SO_2 permits are tradable goods. Firms can sell or buy permits with other firms, including financial companies or brokers that do not own any generating units and thus are not required to comply with the regulation. Although the EPA also holds an annual auction to distribute around 2.7% of the yearly allocation, a centralized trading exchange does not exist. Bilateral trading, which is often mediated by brokers, is the primary way to trade emissions permits with other participants.

Operation of each regulated unit, especially emissions levels of SO_2, is recorded through the Continuous Emissions Monitoring System.[10] At the end of the calendar year, the annual level of SO_2 emissions is finalized, and each regulated unit is required to surrender emissions permits within a grace period of 60 days. The remaining permits are carried over to the next year, which is called banking of emissions permits.[11] As I discuss in section 1.2.3, regulated

[9] See U.S. Environmental Protection Agency (1993a,b) for the details.

[10] There should be no concern about manipulating the measurement of emissions because the operators are required to perform periodic performance evaluations of the monitoring system. These evaluations include daily calibration error tests, daily interference tests for flow monitors, and semi-annual (or annual) relative accuracy test audit and bias tests. See `https://web.archive.org/web/20090211082920/http://epa.gov/airmarkets/emissions/continuous-factsheet.html` for the details.

[11] If an affected unit does not hold sufficient permits to offset the emissions at the end of the compliance deadline, unit operators are required to pay the penalty of \$2000 per SO_2 ton. However, compliance was nearly 100% during the period of my analysis.

firms had a significant amount of banked permits in Phase I when the annual allocation was more generous than in Phase II.

Although emissions permits were allocated to the existing units for free, most of them still needed to decrease their emissions from their business-as-usual level to comply with the regulation. The regulated units were able to reduce emissions by either lowering utilization (output) or emissions per output (emissions rate). The latter option of reducing emissions rates was the primary channel of abatement, which I will explain in section 1.2.3 in detail.

1.2.2. Data

In this paper, I focus on the period from 1995 to 2003. Although the ARP continued after 2004, the proposal of the Clean Air Interstate Rule, announced in December 2003, had a large impact on the regulated firms' expectation over the future regulatory environment. The proposed regulation aimed to strengthen the stringency of the SO_2 regulation from 2010 in the framework of the ARP. After the announcement, the permit price started to rise dramatically, primarily because the value of emissions permits issued before 2010 would be higher than those issued after 2010 in the proposed regulation. Firms also started to invest in scrubbers in anticipation of more strict intensity of the proposed regulation.[12] Thus, I do not include the data after 2004 and rather focus on the periods when the regulatory environment regarding SO_2 emissions was stable.

The data I use are a combination of transaction data for emissions permits and various data on electricity production. The data on permit transactions are from the Allowance Tracking System (ATS) operated by the EPA. The latter data are a compilation of various

[12]See Schmalensee and Stavins (2013) for a detailed review on how the regulatory environment regarding SO_2 emissions has been changing since 2004.

databases from the EPA and the US Energy Information Administration (EIA). I explain these two type of data in turn.

First, the EPA uses the ATS to manage permit allocation and track private transactions and surrender of permits for compliance, and makes the data public. Each transaction record in the tracking system contains the account name of a transferor, a transferee, vintage of permits, quantity of transferred permits, and the confirmation date of the transaction.[13] I constructed the transaction data at the firm and year level from the database. Specifically, I aggregated the account-level information into the firm-level information by using ownership information constructed from various sources including EGrid database and EIA-860. The final data set includes (1) permit holding at the beginning of the year, (2) annual allocation, (3) volume of permit transaction (net purchase of emissions permits), and (4) banking volume. Note that I only used information on transactions of permits whose vintage is current or old.

The ATS does not collect any information on transaction prices. I instead collected the market-price index of SO_2 permits provided by Cantor Fitzgerald, one of the biggest brokers in SO_2 permit markets. The frequency of the price data is monthly. I explain the details of the price data in section 1.2.3.5.

The second piece of my data set is production information of electricity companies. I combined multiple databases to construct the data set. These databases include EPA data as well as EIA survey data. First, the EPA makes public unit-level operation data of generating units collected by the Continuous Emissions Monitoring System (hereafter CEMS). The CEMS data include gross generation (in MWh), heat inputs (in MMBtu), and

[13]The confirmation date must lag behind the actual transaction date to some extent, although the prompt recording of private trading was considered the rule rather than the exception according to the EPA staff and industry experts. See Joskow et al. (1998) for details.

SO_2 emissions. In addition, the EIA conducts various surveys on operation of power plants. Specifically, the Form EIA-767 "Steam-Electric Plant Operation and Design Report" gives me information on fuel usage (sulfur content, ash content, heat inputs), net generation, and generation capacity at the unit and month level. Also, the Form FERC No. 423 (EIA-423) "Monthly Report of Cost and Quality of Fuels for Electric Plants" provides plant- and month-level information on fuel procurement, including fuel type, sulfur contents, heat contents, and purchase costs.

1.2.3. Descriptive Analysis

I now provide a descriptive analysis on the data set I constructed. I focus on various aspects of the ARP including banking of emissions permits, abatement decisions of regulated sources, and market of emissions permits. These descriptive findings motivate the modeling approach I introduce in section 1.3.

1.2.3.1. Banking of Emissions Permits. Figure 1.1 shows the aggregate SO_2 emissions level and emissions caps under the ARP from 1990 to 2003. The bars show emissions levels each year, and the dashed lines show the emissions cap. As I mentioned in section 1.2, the timing of the regulation was different across electricity generating units. I denote those units that are regulated beginning 1995 as group I units and those regulated since 2000 as Group II units. The blue bar in the figure corresponds to emissions from Group I units, and the orange bar corresponds to those from group II units. The blue dashed line shows the allocation for Group I units, and the black dashed line from 2000 shows the total cap of emissions including both Group I and II units.

The figure shows that Group I units reduced their emissions almost by half compared to their 1980 level once Phase I started in 1995. While both Group I and II units reduced

Figure 1.1. Aggregate Volume of SO_2 Emissions and Caps (1990 - 2003)

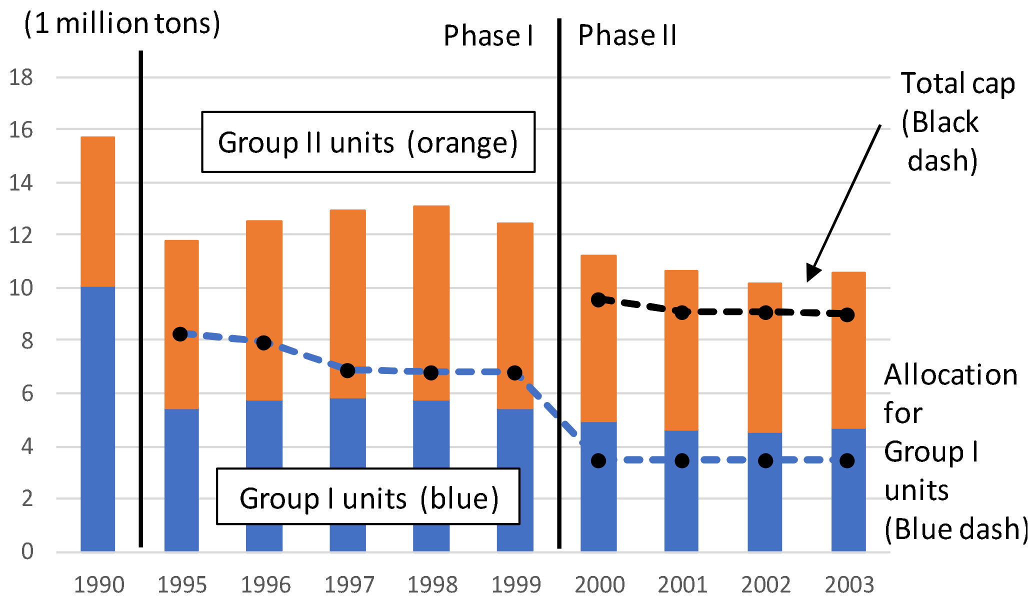

Notes: The blue (orange) bar corresponds to emissions from Group I (Group II) sources. The blue dashed line shows permit allocation for Group I units, and the black dashed line from 2000 shows the total cap including allocation for both Group I and II units.

emissions further in 2000, the first year of Phase II, Group I units did not reduce emissions as much as in 1995. Emissions before 1999 were significantly lower than the emissions cap, though the aggregate emissions exceed the allocation of emissions permits after 2000. These observations imply that Group I units saved their permits in Phase I and then started to use them after 2000 for the purpose of compliance.

1.2.3.2. Abatement Strategy for Coal units. Emissions from electricity generation can be reduced either by (1) reducing the emissions rate (i.e., emissions per output) or (2) reducing output (lower utilization of generating units). In this subsection, I explain the former abatement strategy for generating units whose primary fuel type is coal. Although the target of the ARP includes all types of fossil fuel units (coal, gas, and oil), SO_2 emissions from gas and oil units are relatively small and no room remains for lowering the emissions rate of these units.

Two common options are available to reduce the emissions rate of coal units. The first option is called fuel switching. An operator of coal units can switch the type of coal from dirty (e.g., high-sulfur bituminous coal) to cleaner (e.g., subbituminous coal or low-sulfur bituminous coal). The fuel costs of cleaner coals are higher than the fuel costs of dirty coals. Also, switching fuel types requires retrofitting the boiler to make it compatible with the new type of coal, which incurs fixed costs. Another abatement option is installing flue-gas desulfurization equipment (a scrubber). This equipment is installed at the stack of generation units and eliminates more than 80% of SO_2 emissions. This option, however, incurs large investment costs as well as a long lead time (2 to 3 years on average).

Figure 1.2 shows the distribution of unit-level SO_2 emissions rates (measured in pounds per MMBtu) for each group in selected years. The left panel shows the distribution for group I sources. The emissions rates of these sources decreased between 1990 and 1995, the beginning of Phase I. The emissions rates stayed almost constant within Phase I, and it decreased further in 1999, which anticipates the beginning of Phase II. For generating units in group II, their emissions rates did not change until 1999, and then decreased in 2000, the first year of the cap-and-trade program for these units. These observations imply that firms adjusted their emissions rates at the beginning of each phase, but emissions rates remain almost constant within the phase. This observation motivates my model of abatement investment in the structural model I introduce in section 1.3.

Figure 1.2. Distribution of Unit-level SO_2 Emissions Rate

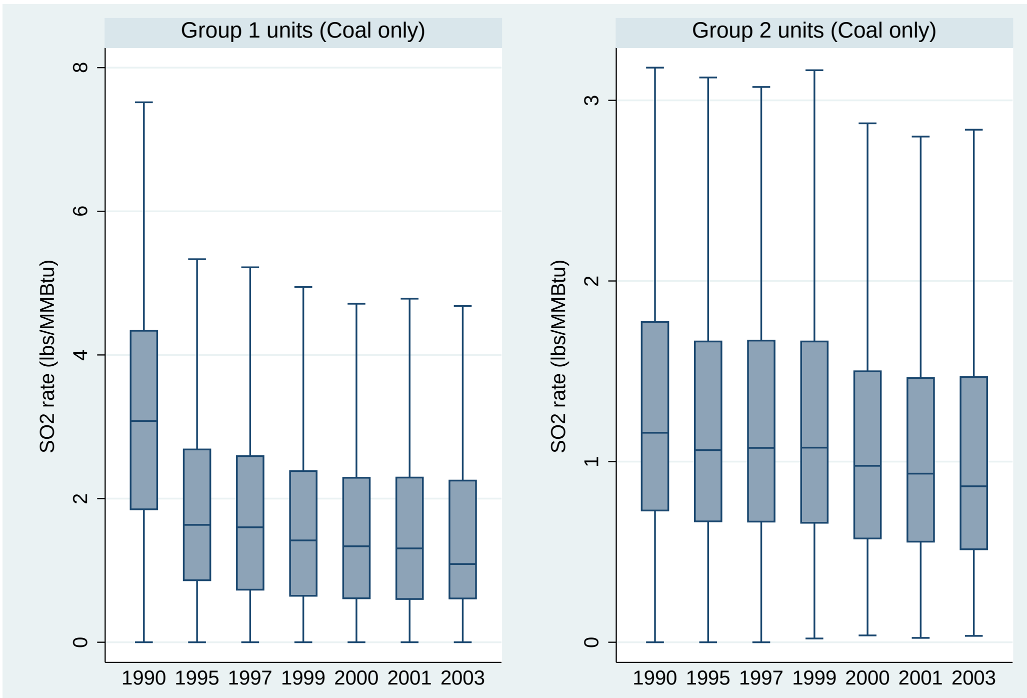

Note: The box shows the interquartile range of the distribution. The two lines correspond to the upper and lower adjacent values of the distribution.

1.2.3.3. Effects of a cap-and-trade on Output. I now examine whether firms reduced outputs in response to the regulation, which is another margin of emissions abatement. Here, I focus on the intensive margin of operation and treat entry/exit as given. Although retirement of coal units could be a potential option for emissions abatement, the data shows that this margin is small. Among the 263 EGUs in the "Table 1" group, only seven units retired before 1995, and two additional units retired before 2003. Regarding other coal units, around 6% of EGUs retired between 1990 and 2003.

Figure 1.3. Trend of Capacity Factor of Group I and Group II units

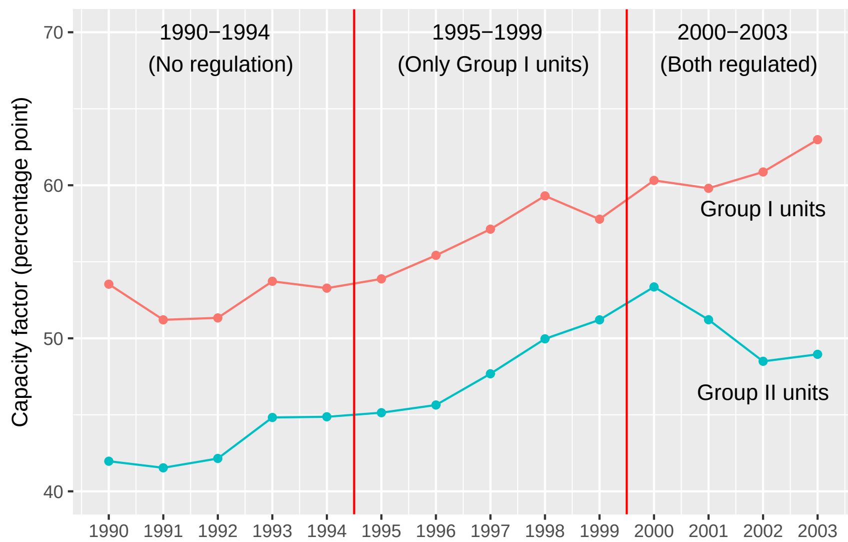

To estimate the effects of the ARP on production output, I exploit the variation of the timing of the regulation across units in a difference-in-differences (DID) framework. Figure 1.3 shows the trend of the capacity factor, defined by the ratio of net generation (output) to generation capacity, over time. I calculate the mean of the monthly level capacity factor in each year for two groups: those that are regulated from 1995 (Group I units) and those that are regulated from 2000 (Group II units). The figure shows that these two groups have a similar trend in their capacity factor from 1990 to 1994, which supports the parallel-trend assumption in the DID framework.

The regression equation I estimate is given by

$$cf_{jm} = \alpha_1 \text{GroupI}_j \cdot 1\{t \geq 1995\}_t + \alpha_2 \text{GroupII}_j \cdot 1\{t \geq 2000\}_t + \gamma X_{jm} + u_j + u_m + u_{jm},$$

where cf_{jm} is capacity factor of unit j in month m. The capacity factor is defined by $cf_{jm} = q_{jm}/k_j$, where q_j is net-generation and k_j is nameplate capacity. GroupI and GroupII are the dummy variables for each group. X_{jm} includes control variables such as fuel costs. Unit and time fixed effects are captured by u_j and u_m.

Regression results are shown in Table 1.1. I find that the introduction of the ARP decreased the capacity factor by 1 to 2.5 percentage points, which is statistically significant. This finding is consistent with the idea that introducing a cap-and-trade program increases marginal costs of production, because firms are facing opportunity costs of emissions under a cap-and-trade program. The increase in marginal costs thus decreases outputs of generating units under a cap-and-trade regulation. Although the effects are statistically significant, the economic significance of the effects seems to be limited. Because the mean of the capacity factor is within the range of 40-60 percentage points in my sample, electricity generation decreased by around 2%-6% due to the introduction of a cap-and-trade program. This magnitude is not as much as the decrease in emissions over time, as shown in section 1.2.3.1. Combined with the findings from the previous subsections, this regression analysis indicates that the abatement of SO_2 emissions was achieved primarily through the adjustment of emissions rates.

Table 1.1. Reduced-Form Model of Capacity Factor

	Dependent variable:			
	Capacity factor in pct-point			
	(1)	(2)	(3)	(4)
Treatment (Group I units)	−0.656	−2.115***	−1.141**	−2.517***
	(0.575)	(0.668)	(0.562)	(0.693)
Treatment (Group II units)	−4.002***	−2.615***	−2.363***	−1.082*
	(0.548)	(0.592)	(0.564)	(0.619)
log(fuel costs)			−11.214***	−11.168***
			(0.479)	(0.479)
log(electricity demand)	40.937***	40.986***	42.479***	42.494***
	(1.194)	(1.195)	(1.229)	(1.231)
Group-trend	No	Yes	No	Yes
Observations	306,727	306,727	252,223	252,223
Adjusted R^2	0.635	0.635	0.600	0.600

Notes: Unit-level dummies, year dummies, and month-of-year dummies are included. Standard errors are clustered at the unit level. $*p < 0.1; **p < 0.05; ***p < 0.01$

1.2.3.4. Firm-level Trading Information. I now explain how firms behaved in the market of emissions permits. Figure 1.4 shows the correlation between trading decisions and firm size, measured by the sum of nameplate capacity of units under the ARP. The left panel shows the unconditional probability of market transaction at the firm-year level, and the right panel shows trading experience in the sample period at the firm level.

The left panel shows that firms did not necessarily trade every year. The unconditional probability of conducting permit trading was 72%. The trading probability was positively correlated with firm size. This observation is also found in the context of the EU-ETS scheme (see, e.g., Jaraitė-Kažukauskė and Kažukauskas, 2015). Although this finding can be interpreted as suggestive evidence for the presence of fixed costs of transaction, firms do not need to conduct a transaction in every period, due to the storability of emissions permits. In the right panel, I show firm-level experience of market trading during the sample period. 86% of firms had at least one experience of trading with other firms in the sample period, although some firms, most of which are small, did not trade at all.

Figure 1.4. Trading Pattern at Firm Level

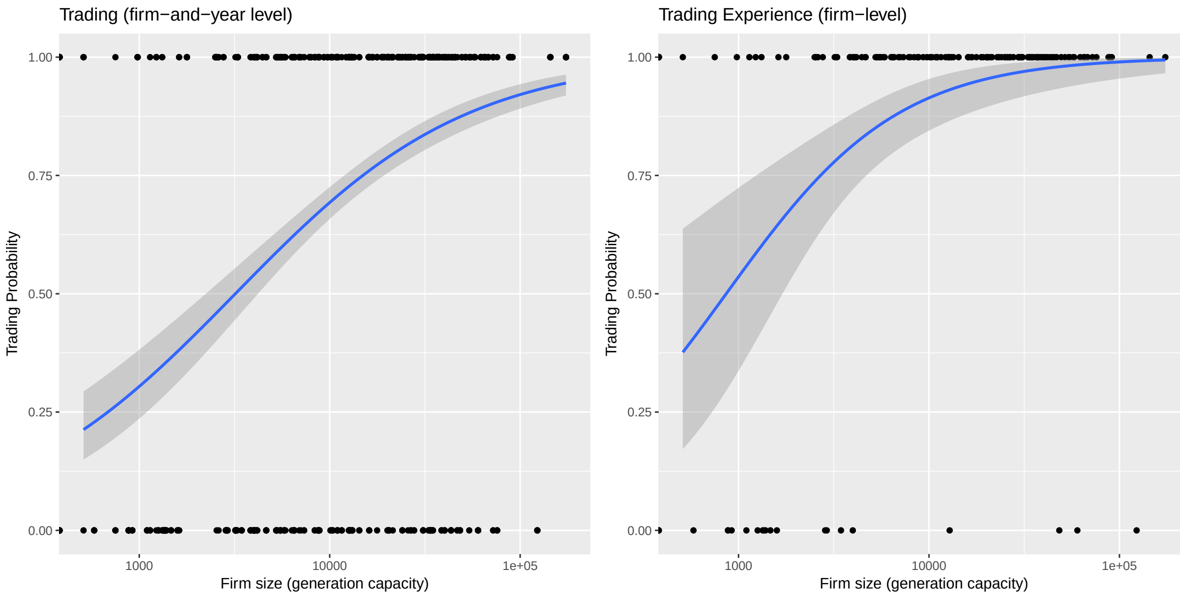

1.2.3.5. Price Data from a Broker. As I discussed in section 1.2.1, there is no centralized trading exchange for emissions permits under the Acid Rain Program. Although regulated firms need to have bilateral trades with other firms, brokers act as an intermediary for those transactions. Brokers also provide information about permit prices. Figure 1.5 shows price information provided by Cantor Fitzgerald, a broker in this market. I use the monthly SO_2 price index as a price measure in this paper. Cantor Fitzgerald constructs this index based on various trading information including the allowance bids (to buy), the allowance offers (to sell), and the actual trade prices. The company also posts this price information on the website in every month. I aggregate the monthly price index by taking the median for each year. Note that the price is also normalized to the level of 2000 by the Producer Price Index.

The price at the beginning was around 150 USD, and fell below 100 USD in 1996 and 1997. It increased to 200 USD in 1999, and fluctuated in the range of 120-200 USD after 2000. The figure suggests that the market price reflects the availability of banking. In the absence of permit banking, I would expect to see a spike in the permit price between Phase I and II, because the target emissions rate in Phase II is much stricter than in Phase I. Instead, the permit price has been gradually increasing over time, though it is volatile to some extent. [14]

Figure 1.5. Price of Emissions Permits by a Broker

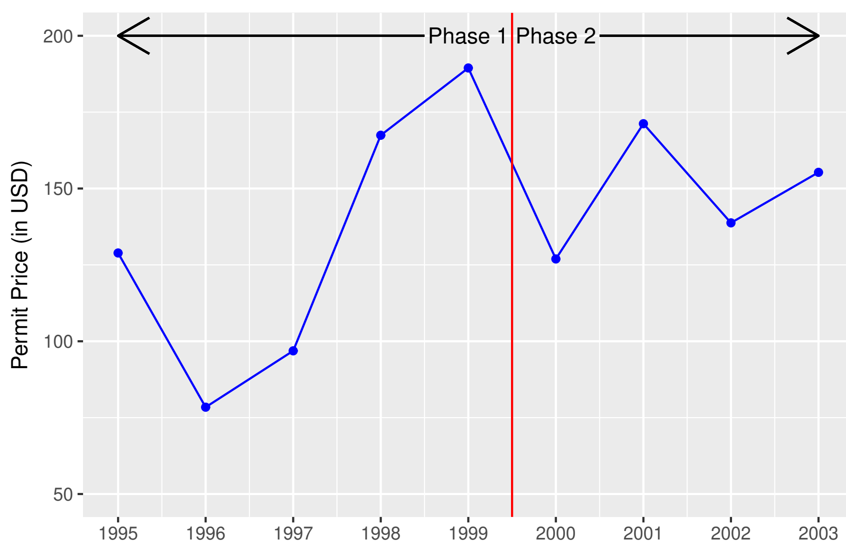

Note: Price is normalized to January 2000 by the Producer Price Index.

[14] A key theoretical prediction regarding permit prices is the Hotelling rule: permit prices should increase with the risk-free interest rate if the market is efficient and there are no transaction costs. Helfand et al. (2006) test the Hotelling rule by using the monthly prices of emissions permits in the same period, and reject the rule after controlling for structural changes and market shocks.

1.3. Model of Cap-and-Trade Program

1.3.1. Overview of the Model

This section introduces a structural model of the cap-and-trade program. My model is a discrete- and finite-horizon model indexed by $t = 1995, \ldots, 2003(\equiv T)$, and each discrete decision period corresponds to one compliance year. Firms have the common discount factor of β.

The overview of the model is summarized in Figure 1.6. The model has two building blocks: (i) investment in abatement options at the beginning of each phase (1995 and 2000), and (ii) decisions on production, trading, and banking in each year. At the beginning of each phase (1995 and 2000), firms make an investment decision on the abatement option and determine the emissions rate (R_i^1, R_i^2). The emissions rates are assumed to be fixed within each phase. This assumption reflects the observation from section 1.2.3 that the emissions rate changes at the beginning of each phase and stays constant within the phase.

Given the emissions rate, a firm makes decisions on production, permit trading, and banking. The timeline of each period is as follows:

(1) Firm i holds permits that are carried over from the previous period, denoted by h_{it}. A firm also receives annual allocation of permits denoted by a_{it}.

(2) Participation decision: Denote firm i's experience of market trading by I_{it}; i.e., $I_{it} = 1$ if a firm has experience in market trading, and 0 otherwise. If $I_{it} = 0$, a firm can pay the one-time sunk cost F_{it} to participate.

(3) A firm chooses (i) production quantity of each generating unit $\{q_{jt}\}_j$, (2) net volume of trading b_{it} if a firm already participated in the market, and (3) banking of permits $h_{i,t+1}$.

Figure 1.6. Model Overview

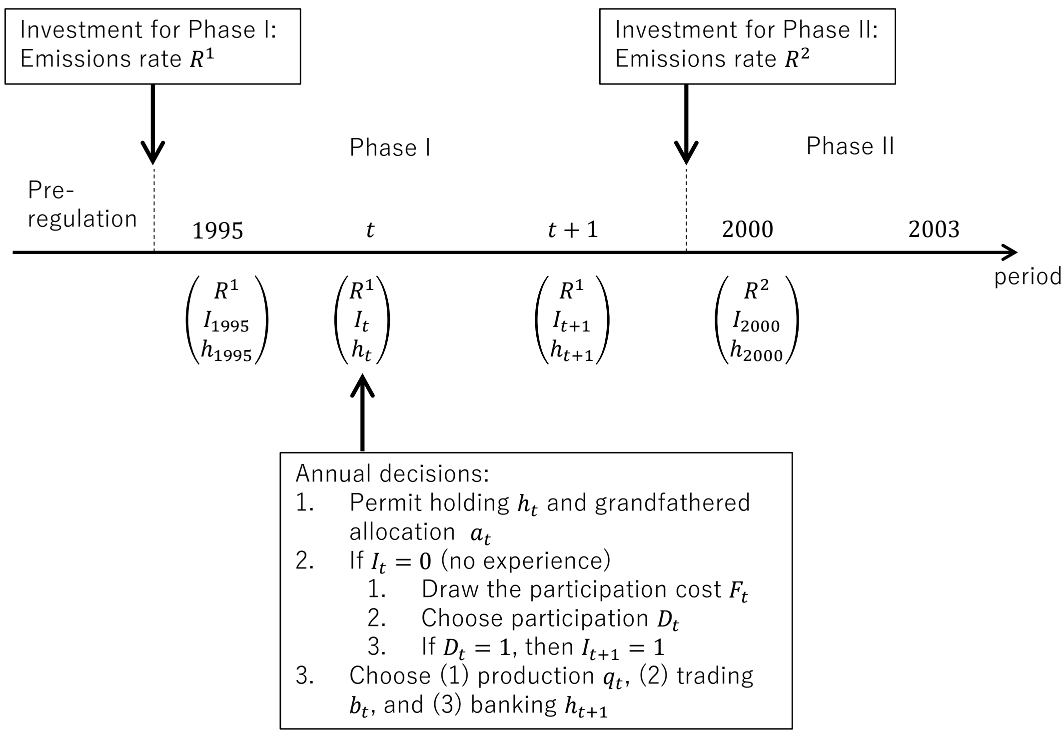

(4) A firm obtains profits from electricity generation and pays the costs of permits (or obtains the revenue from selling permits).

(5) Move to the next period with the holding $h_{i,t+1}$.

I now turn to explain each component of the structural model.

1.3.2. Electricity Production and SO_2 Emissions

Firms are earning profits from electricity production in a competitive electricity market. Firm i holds J_{it} units of the regulated sources and chooses the production quantity q_{jt} for each generating unit j. The profit is given by

$$\pi_{it}\left(\{q_{jt}\}_j\right) = \sum_{j \in J_{it}} \left\{ (\tau_{st}^{elec} - c_{jt}^{fuel}) \cdot q_{jt} - g(q_{jt}, k_j) \right\},$$

where τ_{st}^{elec} is the electricity price in state s where unit j is located, and c_{jt}^{fuel} is the unit-specific fuel costs of production. Fuel costs account for around 75% of total operating expenses (see EIA, 2012). $g(q_{jt}, k_j)$ is the convex cost of production. This term captures the increasing costs of operation near the capacity constraint (see, e.g, Ryan, 2012).

Electricity production is associated with SO_2 emissions. Firm-level emissions are given by

$$e_{it}\left(\{q_{jt}, \rho_{jt}\}_j\right) = \sum_{j \in J_{it}} \rho_{jt} q_{jt}, \tag{1.3.1}$$

where ρ_{jt} is the unit-level SO_2 emissions rate per production. I assume that ρ_{jt} is given by

$$\rho_{jt} = \begin{cases} HR_j \cdot R_i^1 & if\ j = coal\ \&\ t \in [1995, 1999] \\ HR_j \cdot R_i^2 & if\ j = coal\ \&\ t \in [2000, 2003] \\ HR_j \cdot R_{j,t}^{gas} & if\ j = gas\ or\ oil. \end{cases},$$

The unit-specific heat rate HR_j is an inverse of the production efficiency measure. HR_j represents how much fuel (in MMBtu) is needed to produce 1 MWh of electricity. R_i^1 and R_i^2 are the firm- and phase-specific SO_2 emissions rate, the emissions level per 1 MMBtu units of fuel. These emissions rates are endogeneously determined by the investment decisions at the beginning of each phase. I treat gas and oil units separately from coal units because these units have already low SO_2 emissions rates.

Note that the profit $\pi_{it}(\cdot)$ does not include the costs associated with emissions. Firms should take into account the cost of using emissions permits in their production decisions. As I show in section 1.3.4, the optimal decision on production quantity reflects emissions costs as well as the output price τ_{st}^{elec} and the fuel cost c_{jt}^{fuel}. The profit function $\pi(\cdot)$ is interpreted as the gross profit from electricity production that excludes costs associated with the permit trading.

1.3.3. Structure of Permit Trading and Transaction Costs

The role of cap-and-trade regulation is to penalize emissions from production activity and incentivize firms to reduce emissions. This subsection introduces the regulation into my model.

Each firm is allocated the annual allocation of permits a_{it} in each period. Because the allocation plan was announced before the regulation, the sequence of $\{a_{it}\}_t$ is exogenous in the model. The firm also holds the emissions permits that are carried over from the previous period, denoted by h_{it}. A firm decides emissions level e_{it}, which is determined by production quantity $\{q_{jt}\}$ as given by equation (1.3.1), net purchase volume b_{it}, and banking volume $h_{i,t+1}$. b_{it} is positive (or negative) if firm i is a buyer (or a seller), implying that she is buying (or selling) $|b_{it}|$ units of permits.

The transition of permit holding is given by

$$e_{it} + h_{i,t+1} = a_{it} + h_{it} + b_{it}, \tag{1.3.2}$$

$$h_{i,t+1} \geq 0. \tag{1.3.3}$$

Note that equation (1.3.3) is the non-negativity constraint of banking and excludes the possibility of borrowing of permits from future allocation of permits. I assume that firms achieve perfect compliance in my model. This is based on the fact that the compliance rate under this regulation is nearly perfect.

I model the permit market as a competitive market with transaction costs. The Acid Rain Program was a federal-wide program where many electric utilities, as well as financial companies, were participating. Exercising market power in the permit market was limited. [15] The presence of transaction costs reflects the fact that the vast majority of permit transactions were bilateral because there was virtually no centralized exchange for emissions permits. Incorporating bilateral trading of emissions permits into my model, however, is quite difficult because emissions permits are divisible objects and my model also features dynamic investment in clean technology and permit banking. I thus capture the nature of permit market by introducing transaction costs in a reduced form way.

Firms are price-takers in the permit market and face the market price P_t. In addition, they have to pay two types of transaction costs (see, e.g., Stavins, 1995). First, when a firm trades for the first time, it has to pay a sunk cost of participation F_{it}. This cost is motivated by the observation that some firms did not participate in the permit trading. An interpretation of F_{it} includes the costs associated with setting up a trading desk at the company and hiring a financial-trading expert. I specify F_{it} as the sum of a fixed cost F and an idiosyncratic cost ϵ_{it} given by

$$F_{it} = F + \epsilon_{it}, \epsilon_{it} \sim G(\cdot; \sigma_F),$$

[15] Liski and Montero (2011) examined how the four biggest electric utilities (in terms of initial allocation) trades in the permit market. They found that their behavior is not consistent with the model of market power in a storable commodity market.

where $G(\cdot;\sigma_F)$ is the cumulative distribution function of type I extreme value distribution.

Second, firms have to pay variable transaction costs associated with net purchase of permits b_{it}. This cost is given by

$$TC(|b_{it}|),$$

where $TC(\cdot)$ is a differentiable and strictly convex function. Variable transaction costs include brokerage commissions and bid-ask spreads. The convex nature of the cost function also captures the difficulty of large-scale transactions of emissions permits. Suppose that a firm wants to buy a certain amount of permits, but its trading partner cannot meet the demand. In such a case, a firm has to find another trading partner to buy more permits, and hence incurs a costly search process in a bilateral market. Convex transaction costs are employed in the theoretical literature in finance (e.g., Gârleanu and Pedersen, 2013, and Dávila and Parlatore, 2017) and also motivated by empirical findings (see, e.g., Breen et al., 2002, Lillo et al., 2003, and Robert et al., 2012).

In summary, the compliance costs (or revenue) from trading b_{it} units of permits is given by

$$P_t b_{it} + TC(|b_{it}|).$$

Introducing Fringe Firms. The sample I use does not cover all firms participating in permit trading. For example, financial companies or brokers do not have any generation facilities, thus I do not include them. Also some electricity companies are excluded from the sample in the process of data cleaning. I call them fringe firms in permit trading. To deal with the presence of fringe firms, I introduce the demand function of firms outside my sample. I denote the total net purchase from fringe firms by $\bar{B}_t^{fringe}(P_t)$. I explain a specification and estimation approach of the fringe demand function in section 1.4.3.

1.3.4. Optimal choices on production, trading, and banking

I now consider the optimization problems in year t. A firm makes both discrete (participation) and continuous decisions on production, trading, and banking. I first explain the decision problems conditional on the status of trading participation. These problems characterize the values from participation and non-participation, which determines the optimal participation decision.

Let V_{it}^1 and V_{it}^0 be the optimal values when a firm participates in trading ("trader") and does not ("non-trader"). The Bellman equation for the "trader" is given by

$$\begin{aligned} V_{it}^1(h_{it}, R_{it}) = \max_{\{q_{jt}\}_{j\in J_i}, b_{it}, h_{i,t+1}} \quad & \pi_{it}\left(\{q_{jt}\}_j\right) - \left(P_t b_{it} + TC(b_{it})\right) + \beta EV_{i,t+1}(h_{i,t+1}, 1, R_{i,t+1}) \\ \text{s.t.} \quad & e_{it}\left(\{q_{jt}, \rho_{jt}\}_j\right) + h_{i,t+1} = a_{it} + h_{it} + b_{it}, \\ & h_{i,t+1} \geq 0. \end{aligned} \tag{1.3.4}$$

$EV_{it}(h_{it}, I_{it}, R_{it})$ denotes the ex-ante value function for firm i in period t when the firm holds h_{it} units of emissions permits, the trading experience is I_{it}, and the emissions rate is R_{it}. Recall that $I_{it} = 1$ if firm i already participated in the market trading by paying the participation cost.

When a firm is a non-trader, it does not choose the trading volume b_{it} by definition. The Bellman equation in this case is

$$\begin{aligned} V_{it}^0(h_{it}, R_{it}) = \max_{\{q_{jt}\}_{j\in J_i}, h_{i,t+1}} \quad & \pi_{it}\left(\{q_{jt}\}_j\right) + \beta EV_{i,t+1}(h_{i,t+1}, 0, R_{i,t+1}) \\ \text{s.t.} \quad & e_{it}\left(\{q_{jt}, \rho_{jt}\}_j\right) + h_{i,t+1} = a_{it} + h_{it}, \\ & h_{i,t+1} \geq 0. \end{aligned} \tag{1.3.5}$$

Note that the value functions $\{V_{it}^0(\cdot), V_{it}^1(\cdot)\}$ are indexed by t, which is meant to include all state variables except for h_{it}, I_{it}, R_{it}, and ϵ_{it}. I assume perfect foresight over the state variable in the next period except for the shock to the participation cost ϵ_{it}.

The optimality conditions for the traders are given by

$$\tau_{st}^{elec} - c_{jt}^{fuel} - g'(q_{jt}) - \lambda_{it}\rho_{jt} = 0 \tag{1.3.6}$$

$$\lambda_{it} = P_t + TC'(b_{it}) \tag{1.3.7}$$

$$\lambda_{it} = \beta\frac{dEV_{i,t+1}(h_{i,t+1}, I_{i,t+1}, R_{i,t+1})}{dh_{i,t+1}} + \mu_{it}, \tag{1.3.8}$$

$$\mu_{it} \geq 0 \perp h_{i,t+1} \geq 0, \tag{1.3.9}$$

where λ_{it} denotes the Lagrange multiplier on the transition of permit holding (1.3.2) and μ_{it} denotes the Lagrange multiplier on the non-borrowing constraint (1.3.3). Note that λ_{it} is interpreted as the shadow price of emissions permits for firm i.

Equation (1.3.6) determines the optimal production decision given the shadow costs of emissions permits. The left-hand-side is the marginal profit that accounts for the cost associated with emissions $\lambda_{it}\rho_{jt}$. This condition can be also written as

$$\frac{\tau_{st}^{elec} - c_{jt}^{fuel} - g'(q_{jt})}{\rho_{jt}} = \lambda_{it},$$

implying that the marginal profit from additional emissions should be equal to the shadow costs of emissions permits λ_{it}.

Equations (1.3.7) and (1.3.9) determine the shadow costs λ_{it} and μ_{it} from the trading and banking decisions. Equation (1.3.7) says that the shadow price is equal to the sum of the market price and the marginal trading costs $TC'(b_{it})$. Equations (1.3.8) and (1.3.9) show that the shadow value of an emissions permit today is equal to the sum of the discounted marginal value of holding an additional permit tomorrow and the shadow value of borrowing (when it is binding). These conditions along with the transition equation of permit holdings determine the optimal choices for production $\{q_{jt}\}_j$, trading b_{it}, and the banking $h_{i,t+1}$.

The optimality conditions for the non-trader are the same as above except we do not have equation (1.3.7), and $b_{it} = 0$. These conditions implies that the shadow value of an emissions permit is not directly related to today's permit price in this case. Rather, the shadow value is given by the discounted marginal value from equation (1.3.8).

Next, I consider the participation decision. If a firm has no prior trading experience (i.e., $I_{it} = 0$), it can choose whether to participate in the market by paying $F_{it}(= F + \epsilon_{it})$. The optimal participation decision is given by

$$D_{it} = \mathbf{1}\left\{V_{it}^1(h_{it}, R_{it}) - (F + \epsilon_{it}) > V_{it}^0(h_{it}, R_{it})\right\},$$

and the participation probability is

$$\mathbb{P}_{it}(h_{it}, R_{it}) = \int \mathbf{1}\left\{V_{it}^1(h_{it}, R_{it}) - (F + \epsilon_{it}) > V_{it}^0(h_{it}, R_{it})\right\} dG(\epsilon_{it}).$$

If a firm already participated in trading (i.e., $I_{it} = 1$), it does not have to pay the participation costs.

Based on the optimal choices for traders and non-traders, I now provide the value function. Let $V_{it}(h_{it}, I_{it}, R_{it}, \epsilon_{it})$ be the value function after observing the random draw of the participation costs. The value function is given by

$$V_{it}(h_{it}, I_{it}, R_{it}, \epsilon_{it}) = \begin{cases} \max\left\{V_{it}^{0}(h_{it}, R_{it}), V_{it}^{1}(h_{it}, R_{it}) - (F + \epsilon_{it})\right\} & if\ I_{it} = 0 \\ V_{it}^{1}(h_{it}, R_{it}) & if\ I_{it} = 1. \end{cases}$$

Also, the ex-ante value functions (before observing ϵ_{it}) are

$$EV_{it}(h_{it}, I_{it}, R_{it}) = \begin{cases} \int \max\left\{V_{it}^{0}(h_{it}, R_{it}), V_{it}^{1}(h_{it}, R_{it}) - (F + \epsilon)\right\} dG(\epsilon) & if\ I_{t} = 0 \\ V_{it}^{1}(h_{it}, R_{it}) & if\ I_{t} = 1. \end{cases}$$

By applying the Williams-Daly-Zachary theorem and the envelope theorem (see Appendix A.3.1 for the derivation), the derivative of the expected value function with respect to the state variable h_{it} can be expressed as follows:

$$\frac{dEV_t(h_{it}, 0, R_{it})}{dh_{it}} = \mathbb{P}_{it}(h_{it}, R_{it})\lambda_{it}^{1} + \left(1 - \mathbb{P}_{t}(h_{it}, R_{it})\right)\lambda_{it}^{0}. \tag{1.3.10}$$

$$\frac{dEV_t(h_{it}, 1, R_{it})}{dh_{it}} = \lambda_{it}^{1}, \tag{1.3.11}$$

where λ_{it}^{1} and λ_{it}^{0} are the Lagrange multipliers on the transition constraint in the optimization problems (1.3.4) and (1.3.5), respectively.

Continuation Value at the Terminal Period. My model has a finite time period, and the terminal period T corresponds to the year 2003, which is the last period of my sample. However, the cap-and-trade program continued after 2003, and the banking at the end of 2003 was still substantial in the data. To deal with this issue, I introduce the reduced-form continuation value function $CV_{T+1}(h_{i,T+1}, R_i^2)$ in the model. This term captures the banking incentive at the terminal period $T(= 2003)$. In section 1.4.2, I provide the functional form of $CV_{T+1}(h_{i,T+1}, R_i^2)$ and and estimate it along with the other parameters.

1.3.5. Investment Decisions on Emissions Rate

I now introduce the investment decision on abatement options. In my model, a firm determines phase-specific emissions rates at the beginning of each phase:

$$R_{it} = \begin{cases} R_i^1 & t = 1995, \cdots, 1999 \\ R_i^2 & t = 2000, \cdots, 2003 \end{cases}.$$

The lower the emissions rate, the higher the level of investment in my model. I also assume that the emissions rate is a continuous choice variable. I denote the cost function of investment by $\Gamma(\bar{R} - R)$, where R is the emissions-rate level a firm chooses and $\bar{R}$ is the emissions rate before the investment.

The investment problem for Phase I is given by

$$\begin{aligned} \max_{R_{i,P1}} \quad & EV_{i,1995}(0, 0, R_i^1) - \Gamma(R_i^0 - R_i^1) \\ \text{s.t.} \quad & R_i^1 \leq R_i^0, \end{aligned} \tag{1.3.12}$$

where R_i^0 is the emissions rate in 1990, that is, before the regulation. I incorporate the adjustment costs and the irreversibility of investment by allowing R_i^0 to affect both the investment cost and the choice set of emissions rate R_i^1. Note that $h_{i,1995} = 0, I_{i,1995} = 0$ by definition.

The problem for Phase II is similarly defined as

$$\begin{aligned} \max_{R_{i,P2}} \quad & EV_{i,2000}(h_{i,2000}, I_{i,2000}, R_i^2) - \Gamma(R_i^1 - R_i^2). \\ \text{s.t.} \quad & R_i^2 \leq R_i^1 \end{aligned} \tag{1.3.13}$$

The investment cost now depends on R_i^1, which is endogenously determined in Phase I.[16]

1.3.6. Dynamic Competitive Equilibrium with Perfect Foresight

I now define an equilibrium for the permit market. I assume that firms have perfect foresight over the future environment and the only stochastic shock is the participation cost ϵ_{it}. [17]

Definition 1. In a finite-period competitive equilibrium with perfect foresight, a sequence of permit prices $\{P_t\}_{t=1995}^{2003}$ is determined such that

(1) [Optimization] Each firm i optimally chooses $\{\{q_{jt}^*\}_j, b_{it}^*, h_{i,t+1}^*\}_{t=1995}^{2003}$ and $\{R_i^{1*}, R_i^{2*}\}$ given a sequence of permit prices, and

(2) [Market Clearing] $\sum_i b_{it}^* + \bar{B}_t^{fringe}(P_t) = 0$ for $t = 1995, \cdots, 2003$.

I plan to provide a formal argument for the existence of a dynamic competitive equilibrium. My model is close to models of a dynamic competitive market as in Jovanovic (1982), Hopenhayn (1990), Hopenhayn (1992), and Cullen and Reynolds (2017). They show the existence of an equilibrium by providing a correspondence between the social planner's solution and a competitive equilibrium. Regarding uniqueness of equilibrium, I try different initial prices of emissions permits when I numerically solve a dynamic competitive equilibrium and

[16] Four types of firms exist: (1) those that own coal units in only the Phase I group, (2) those that own coal units in Phase I and II groups, (3) those that own coal units in only the Phase II group, and (4) those that only own gas or oil units. Equations (1.3.12) and (1.3.13) are the problems for the first type of firms. For the second type of firms, the pre-investment emissions level in Phase II is given by $\omega R_i^1 + (1-\omega) R_i^{0,Phase2}$, where $R_i^{0,Phase2}$ is the emissions rate of Phase II units in 1990 and ω is the share of Phase I units in firm i in terms of generation capacity. In the case of the third type of firm, the pre-investment emissions level in Phase II is $R_i^{0,Phase2}$. The last type of firms does not choose the emissions rate in the model.

[17] Incorporating aggregate uncertainty (i.e., allowing aggregate state variables such as fuel price to stochastically evolve) is a challenging task. Under this setting, the permit price P_t also becomes a random variable, and firms have to form an expectation over future price. Specifying a form of expectation is a major challenge because the permit price itself is an equilibrium object in my model. This setting is close to Krusell and Smith (1998) on a heterogeneous macro model, Cullen (2015) on a dynamic competitive equilibrium in electricity competition, Lee and Wolpin (2006) on structural estimation of a general equilibrium labor model, and Richards-Shubik (2015) on structural estimation of a peer-effect model.

find that these initial values converge to the same equilibrium prices. I leave a formal proof of equilibrium uniqueness to future revisions.

1.3.7. Discussions of Model

Output price τ_{st}. I assume that output price τ_{st} is given as exogenous throughout the analysis. Since the main target of the Acid Rain Program is coal units which are typically infra-marginal units, the program will not affect the electricity price in the wholesale market (See Fowlie, 2010b).

Interpretation of Continuation Value Function $CV_{T+1}()$. The continuation value function $CV_{T+1}(h_{T+1})$ captures the incentive to bank emissions permit at the terminal period in my model, namely 2003. To be more precise, I assume that $CV_{T+1}(\cdot)$ captures firms' incentive to bank under the expectation that the Acid Rain Program continues after 2004 without any additional regulations. This is a reasonable assumption given that the CAIR was announced in the last month of 2003 (i.e., December 2003), implying that, in 2003, firms were expecting that the same regulatory environment would continue.[18]

1.3.8. Model Implications

I discuss several implications from my structural model.

1.3.8.1. Role of Transaction Costs. I discuss three implications of transaction costs $TC(\cdot)$ and F that I introduced in the model: (i) the efficiency property of a cap-and-trade program, (ii) the independence property, and (iii) dynamic implications.

[18]Alternatively, I could model the terminal period as a stationary and infinite-period dynamic programming problem by assuming that the same regulatory environment continues afterwords and the CAIR would not be introduced. This approach allows me to avoid specifying a parametric form of the continuation value function.

To discuss these implications, I first provide optimality conditions whereby no transaction costs exist, namely, $T(\cdot) = 0$ and $F_{it} + \epsilon_{it} = 0$. In this case, equation (1.3.7) imply that $\lambda_{it} = P_t$ holds for all i, meaning all the firms have the same shadow price, which is given by the market price. Also, equation (1.3.8), along with envelope conditions (1.3.10) and (1.3.11), implies that $\frac{dEV_{i,t+1}(h_{i,t+1}, I_{i,t+1})}{dh_{i,t+1}} = P_{t+1}$. Summarizing these optimality conditions , we have

$$\frac{\tau_t^{elec} - c_{jt}^{fuel} - \frac{\partial g(q_{jt}, k_j)}{\partial q_{jt}}}{\rho_{jt}} = P_t. \tag{1.3.14}$$

$$P_t = \beta P_{t+1} + \mu_{it} \tag{1.3.15}$$

$$\mu_{it} \geq 0 \perp h_{i,t+1} \geq 0.$$

I now explain the three implications in turn.

Efficiency Property. One of the virtues of cap-and-trade regulation is that the allocation of emissions, given the emissions cap, is efficient in the absence of transaction costs. This assertion is confirmed by equation (1.3.14)'s implication that the marginal profit from producing one unit of emissions is equalized across firms at the level of permit price P_t. The key mechanism is that all firms are facing the same shadow value given by the market price P_t.

I now examine how the trading behavior affects the shadow costs of emissions permits and thus creates an inefficient allocation of emissions in the presence of transaction costs. Consider the case in which three firms exist: one is a buyer (i.e., $b_{buyer,t} > 0$), the other is a seller (i.e., $b_{seller,t} < 0$), and the last one is a non-trader. Equation (1.3.6) implies that

$$\begin{aligned}
\lambda_{buyer,t} &= P_t + TC'(b_{buyer,t}) > P_t \\
\lambda_{seller,t} &= P_t + TC'(b_{seller,t}) < P_t, \\
\lambda_{nontrader,t} &= \beta \frac{\partial EV_{t+1}(h_{nontrader,t+1}, I_{nontrader,t+1} = 0)}{\partial h_{nontrader,t+1}} + \mu_{nontrader,t}.
\end{aligned}$$

The inequalities in the first two lines hold because $TC'(b) > 0$ for $b > 0$ and $TC'(b) < 0$ for $b < 0$. Intuitively, in the presence of variable transaction cost, buyers have to pay additional costs to purchase emissions permits. By contrast, the revenue from selling a unit of emissions permits is the market price minus the marginal transaction costs. Thus, the marginal profit of emissions for the buyer is strictly higher than that for the seller. In other words, buyers produce less and sellers produce more than the efficient level at which the marginal profits of two firms are equalized. When a firm does not trade, the shadow cost of emissions permits is given by the discounted marginal value of permits tomorrow.

Independence Property. Another important property of a cap-and-trade program without any frictions (e.g., transaction costs) is that how the regulator allocates emissions permits has no effect on the pattern of emissions in an equilibrium. In equation (1.3.14), the initial allocation of permits a_{it} and permit holding h_{it} has no role in determining the production q_{jt}. This property is called the independence property of initial allocation of emission permits, which is an implication of the Coase theorem.

Once I introduce transaction costs, the independence property no longer holds: regardless of whether a firm participates in trading, h_{it} increases e_{it} and h_{it+1} and decreases b_{it} when a firm participates in trade. [19] Consider first the case in which a firm does not trade. When

[19] Note that both a_{it} and h_{it} have the same implication on the endogenous variables $\{e_{it}, b_{it}, h_{i,t+1}\}$ in period t.

the permit holding h_{it} increases marginally, the volume of banking $h_{i,t+1}$ also increases, which lowers the discounted marginal value of permit holding in the next period. Because a firm now has a lower shadow cost of emissions, it has an incentive to produce more emissions. Next, consider the case in which a firm participates in the permit market. An increase in permit holding decreases the volume of net purchase b_{it}. Because the variable transaction cost $TC(b_{it})$ is convex in b_{it}, the decrease in b_{it} lowers the marginal transaction costs a firm faces at the margin. Thus, the emission level e_{it} increases. [20] I discuss the detailed derivation in Appendix A.3.2.

Dynamic Implications. Dynamic implications are also different once I introduce transaction costs. I first explain the case without any transaction costs as a benchmark case. Equation (1.3.15) implies that the equilibrium permit prices P_t should increase at the rate of β^{-1} over time as long as banking volume is positive and no transaction costs exist. This property is known as the Hotelling $r-$percent rule, in which the price of exhaustible resources should increase at the rate of the inverse of the interest rate (see, e.g., Rubin, 1996). Another important implication is that the model does not pin down the individual optimal behavior for trading b_{it} and banking $h_{i,t+1}$ in the absence of the transaction costs, because the discounted marginal value from banking is constant and given by βP_{t+1}, which is equal to the current shadow value P_t in an equilibrium. Thus, marginal values of net purchase b_{it} and banking $h_{i,t+1}$ are always the same, and any choices are equivalent for individual firms as long as a firm can produce the level of emissions given by the optimality condition on production quantity.

[20] Stavins (1995) discusses the effect of a_t on emissions level e_t in other functional forms of the transaction cost function in a static model of emissions trading.

I now consider the case when transaction costs are present. Combining optimality conditions 1.3.7 and 1.3.8 and using envelope condition 1.3.11, I obtain the following condition:

$$P_t + TC'(b_{it}) = \beta \left\{ P_{t+1} + TC'(b_{it+1}) \right\} + \mu_{it},$$

which implies that the permit price does not necessarily increase at the rate of β^{-1}.

More importantly, the marginal values of net purchase b_{it} and banking $h_{i,t+1}$ are no longer constant in this setting. The marginal cost from net purchase is increasing due to the convex transaction costs $TC(b_{it})$. Intuitively, buying more permits becomes more difficult. The discounted marginal value from banking is decreasing because it is given by $\beta\{P_{t+1}+TC'(b_{it+1})\}$, and $b_{i,t+1}$ decreases in $h_{i,t+1}$. In other words, the value of permit holding in the next period is not constant, because firms have to pay the transaction costs so that their marginal revenue from selling is decreasing as they try to sell more. These observations imply that the model now pins down the optimal decisions on both net purchase $b_{it}$and banking volume $h_{i,t+1}$.

1.3.8.2. Incentives to Invest. I now examine how the incentive to invest is determined in my model. Using the envelope theorem, I calculate the marginal returns from reducing emissions rate R_{P1} as follows:

$$\begin{aligned} -\frac{\partial EV_{1995}}{\partial R^1} = &\sum_{t=1995}^{1999} \beta^{t-1995} \left(\lambda_{it} \cdot \sum_j HR_{jt} q_{jt}^* \right) + \sum_{t=1995}^{1999} \beta^{t-1995} \left(\sum_j \frac{\partial c_{jt}}{\partial R^1} q_{jt}^* \right) \\ &+ \beta^{2000-1995} \frac{\partial}{\partial R^1} \Gamma(R^1 - R^2). \end{aligned}$$

The first component is the returns from reducing emissions evaluated at the shadow value λ_{it}. The second component is the additional costs of using a cleaner fuel. Note that $\frac{\partial c_{jt}}{\partial R^1} < 0$ because fuel costs are higher for low-sulfur coals. The last component is the saving of investment costs in Phase II by investment in Phase I.

The primary component in the returns from investment is the first term. By reducing the emissions rate, the firm can marginally reduce emissions by $\sum_j HR_{jt}q^*_{jt}$. This marginal abatement is evaluated at the shadow value of λ_{it}. The returns from investment is thus given by the discounted sum of the returns from marginal abatement. The path of shadow values λ_{it} is key for investment incentives. As I discussed above, λ_{it} and equilibrium permit price P_t are affected by both banking and transaction costs.

First, the path of the shadow prices will be smoothed over the periods when banking is allowed, due to the optimality conditions on banking (1.3.8) that the current shadow value of emissions is the discounted shadow value in the future period. The smoother path of $\{\lambda_{it}\}$ implies that firms would invest more in the periods with generous emissions caps (Phase I) and then invest less later (Phase II) when banking is available.

Second, as I have shown above, the shadow value of emissions permits for buyers is higher than for sellers:

$$\lambda_{buyer} > P_t > \lambda_{seller}.$$

Thus, buyers have a higher incentive to invest, whereas sellers have a lower incentive. Intuitively, under the presence of transaction costs, buyers face higher shadow costs of emissions permits, making them prefer investing in technology rather than buying permits. Sellers, on the other hand, obtain lower revenue due to the transaction costs, which gives them lower incentives to invest.

1.4. Estimation Strategy

This section introduces the estimation strategy for the structural model. I conduct estimation in three steps. First, I estimate a reduced-form model for the capacity factor based on the optimality condition for production quantity q_{jt}. Using the estimated reduced-form model, I next estimate the variable transaction costs, $TC(b)$, the distribution of the fixed transaction costs, F_{it}, the continuation value at the terminal period, $CV_{T+1}(h_{i,T+1}, R_i^2)$, and costs of abatement investment, $\Gamma(\bar{R} - R)$. Note that I fix the annual discount factor at $\beta = 0.95$ throughout the paper. Finally, I estimate the fringe demand.

1.4.1. Step 1: Reduced-Form Model for the Capacity Factor

Recall that the FOC for unit-level production quantity q_{jt} is

$$\frac{\tau_t^{elec} - c_{jt}^{fuel} - \frac{\partial g(q_{jt},k_j)}{\partial q_{jt}}}{\rho_{jt}} = \lambda_{it},$$

which can be written as

$$\frac{\partial g(q_{jt}, k_j)}{\partial q_{jt}} = \tau_t^{elec} - c_{jt}^{fuel} - \lambda_{it}\rho_{jt}.$$

I consider the following reduced-form model for the optimal choice of q_{jt}

$$q_{jt} = \frac{\exp(\gamma(\tau_t^{elec} - c_{jt}^{fuel} - \lambda_{it}\rho_{jt}))}{1 + \exp(\gamma(\tau_t^{elec} - c_{jt}^{fuel} - \lambda_{it}\rho_{jt}))} \cdot k_j.$$

The first component of the right-hand side is the capacity factor as a function of the markup of electricity production, $\tau_t^{elec} - c_{jt}^{fuel} - \lambda_{it}\rho_{jt}$. [21]

By transforming the above model, we have

[21] The reduced-form model can be derived from the following functional form: $g(q, k) = \frac{1}{\gamma}\left(q \log(q) + (k - q)\log(k - q)\right).$

$$\log \frac{cf_{jt}}{1 - cf_{jt}} = \gamma \left(\tau_t^{elec} - c_{jt}^{fuel} - \lambda_{it}\rho_{jt} \right),$$

where $cf_{jt} \equiv q_{jt}/k_j$ is the capacity factor.

In empirical implementation, I use month-level observations instead of year-level observations. Also, the sample includes generation units that are not affected by the SO_2 regulation, because I include the data before 1995 (before the ARP started) as well as data for units that were not affected at that point (e.g., data of Phase II units before 2000). To accommodate these observations, I consider the following form of the regression equation indexed by month m:

$$\log \frac{cf_{jm}}{1 - cf_{jm}} = \gamma \left(\tau_m^{elec} - c_{jm}^{fuel} - \mathbf{1}\{SO_2reg\}_{jt} \cdot \lambda_{it}\rho_{jt} \right) + u_j + u_m + u_{jm}, \tag{1.4.1}$$

where u_j is a unit fixed effect, u_m is time fixed effects, and u_{jm} is an error term. The dummy variable $\mathbf{1}\{SO_2reg\}_{jt}$ takes the value of 1 if unit j is under the ARP in year t.

In equation (1.4.1), I observe output price τ_{jm}^{elec}, fuel costs c_{jt}, and emissions rate per production ρ_{jt} in the right-hand-side . However, I cannot directly observe the firm-level shadow costs λ_{it}, which are endogenously determined in the structural model. I thus proxy λ_{it} by using the optimality condition from the model. Equation (1.3.7) implies that if a firm i already participated in permit trading (i.e., $I_{it} = 1$ or $D_{it} = 1$), λ_{it} is given by

$$\lambda_{it} = P_t + TC'(b_{it}).$$

Here I consider a quadratic specification of $TC(b_{it})$ and specify $TC'(b_{it})$ as a linear function: $\theta_0 b_{it} + \theta_i size_i b_{it}$, where $size_i$ is the firm size measured by the sum of the generation capacity of firm i. Then, λ_{it} is written as $\lambda_{it} = P_t + \theta_0 b_{it} + \theta_1 size_i b_{it}$. Putting this equation into equation (1.4.1), I obtain

$$\log \frac{cf_{jm}}{1-cf_{jm}} = \tilde{\theta}_1(\tau_m^{elec} - c_{jm}^{fuel}) + \tilde{\theta}_2 \mathbf{1}\{SO_2 reg\}_{jt} P_t \rho_{jt}$$
$$+\tilde{\theta}_3 \mathbf{1}\{SO_2 reg\}_{jt} b_{it}\rho_{jt} + \tilde{\theta}_4 \mathbf{1}\{SO_2 reg\}_{jt} \cdot size_i b_{it}\rho_{jt}$$
$$+u_j + u_m + u_{jm},$$

where $\tilde{\theta}_1 = \gamma, \tilde{\theta}_2 = -\gamma, \tilde{\theta}_3 = -\gamma\theta_0$, and $\tilde{\theta}_4 = -\gamma\theta_1$.

The remaining concern is the endogeneity of ρ_{jt} and b_{it}, both of which are choice variables in the structural model. I use the pre-regulation SO_2 emissions rate $\rho_j^{1990} \equiv HR_j \cdot R_{j,1990}$ as an instrument for ρ_{jt}, where $R_{j,1990}$ is the emissions rate in 1990. Another instrument is the sum of initial permit allocation for other units within the same firm, $\sum_{k\in J_i, k\neq j} a_{kt}$. I exclude the initial allocation of unit j because the unit-level allocation might depend on the unobserved characteristics of that unit. I use two-stage least squares to estimate the above parameters.

1.4.2. Step 2: Estimation of Remaining Parameters

Estimation in step 1 gives me the profit function $\pi_{it}(\{q_{jt}\}_j)$. The next step is to estimate the remaining parameters including transaction costs, the continuation value, and investment costs. I first provide specifications for these primitives.

My model contains two type of transaction costs: variable costs and participation costs. The variable transaction cost function $TC(|b|)$ is specified as follows:where $size_i$ denotes firm i's size measured by the sum of the generation capacity of firm i. Participation cost F_{it} is specified as $F_{it} = F + \epsilon_{it}$ where ϵ_{it} follows i.i.d. type 1 extreme value distribution with standard deviation σ_F.

I consider the following parameterization of the continuation value in the terminal period:

$$CV(h_{i,T+1}, R_i^2) = \exp\left(\alpha_0 + \alpha_1 \log(size_i) + \alpha_2 R_i^2\right) h_{i,T+1}^{\alpha_3}.$$

The coefficient depends on the firm size, $size_i$, and the emissions rate in Phase II, R_i^2. These variables capture the heterogeneity in the incentives to bank in the terminal period.

The specification for the investment cost $\Gamma(\cdot)$ is given by

$$\Gamma(\bar{R} - R) = \frac{exp(\zeta_0 + \zeta_1 \log(K_{i,\tau}))}{2}(\bar{R} - R)^2,$$

where $K_{i,\tau}$ is the generation capacity of units regulated in Phase $\tau \in \{I, II\}$. The parameters I estimate in this step are summarized by $\boldsymbol{\theta} = (\eta_0, \eta_1, F, \sigma_F, \alpha_0, \alpha_1, \alpha_2, \alpha_3, \zeta_0, \zeta_1)$.

I use simulated nonlinear least squares to estimate the model parameters. For a given candidate of parameter $\boldsymbol{\theta}$, I solve the model to obtain the prediction of choice variables and match the prediction with the data. However, solving a dynamic competitive equilibrium for each candidate of parameters is computationally infeasible. Instead, I use the observed prices of emissions permits as an equilibrium price and solve only the single-agent dynamic problems to get the model predictions of individual behavior.

This estimation approach builds on the literature of estimation of dynamic structural models in industrial organization and labor economics.[22] My structural model of a cap-and-trade regulation belongs to the class of dynamic competitive equilibrium models. Unlike Lee and Wolpin (2006), however, I can avoid solving for a dynamic competitive equilibrium in estimation, because I can use the observed prices of emissions permits as a sequence of

[22]Examples are Rust (1987), Hotz and Miller (1993), and Aguirregabiria and Mira (2002) for a single-agent dynamic discrete choice model, Lee and Wolpin (2006) for a dynamic competitive equilibrium model, and Aguirregabiria and Mira (2007) and Bajari et al. (2007) for dynamic Markov games. See Aguirregabiria and Mira (2010) for a survey of this literature.

equilibrium prices. The observed prices are fed into solving the single-agent optimization problems, which are much easier to solve than the dynamic competitive equilibrium, to construct the objective function in estimation. [23]

The procedure of obtaining the model prediction is as follows.

(1) Fix a candidate of parameter $\boldsymbol{\theta}$ and the observed permit prices $\{P_t\}_{t=1995}^{2003}$.
(2) For each firm i , solve the optimization problem by backward induction and obtain the policy functions for emissions $e_{it}(h_{it}, I_{it}, R_{it})$, trading $b_{it}(h_{it}, I_{it}, R_{it})$, and banking $h_{it+1}(h_{it}, I_{it}, R_{it})$, the participation probability $P_{it}(h_{it}, R_{it})$, and the investment decisions $R_i^1(h_{i1995}, I_{i1995})$, $R_i^2(h_{i,2000}, I_{i,2000}, R_i^1)$.
(3) Using the policy functions, simulate the optimal decisions for each pattern of participation in permit trading. I denote the year of participation by $s \in \{\emptyset, 1995, \cdots, 2003\}$, where $s = \emptyset$ means that a firm does not trade at all in the period. Denote the optimal decision for pattern s by $\hat{x}_{it}(s)$.
(4) Calculate the probability that each pattern of participation timing is realized. Denote this probability by $Prob_{it}^{enter}(s)$.
(5) The prediction for firm i in year t is then given as

$$\hat{x}_{it} = \sum_{s \in \{\emptyset, 1995, \cdots, 2003\}} Prob_i^{enter}(s) \hat{x}_{it}(s). \tag{1.4.2}$$

Using the simulated choices, I construct the following objective function given by

$$J(\boldsymbol{\theta}) = J_1(\boldsymbol{\theta}) + J_2(\boldsymbol{\theta}).$$

[23]This empirical strategy is similar in spirit to that in the two-step estimation of a dynamic game, in which the equilibrium objects are directly recovered from the observed data. For example, Aguirregabiria and Mira (2007) estimates players' beliefs over other players' policies from the observed data and solve the optimal response of a player given the estimated beliefs to construct the pseudo-likelihood function.

The first component $J_1(\boldsymbol{\theta})$ is minimizing the distance between the prediction and the data at the firm-and-year level

$$J_1(\boldsymbol{\theta}) = \sum_{i=1}^{N} \left(\mathbf{x}_i^{data} - \hat{\mathbf{x}}_i(\boldsymbol{\theta})\right)' \Omega_i \left(\mathbf{x}_i^{data} - \hat{\mathbf{x}}_i(\boldsymbol{\theta})\right),$$

where

$$\mathbf{x}_i^{data} = (e_{i,t_{i1}}, \cdots, e_{i,2003}, b_{i,t_{i1}} \cdots, b_{i,2003}, h_{i,t_{i1}+1} \cdots, h_{i,2004}, D_{i,\emptyset}, D_{i,t_{i1}} \cdots, D_{i,2003})$$

and $\hat{\mathbf{x}}_i(\boldsymbol{\theta})$ is the corresponding vector for the model prediction given parameter $\boldsymbol{\theta}$. t_{i1} is the first year of observations for firm i; that is $t_{i1} = 1995$ if firm i owns Group I units (those that are under the ARP from 1995), and $t_{i1} = 2000$ if firm i owns only Group II units (those that are under the ARP from 2000). $D_{i,s}$ is a dummy variable that indicates the timing of firm i's participation; $D_{i,s} = 1$ if firm i enters in year s, and, otherwise, 0; and $D_{i,\emptyset} = 1$ if firm i does not trade at all in the sample period. The weighting matrix Ω_i is a diagonal matrix to adjust for differences in scaling.[24]

The second component $J_2(\boldsymbol{\theta})$ incorporates the market clearing conditions, and is given by

$$J_2(\boldsymbol{\theta}) = \frac{1}{T} \sum_{t=1995}^{2003} \left(\sum_i \hat{b}_{it} - \sum_i b_{it}^{data} \right)^2.$$

[24]The weighting matrix is set as

$$\Omega_i \equiv \text{diag}\left(\underbrace{\hat{\mathbf{V}}(e), \cdots, \hat{\mathbf{V}}(e)}_{T_i}, \underbrace{\hat{\mathbf{V}}(b), \cdots, \hat{\mathbf{V}}(b)}_{T_i}, \underbrace{\hat{\mathbf{V}}(h), \cdots, \hat{\mathbf{V}}(h)}_{T_i}, \underbrace{\hat{\mathbf{V}}(D), \cdots, \hat{\mathbf{V}}(D)}_{T_i+1} \right),$$

where $\hat{\mathbf{V}}(e)$, $\hat{\mathbf{V}}(b)$,$\hat{\mathbf{V}}(h)$, and $\hat{\mathbf{V}}(D)$ are sample variances of emissions, net purchase, banking, and participation in the sample, respectively.

This component of the objective function requires that the estimated parameters are such that the observed prices are close to clearing the market in each period. Note that the sum of net purchases in the data $\sum_i b_{it}^{data}$ may not necessarily be equal to zero, since the sample does not cover all the firms participating in the permit trading.

Standard errors are calculated by bootstrap at the firm-history level. I randomly draw samples of 85 firms with replacement and construct 40 bootstrap samples.

1.4.3. Step 3: Estimation of Fringe Demand

I now estimate the fringe-demand function. I consider the following specification with constant elasticity:

$$\bar{B}_t^{fringe} = -\alpha_t P_t^{\epsilon}$$
$$\Longleftrightarrow \log\left(-\bar{B}_t^{fringe}\right) = \epsilon \log P_t + \log(\alpha_t).$$

One of the difficulties in estimating the fringe function is the paucity of data points. I have only 9 data points, because the sample is between 1995 to 2003, as shown in Figure 1.7, where I plot permit prices against the fringe demand for each year. I estimate the model by OLS and IV. I use the sum of initial allocations for firms in my sample as an instrument for P_t. I plan to provide a robustness check by using different levels of price elasticity and using different specifications for the fringe function (including a linear specification and a semi-log specification).

Figure 1.7. Plot of Permit Prices P_t and Fringe Demand $\bar{B}_t^{fringe}$

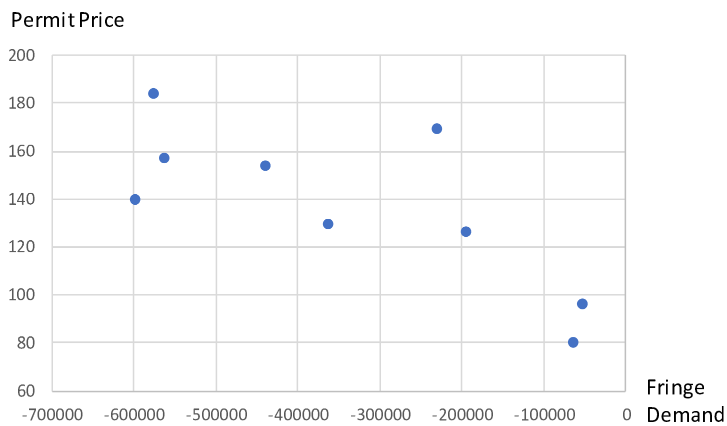

1.5. Estimation Results

1.5.1. Parameter Estimates

Table 1.2 presents the parameter estimates of the structural model. The model incorporates two types of transaction costs: participation and variable costs. Regarding the variable transaction costs, the coefficient on the firm size is negative but small. This estimate implies that although bigger firms tend to have lower transaction costs, heterogeneity across firms is negligible. Based on the parameter estimates, I calculate the marginal transaction cost, given by $exp(\eta_0 + \eta_1 \log(size_i))|b_{it}|$. The mean of the costs is \$193 and the median is \$98. Considering that the range of the permit prices is within \$100 to \$200 , these numbers are substantial. The mean of the fixed costs of participation is around \$38,000, which is quite small. The parameters in the continuation value function at the terminal period have signs that are intuitively reasonable. Bigger and dirtier firms obtain a higher value from banking. Estimates of the investment cost also have the reasonable sign. The bigger the capacity, the higher the investment costs.

Table 1.2. Parameter Estimates

	Parameter	Description	Estimate	Standard Errors
Production Parameters:	γ	Curvature	4.333e-03	1.149e-03
$g(q,k) = \frac{1}{\gamma}\left(q\log(q) + (k-q)\log(k-q)\right)$				
Variable Costs: $TC(b) =$	η_0	Constant	-3.516	0.343
$\frac{1}{2}exp(\eta_0 + \eta_1 log(size))\|b\|^2$	η_1	Firm size	-0.011	0.011
Participation Costs: F_{it}	F	mean (1 USD)	38,379	148,846
$logit(F, \sigma_F)$	σ_F	SD (1 USD)	1,259	12,405,130
Continuation Value:$CV(h_{i,T+1}, R_i^2)$	α_0	Constant	4.165	2.267
$\exp(\alpha_0 + \alpha_1\log(size_i) + \alpha_2 R_i^2)h_{i,T+1}^{\alpha_3}$	α_1	Firm Size	0.898	0.197
	α_2	Emissions Rate	0.036	0.183
	α_3	Curvature	0.314	0.053
Investment Costs:$\Gamma(R_{i,\tau}, \bar{R}_{i,\tau})$	ζ_0	Constant	11.725	1.137
$\frac{1}{2}exp(\zeta_0 + \zeta_1\log(K_{i,\tau}))(\bar{R}_{i,\tau} - R_{i,\tau})^2$	ζ_1	Capacity	0.697	0.251

Table 1.3. Parameter Estimates on Fringe Demand

	[1] log-log		[2] log-log		[3] linear		[4] linear	
Price	2.72	(0.52)	2.39	(0.77)	4386.5	(1184.91)	5106.0	(1275.42)
Phase II dummy	0.34	(0.33)	0.40	(0.43)	83197.5	(108217)	68423.0	(108363)
Constant	-0.99	(2.58)	0.57	(3.54)	-297448.9	(150911)	-389813.9	(170423)
Method	OLS		IV		OLS		IV	

Note: Standard errors are shown in the brackets.

Table 1.3 shows the estimates of the fringe elasticity. In OLS, the elasticity is estimated to be 2.72, while the estimate in the IV specification is 2.39. In the counterfactual analysis, I use 2.39 as a benchmark parameter. I plan to provide a robustness check on this parameter.

Figure 1.8. Model Fit of Equilibrium Permit Prices

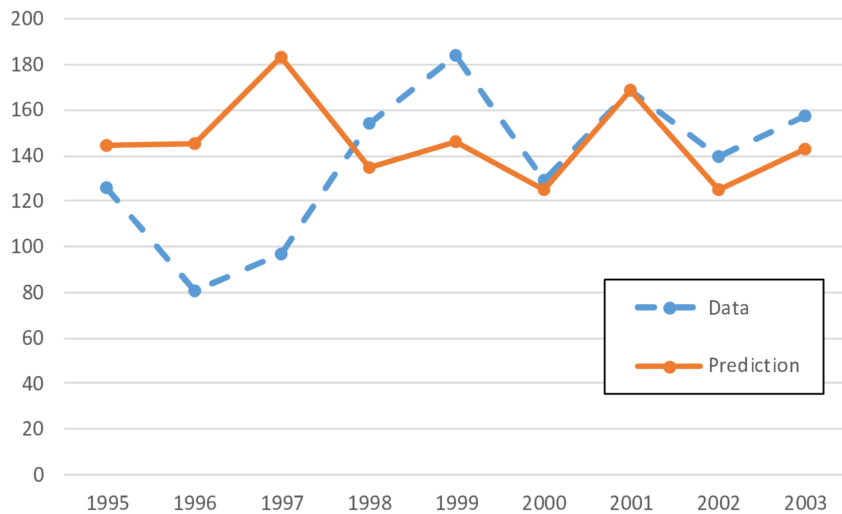

1.5.2. Model Fit

This subsection discusses the model fit under the estimated parameters. I first solve a dynamic competitive equilibrium under the estimated parameters and obtain model predictions on permit prices and individual behavior. Appendix A.1 explains the algorithm for solving a competitive equilibrium in detail.

Figure 1.8 shows the predicted and observed prices in the data. The blue dashed line corresponds to the data and the orange real line to the model prediction. Although the path of the predicted equilibrium prices deviates from the observed one in several years, 1996, 1997, and 1999, my model predicts the price quite well in Phase II (2000-2003). One potential reason for the worse model fit in the early periods is the lack of experiences in permit trading. Once firms accumulate sufficient experiences, the market price of emissions permits in Phase II becomes close to the prices predicted by my equilibrium model.

Table 1.4 and Figure 1.9 show the model fit in terms of individual behavior under the predicted equilibrium prices. The predictions on emissions e_{it} and banking h_{it+1} are quite

Table 1.4. Model Fit

	Emissions e_{it}		Net Purchase b_{it}		Banking $h_{i,t+1}$	
	Data	Prediction	Data	Prediction	Data	Prediction
Min	849	0	-335,072	-55,375	7	0
1st quantile	18,685	20,232	-3,916	-1,440	7,652	6,751
Median	47,283	48,222	0	2,131	30,283	24,487
Mean	89,596	85,116	5,764	5,411	75,251	83,183
3rd quantile	109,114	104,339	12,652	11,470	79,356	84,081
Max	727,040	910,781	351,702	46,039	1,204,817	1,711,434

Figure 1.9. Histogram of the Model Prediction and the Data

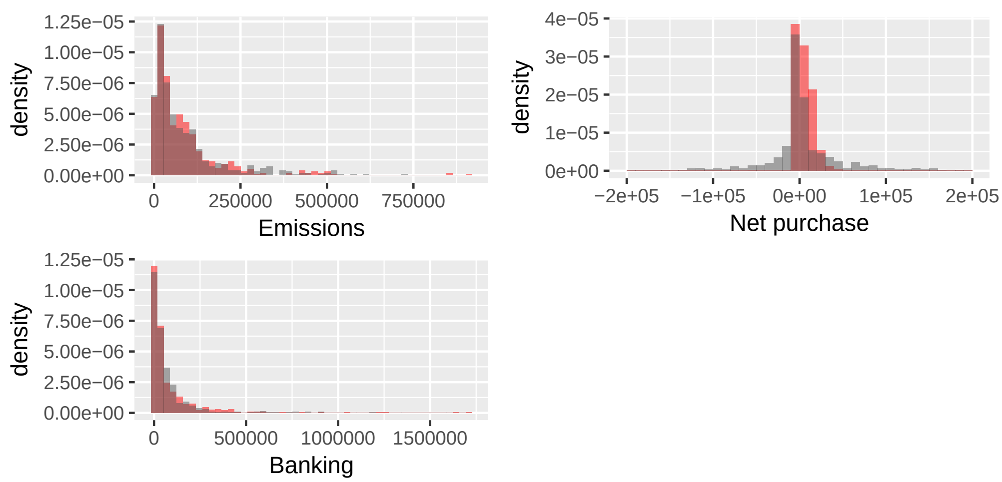

Note: Gray histogram corresponds to the data, and red histogram corresponds to the model predictions.

close to the observed distribution. Although both lower and upper tails of the distribution of net-purchase are different between the prediction and the data, my model predicts the median and mean net purchase quite well.

In the counterfactual simulations in section 1.6, I refer to the equilibrium outcome I solved here as a baseline outcome.

1.6. Counterfactual Experiments

This section discusses a series of counterfactual exercises using the model with the estimated parameters from section 1.5. In section 1.6.1, I first examine the effect of the cap-and-trade program and decompose it into permit trading and permit banking. I then quantify potential gains from trade by simulating the market outcome in the absence of transaction costs in section 1.6.2.

In counterfactual simulations, I fix model primitives at the estimated values. This assumption could be problematic for the continuation value at the terminal period $CV_{T+1}(\cdot)$. The continuation value captures the incentive of permit banking at the terminal period, which could change in the counterfactual situations. An alternative approach would be to fix the volume of permit banking at the terminal period, instead of the continuation value function. I will work on this approach in the future revision as a robustness check.[25]

The supply function from fringe firms $\bar{B}_t^{fringe}(\cdot)$ could potentially change in the counterfactual scenarios as well. In the simulation below, I fix the *level* of the fringe supply, instead of the function itself, in each year at the level under the equilibrium I solved in section 1.5.2. Under this assumption, I fix the total number of emissions permits available for firms in my sample, including their initial allocation and the fringe supply.

1.6.1. Experiment 1: Effects of a Cap-and-Trade Program and its Decomposition

I simulate the outcomes in the following two situations. First, I simulate the case in which all firms are required to achieve the uniformly determined level of emissions rate in each phase.

[25] Another approach is to model the terminal period as a stationary and infinite-period dynamic programming problem. This approach allows me to avoid specifying a parametric form of the continuation value. However, this approach is computationally more demanding because I have to solve the value function by a contraction mapping for each firm. I plan to take these approaches in the future revision as a robustness check.

The target emissions rates are defined as the rates under which the aggregate emissions are equal to the level of the baseline outcome, which I calculated in section 1.5.2. This simulation outcome corresponds to the one under uniform standard regulation. The difference between this outcome and the baseline outcome quantifies the effects of a cap-and-trade program. Second, I simulate a market outcome whereby permit banking is completely banned. This simulation allows me to decompose the effects of a cap-and-trade program into the effects of permit trading and the effects of permit banking. Appendix A.2 explains how to simulate the equilibrium outcome in each case.

Effects of a Cap-and-Trade in Comparison to Uniform Standards. Table 1.5 presents the results of counterfactual experiments. The numbers shown in the table are the totals in my sample period. The upper panel shows the aggregate levels of emissions, left-over permits, and banking at the terminal period. In the absence of a permit banking system, emissions permits that firms have at the end of the period would expire. I call these permit left-over permits. The lower panel shows the welfare measures including abatement costs and health and environmental damages. I explain the definition of health and environmental damages below.

I first discuss the effects of the cap-and-trade program in comparison to the uniform standard. The simulation results show that the cap-and-trade decreases the total costs of abatement, including investment costs and loss of profit from electricity generation, by 3.1 billion USD in total, or by 16.6%. The cost saving is allowed by a more flexible pattern of investment as shown in Figure 1.10. In the case of uniform standard, there is a mass of firms at 2.26 lbs/MMBtu in Phase I and 1.31 lbs/MMbtu in Phase II. These numbers are the emissions rates that achieve the same aggregate emissions as the baseline in Phase I and II,

Table 1.5. Effects of a cap-and-trade and its Decomposition

		Baseline	Uniform Standard	No Banking
Emissions (in 1 million ton)		45.54	45.54	41.60
Left-over permits		0.00	n.a.	9.15
Banking at the terminal period		5.21	n.a.	0.00
Firm costs	Investment Costs	15,179	18,648	15,934
(in million USD)	Loss of Electricity Profit	380	0	869
	Total costs	15,559	18,648	16,803
	change from baseline		3,090	1,245
Health and environmental damages	Total	54,787	54,592	50,692
(in million USD)	change from baseline		-195	-4,095
Total Costs	Total costs	70,346	73,241	67,495
(in million USD)	change from baseline		2,895	-2,850
Firm costs / abatement (in USD)		376	450	371
Health and environmental damage / emissions (in USD)		1,203	1,199	1,219

Note: The numbers are the total from 1995 to 2003. The unit of emissions, left-over permits, and banking at the terminal period is 1 million SO_2 tons. The unit of costs are 1 million USD in 2000.

respectively.[26] Introduction of the cap-and-trade program allows firms to optimally decide their emissions rate based on the returns and costs from investment, which depend on the characteristics of firms, permit prices, and transaction costs in the permit trading. Their optimal decisions lead to a more heterogeneous distribution of emissions rates and significant cost saving, compared to the uniform standard.

Implications for Health and Environmental Damages. A potential concern of a cap-and-trade program is the implication for health and environmental damages. Even though the aggregate level of emissions is fixed, the distribution of emissions under a cap-and-trade might be different from the one under a uniform standard regulation (See, e.g., Muller

[26]Some firms have the lower emissions rate than the uniform rate in the figure. This is because those firms already had satisfied the uniform standard before the regulation.

Figure 1.10. Distribution of Emissions Rate in the baseline and uniform standard

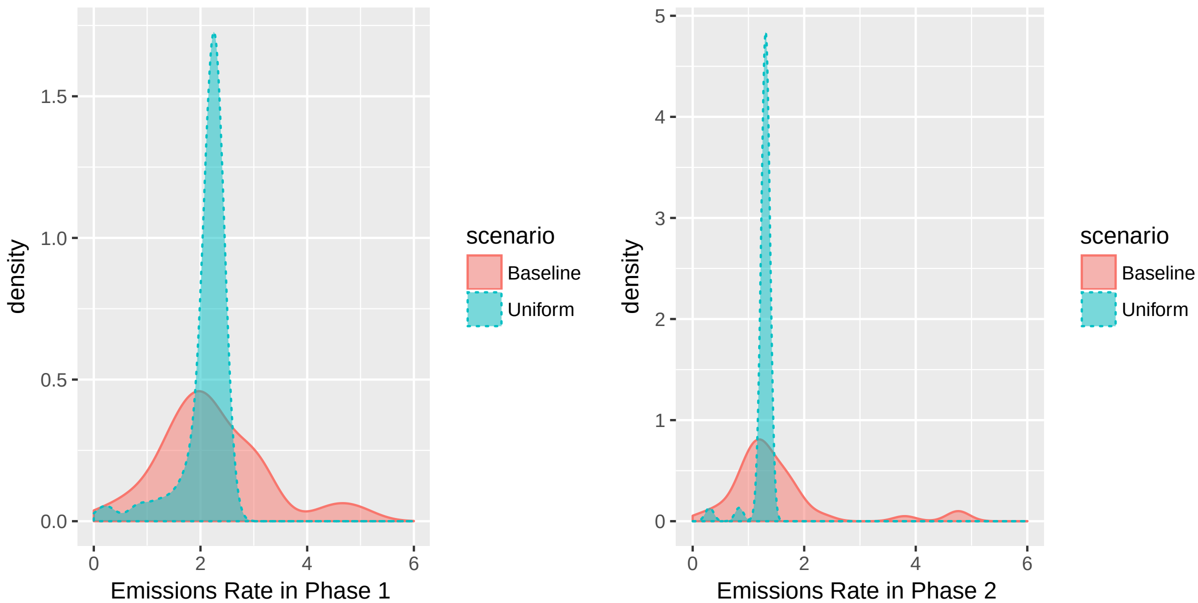

Note: Emissions rate is measured by lbs per MMBtu.

and Mendelsohn 2009, Fowlie et al. 2012, Fowlie and Muller 2013, and Chan et al. 2015). In particular, SO_2 emissions are non-uniformly mixed pollution; health and environmental damages depend on the location of the emissions' source. To the extent that the damage from a particular location is positively correlated with abatement costs of power plants in that location, a cap-and-trade program would increase damages from emissions in comparison to a uniform standard.

To discuss the net benefit of a cap-and-trade program, I calculate health and environmental damages in each case in Table 1.5 using the data from Muller and Mendelsohn (2009). They use the AP2 model, an integrated assessment model, to calculate marginal damages from SO_2 emissions at the county level.[27][28] Following Muller and Mendelsohn (2009), I assume that damages are linear in SO_2 emissions. The emissions damage from a particular county is given by the product of the marginal damage and the total SO_2 emissions from electricity plants that locate in the county.[29]

I find that a cap-and-trade program increases damages from SO_2 emissions slightly by 195 million USD (or 0.3%) in comparison to the uniform standard. This percentage increase of 0.3% is within the range of findings in Chan et al. 2015.[30] Overall, the savings of firm costs are large enough to offset the increase in damages, and thus the cap-and-trade program improves the overall welfare.

Effects of Permit Banking. The third column in Table 1.5 shows the outcome in the absence of a permit banking system. The aggregate emissions are different from the baseline due to the presence of left-over permits. Around 18% of permits would expire if a permit banking system is not available. Note that emissions permits might expire in the following two cases. If a firm does not participate in the permit trading, they cannot sell permits and thus the remaining permits must expire. Even though they participate, the marginal revenue from selling permits could be lower than zero due to transaction costs. In such a case, firms do not sell all of the remaining permits.

[27] I use the marginal damages from point sources with effective height less than 250 meters.

[28] The AP2 model is used in various papers, including Fowlie and Muller (2013), Chan et al. (2015), and Holland et al. (2016), for calculating health and environmental damages of air pollutants.

[29] One important limitation of the marginal damage data from Muller and Mendelsohn (2009) is that information about which location is hurt by emissions from a particular county is not available. To obtain this information, I need to run the AP2 model to simulate the flow of SO_2 emissions damages from one county to another county. This additional data would allow me to discuss distributional implications of a cap-and-trade program. I leave this analysis to future work.

[30] Chan et al. 2015 focus on 2002 to evaluate the net benefit of the Acid Rain Program. They find that the effect of a cap-and-trade program on health and environmental damages is within the range -0.4% to 1.8%, depending on the model specifications and the counterfactual policy in the absence of a cap-and-trade program.

The total abatement costs are estimated to be 16.8 billion USD in the absence of permit banking. Using this estimate, I decompose the savings of abatement costs into permit trading and permit banking. I find that permit banking decreases the abatement costs by 1.25 billion USD, accounting for around 26% of the effect of cap-and-trade. However, I need to take into account the difference in the aggregate emissions. To do so, I calculate the average abatement costs in each case. I find that the average abatement cost in the absence of transaction cost is 1.3% lower than the baseline case.

To see this, I calculate the firm-and-year level outcomes in the baseline and no-banking case in Table 1.6. The table shows that the trading volume (i.e., the absolute value of the net purchase, $|b_{it}|$) is higher by 29% in the absence of permit banking than in the baseline, implying that permit trading is more active without the banking. When banking is not allowed, firms have a higher incentive to trade. Because firms cannot rely on their banked permits for compliance, they have to buy permits from other firms. In the case of sellers, they have to discard emissions permits unless they participate in permit trading to sell. Overall, more active trading of emissions permits leads to efficient allocation of emissions permits. The difference in trading patterns is a potential reason why the average abatement costs are not higher than the baseline case, even though permit banking is not allowed.

1.6.2. Experiment 2: The Potential Gains from Trade

I now examine the implications of transaction costs. As I discussed in section 1.3.8.1, transaction costs are a source of inefficiency in a cap-and-trade program. Estimates of my structural model suggest that although the transaction costs in the form of participation cost are small, the variable transaction costs are substantial.

Table 1.6. Simulation outcomes at the Firm-and-year Level

	Baseline		
	Mean	Std. Dev	Median
Emissions e_{it}	85,116.1	111,237.8	48,221.6
Net purchase b_{it}	5,410.8	9,326.0	2,131.3
Trading volume $\|b_{it}\|$	7,229.0	7,996.9	3,657.0
	No Banking		
	Mean	Std. Dev	Median
Emissions e_{it}	77,753.9	90,146.3	50,683.4
Net purchase b_{it}	5,000.6	12,353.0	0.0
Trading volume $\|b_{it}\|$	9,300.0	9,539.2	5,215.0
Left-over permits	16,698.7	57,228.5	0.0

In this simulation, I shut down both participation and variable transaction costs (while allowing permit banking) and solve a market equilibrium. This simulation quantifies the potential gains from trade in the absence of transaction costs. This simulation outcome is also interpreted as an outcome when the regulator introduces a centralized exchange for emissions permits and runs an auction to allocate all the emissions permits.

Table 1.7 shows the counterfactual result along with the baseline result. The upper panel shows that the level of permit banking at the terminal period is lower in the absence of transaction costs than in the baseline. This pattern can be found in other years, as shown in figure 1.12. Under the presence of transaction costs, firms are less active in permit trading and thus they rather save emissions permits.

Regarding the abatement costs for firms, the table shows that the costs would be lower by 5.8 billion USD in total, or by 37%. Although this partially reflects the higher emissions level in the absence of transaction costs, the average cost of abatement is also lower by 119

USD per SO_2 tons, or by 31.6%. This finding indicates that the "potential" gains from trade, which could be achieved in the absence of transaction costs are significant.

Where does this additional cost saving come from? Figure 1.11 plots the distributions of emissions rates in the case without transaction costs and the baseline case. The distribution is more dispersed in the absence of transaction costs than in the baseline. Shutting transaction costs down makes firms trade more actively, so that firms are more flexible in their compliance. Some firms that are costly to reduce emissions by themselves are more likely to purchase emissions permits, whereas other firms invest more because their revenues from selling permits become higher once transaction costs are removed.

Another source of cost saving is due to the lower level of permit banking. Figure 1.12 shows that the level of banking is higher in the baseline than in the absence of transaction costs. When the transaction costs exist, firms prefer to save emissions permits, instead of selling them in the market. Once transaction costs are shut down, firms have a higher incentive to trade permits with other firms, improving the allocative efficiency of emissions permits.

Health and environmental damages increase by 4.5 billion USD in the absence of transaction costs. This increase reflects both the increase in the aggregate level of SO_2 emissions, as well as the increase in the average health damages. In particular, the average health damages increase by 0.7%, indicating that more active trading of emissions permits leads to more emissions in the region where the health damage is higher. In sum, the total costs including firms' costs of abatement and health damages decreased by 1.3 billion USD (1.8%) in the absence of transaction costs.

Table 1.7. The Potential Gains from Trade

		Baseline	No Transaction Costs
Emissions (in 1 million ton)		45.54	48.99
Dumped permits		0.00	0.00
Banking at the terminal period		5.21	1.76
Firm costs	Investment Costs	15,179	9,740
(in million USD)	Loss of Electricity Profit	380	25
	Total costs	15,559	9,765
	change from baseline		-5,794
Health and environmental damages	Total	54,787	59,321
(in million USD)	change from baseline		4,534
Total Costs	Total costs	70,346	69,086
(in million USD)	change from baseline		-1,259
Firm costs / abatement (in USD)		376	257
Health and environmental damage / emissions (in USD)		1,203	1,211

Note: The numbers are the total from 1995 to 2003. The unit of emissions and banking at the terminal period is 1 million SO_2 tons. The unit of costs are 1 million USD in 2000.

1.7. Conclusion and Further Directions

I study the welfare effects of a cap-and-trade program on air pollutants in the context of the SO_2 emissions regulation in the US electricity industry. I construct a dynamic equilibrium model of a cap-and-trade program in which firms makes decisions on abatement investment, permit trading, and permit banking. By applying the model to the data from the US Acid Rain Program, I find that variable transaction costs associated with permit trading are substantial.

I use the estimated model to quantify the effects of a cap-and-trade program in comparison to the uniform standard on emissions rate as a counterfactual command-and-control type policy. I find that the cap-and-trade program decreased the aggregate costs of reducing

Figure 1.11. Distribution of Emissions Rate in the Absence of Transaction Costs

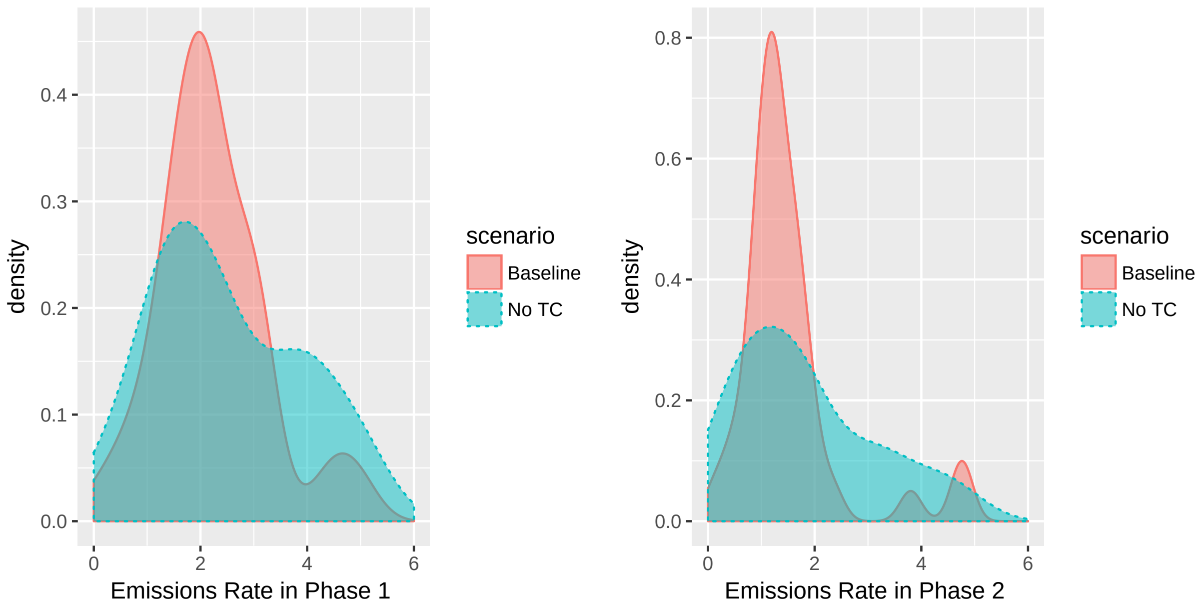

Note: Emissions rate is measured by lbs per MMBtu.

Figure 1.12. Permit Banking in the Absence of Transaction Costs

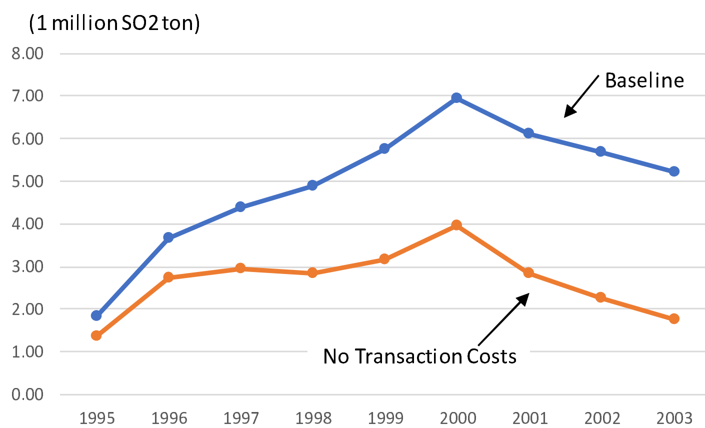

emissions by 340 million USD per year, or 16.6%. Although the damages from SO_2 emissions increased as a result of permit trading, the cost saving is sufficiently high to offset the negative effects on health.

I also simulate the counterfactual outcome in the absence of transaction costs and find that the aggregate costs could be saved further by 643 million USD per year, or by 26%. This additional cost saving is achieved by a more efficient allocation of investment. My findings indicate that the full potential of a cap-and-trade program has not been realized in my sample period.

Several extensions and applications remain for future work. First, it would be interesting to study the implications of policies regarding permit prices in a cap-and-trade program. Recently, the volatile and low permit prices observed in cap-and-trade programs concern policymakers, which leads to the proposal of various measures to stabilize permit prices, including a price floor/ceiling and the Market Stability Reserve. By extending my empirical framework, I could evaluate the effectiveness of these proposed policies and its welfare consequences.

In addition, my framework can be applied beyond a cap-and-trade program on air pollutants. The governments now use a market-based policy in various settings, including credit trading in the CAFE regulation and Renewable Energy Certificates in Renewable Portfolio Standard. Under these policies, firms face a similar problem to the one I study in this paper: a firm can either trade these credits, or invest in technologies (i.e., improving fuel efficiency in the CAFE credit trading, and building renewable generators in the RPS program). My empirical framework can be used to study the effectiveness of these market-based policies and the implications of alternative regulatory designs. I leave these topics for future work.

CHAPTER 2

Voter Turnout and Preference Aggregation

2.1. Introduction

Democracies rely on elections to aggregate the preferences of their citizens. Elections, however, aggregate the preferences of only those that participate. The importance of participation for preference aggregation is documented by studies of suffrage expansion in various contexts, such as the abolition of apartheid in South Africa (Kroth et al., 2013), the passage of the Voting Rights Act of 1965 (Husted and Kenny, 1997; Cascio and Washington, 2013), and the passage of women's suffrage laws (Miller, 2008). Less dramatic measures that have reduced the voting costs of certain groups of voters have also been found to affect policy in important ways (Fujiwara, 2015).

While most democracies now enjoy universal suffrage, participation in elections is far from perfect, given the voluntary nature of voting. To the extent that the preferences of those that turn out are systematically different from those that do not, election outcomes may poorly aggregate the preferences of all citizens. Thus, how well elections aggregate the preferences of citizens and whose preferences are underrepresented are open questions, even in mature democracies.

The issues of preference aggregation and underrepresentation are also relevant from a policy perspective. The concern that the preferences of certain groups of voters are underrepresented has led some to argue for compulsory voting (see, e.g., Lijphart, 1997). More moderate policy proposals, such as introducing Internet voting, relaxing registration requirements, and making election day a holiday, are motivated by similar concerns. Understanding how voter turnout affects preference aggregation can provide a basis for more informed discussions of these policy proposals.

Preference aggregation is also central to the policy debate on partisan redistricting, or gerrymandering. Gerrymandering can be considered as an intentional attempt by one party to aggregate preferences disproportionately in its favor through redistricting. Much of the recent discussion on gerrymandering focuses on how well actual votes are translated into seat shares, but ignores how preferences map to seat shares. Given that turnout is endogenous, it is possible for sophisticated planners to design redistricting plans that map vote shares to seat shares well, but map preferences to seat shares poorly.[1] Studying how underlying preferences, rather than votes, are aggregated into election outcomes provides a coherent alternative even when turnout is endogenous.

In this study, we explore the extent to which preferences are aggregated in elections, which hinges on how the preferences, voting costs, and perceptions of voting efficacy are correlated. We show identification of the joint distribution of these three terms, and estimate it using county-level voting data from the 2004 U.S. presidential election. We find that young, low-income, less-educated, and minority voters have a high cost of voting and that all of these groups tend to prefer the Democrats, except for the less-educated. We then simulate the counterfactual election outcome when all voters vote. The difference between the simulated and actual outcomes allows us to quantify the degree to which preferences are aggregated. In our counterfactual, the two-party vote share of the Democrats increases by about 3.7%, and the Democrats win the plurality of the electoral votes. In our second counterfactual,

[1]In Table 2.5, we illustrate this point using a numerical example.

we compute the election outcome when we eliminate endogeneity in turnout that is state-specific. This allows us to gauge the sensitivity of the efficiency gap, an influential measure of gerrymandering, to endogeneity of turnout.

The key challenge in studying the effect of turnout on preference aggregation is to identify the correlation between preferences and voting costs in the population. In particular, we need to identify how voter characteristics such as race and income simultaneously determine preferences and costs. However, this is not a straightforward task because high preference intensity and high voting cost may predict a similar voting pattern as low preference intensity and low voting cost.

To illustrate, consider a plurality rule election in which voters have private values and choose to vote for candidate A or candidate B or not to turn out. Applying a discrete choice framework to the voter's decision, let $u_A(\mathbf{x})$ and $u_B(\mathbf{x})$ denote the utility of voting for candidates A and B, respectively, and $c(\mathbf{x})$ denote the cost of voting (relative to not voting), where $\mathbf{x}$ is a vector of voter characteristics. Then, the voter's mean utilities are as follows:

$$\begin{aligned} V_A(\mathbf{x}) &= u_A(\mathbf{x}) - c(\mathbf{x}), \\ V_B(\mathbf{x}) &= u_B(\mathbf{x}) - c(\mathbf{x}), \text{ and} \\ V_0(\mathbf{x}) &= 0, \end{aligned}$$

where V_0 represents the mean utility of not turning out. While one can identify $V_A(\mathbf{x}) = u_A(\mathbf{x}) - c(\mathbf{x})$ and $V_B(\mathbf{x}) = u_B(\mathbf{x}) - c(\mathbf{x})$ by using vote share and turnout data (see Berry, 1994; Hotz and Miller, 1993), $u_A(\cdot)$, $u_B(\cdot)$, and $c(\cdot)$ are not separately identified without further restrictions. This is because making a voter care more about the election outcome (say, by adding an arbitrary function, $g(\mathbf{x})$, to both $u_A(\mathbf{x})$ and $u_B(\mathbf{x})$) is observationally equivalent to lowering the voting cost (by subtracting $g(\mathbf{x})$ from $c(\mathbf{x})$). Even if there are exogenous cost shifters $\mathbf{z}$ (e.g., rainfall), they do not help separately identify $u_A(\cdot)$, $u_B(\cdot)$, and $c(\cdot)$.[2] Thus, most existing studies impose ad-hoc exclusion restrictions on the way that $\mathbf{x}$ enters $u_A(\cdot)$, $u_B(\cdot)$, and $c(\cdot)$, assuming that $\mathbf{x}$ is excluded from either $u_k(\cdot)$ or $c(\cdot)$. Imposing such exclusion restrictions assumes away the correlation structure among these terms and precludes the possibility that the preferences of those with high voting costs are different from those with low voting costs. Note that this identification challenge exists regardless of whether the data are available at the individual level or at the aggregate level.

In this paper, we uncover the correlation structure between preferences and costs in a setting in which $\mathbf{x}$ is allowed to enter both $u_k(\cdot)$ and $c(\cdot)$. Our identification is based on the simple observation that, unlike consumer choice problems in which choosing not to buy results in the outcome of not obtaining the good, choosing not to turn out still results in either A or B winning the election. This observation implies that the voter's choice is determined by the utility difference between the two election outcomes rather than by the levels of utility associated with each outcome.[3] Barkume (1976) first used this observation to separately identify $u_k(\cdot)$ and $c(\cdot)$ in the context of property tax referenda for school districts.

To see how this observation leads to the identification of $u_k(\cdot)$ and $c(\cdot)$, consider the calculus of voting models of Downs (1957) and Riker and Ordeshook (1968). In these models, the utility of voting for candidate k can be expressed as $u_k = pb_k$, where p is the voter's beliefs that she is pivotal; b_A is the utility difference between having candidate A in office

[2]Suppose that the cost function is separated into two parts as $c = c_{\mathbf{x}}(\mathbf{x}) + c_{\mathbf{z}}(\mathbf{z})$, where $\mathbf{z}$ is a vector of cost shifters that is excluded from $u_A(\cdot)$ and $u_B(\cdot)$. Then, $u_A(\cdot) - c_{\mathbf{x}}(\cdot)$, $u_B(\cdot) - c_{\mathbf{x}}(\cdot)$ and $c_{\mathbf{z}}(\cdot)$ are all separately identified. However, $u_A(\cdot)$, $u_B(\cdot)$ and $c_{\mathbf{x}}(\cdot)$ are not separately identified. See the subsection titled "Exogenous Cost Shifters" towards the end of Section 4 for details.

[3]This implication holds as long as voters care about the ultimate outcome of the election. However, it may not hold for models in which voters gain utility from the act of voting for a candidate, such as models of expressive voting.

and candidate B in office; and b_B is defined similarly.[4] Hence, we have $b_A = -b_B$. The mean utilities can now be expressed as

$$\begin{aligned} V_A(\mathbf{x}) &= \quad pb_A(\mathbf{x}) - c(\mathbf{x}), \\ V_B(\mathbf{x}) &= -pb_A(\mathbf{x}) - c(\mathbf{x}), \text{ and} \\ V_0(\mathbf{x}) &= 0. \end{aligned}$$

The property $b_A = -b_B$ allows us to separately identify preferences and costs. By adding the first two expressions above, we have $V_A(\mathbf{x})+ V_B(\mathbf{x}) = -2c(\mathbf{x})$ because $pb_A(\mathbf{x})$ cancels out. Given that $V_A(\mathbf{x})$ and $V_B(\mathbf{x})$ are both identified from the vote share and turnout data, $c(\cdot)$ is identified. Similarly, we can identify $pb_A(\cdot)$ because we have $V_A(\mathbf{x})- V_B(\mathbf{x}) = 2pb_A(\mathbf{x})$, and $V_A(\mathbf{x})- V_B(\mathbf{x})$ is identified. As we discuss below, $V_A(x)- V_B(x)$ is identified primarily by the vote share margin and $V_A(x)+ V_B(x)$ is identified primarily by voter turnout. Intuitively, changes in preference intensity in this model necessarily changes the vote share margin, while changes in costs have a similar effect on the vote shares of the two parties. This implies that changes in b_A cannot be undone by changes in c.[5]

In this paper, we retain the basic structure of the calculus of voting model but do not place additional restrictions on p, such as rational expectations, in which p equals the actual pivot probability. In our model, we interpret p more broadly as the voter's perception of voting efficacy, which is allowed to differ across individuals and to be correlated with the true pivot probability in a general manner. In particular, we let p be a function of individual

[4]More precisely, the utility of voting for candidate k *relative to not turning out* can be expressed as $u_k = pb_k$, by normalizing the utility of not turning out to be zero. See Appendix A for details.

[5]See Merlo and de Paula (2016) for identification of voter preferences in a spatial voting model with full turnout.

characteristics and the state in which the voter lives, as $p = p_s \times \tilde{p}(\mathbf{x})$, where p_s is a state-specific coefficient and $\tilde{p}(\cdot)$ is a function of voter characteristics, $\mathbf{x}$. By letting p depend on each state, we can take into account the nature of the electoral college system.[6] We show that the ratios of the state-specific components of efficacy, $p_s/p_{s'}$ $(\forall s, s')$, are identified directly from the data. Moreover, we show that $\tilde{p}(\cdot)$, $b_A(\cdot)$ and $c(\cdot)$ are identified up to a scalar normalization.[7] Our identification discussion does not depend on equilibrium restrictions on p, such as rational expectations. Therefore, our identification and estimation results are agnostic about how voters formulate p.

Given the debate over how to model voter turnout, we briefly review the literature on turnout to situate our model.[8] The model that we estimate in this paper is based on the decision theoretic model of voter turnout introduced by Downs (1957) and Riker and Ordeshook (1968). In their models, a voter turns out and votes for the preferred candidate if $pb - c + d > 0$, where p is the voter's beliefs over the pivot probability; b is the utility difference from having one's preferred candidate in office relative to the other; c is the physical and psychological costs of voting; and d is the benefit from fulfilling one's civic duty to vote. While the original studies do not endogenize any of these terms, the decision theoretic model has provided a basic conceptual framework for much of the subsequent work on voting and turnout.

[6]Under the electoral college system, perceptions of voting efficacy may differ significantly across states. For example, electoral outcomes in battleground states such as Ohio were predicted to be much closer than outcomes in party strongholds such as Texas. Hence, we allow for the possibility that p is higher for voters in Ohio than for voters in Texas.

[7]More precisely, we can identify $p(\cdot)b_A(\cdot)$ state by state given that we have many counties within each state. Assuming that $\tilde{p}(\cdot)$ and $b_A(\cdot)$ are common across states, we can identify $p_s/p_{s'}$. We also show that $\tilde{p}(\cdot)$ and $b_A(\cdot)$ are separately identified up to a scalar multiple in our full specification with county-level shocks to preferences and costs.

[8]For a survey of the literature, see, e.g., Dhillon and Peralta (2002), Feddersen (2004), and Merlo (2006).

Studies subsequent to Riker and Ordeshook (1968) endogenize or micro-found each of the terms in the calculus of voting model in various ways. Ledyard (1984) and Palfrey and Rosenthal (1983, 1985) introduce the pivotal voter model, in which the pivot probability p is endogenized in a rational expectations equilibrium. They show that there exists an equilibrium with positive turnout in which voters have consistent beliefs about the pivot probability. Coate et al. (2008), however, point out that the rational expectations pivotal voter model has difficulties matching the data on either the level of turnout or the winning margin.[9] Moreover, using laboratory experiments, Duffy and Tavits (2008) find that voters' subjective pivot probabilities are much higher than the actual pivot probability, which is at odds with the rational expectations assumption.

More recently, there have been attempts at endogenizing p in ways other than rational expectations. For example, Minozzi (2013) proposes a model based on cognitive dissonance in the spirit of Akerlof and Dickens (1982) and Brunnermeier and Parker (2005). In his model, voters jointly choose p and whether or not to turn out in order to maximize subjective expected utility. Kanazawa (1998) introduces a model of reinforcement learning in which boundedly rational voters, who cannot compute the equilibrium pivot probabilities, form expectations about p from the correlation between their own past voting behavior and past election outcomes (see, also, Bendor et al., 2003; Esponda and Pouzo, 2016, for similar approaches). While these models are based on the basic calculus of voting model, the p term in them no longer carries the interpretation of the actual pivot probability.

[9]Note, however, that with aggregate uncertainty, Myatt (2012) shows that the level of turnout can still be high with rational expectations. Levine and Palfrey (2007) also show that combining the quantal response equilibrium with the pivotal voter model can generate high turnout and finds that the results of laboratory experiments are consistent with the model prediction.

Another strand of the literature endogenizes the c and d terms. Harsanyi (1980) and Feddersen and Sandroni (2006) endogenize the d term by proposing a rule-utilitarian model in which voters receive a warm-glow payoff from voting ethically. Based on their approach, Coate and Conlin (2004) estimate a group-utilitarian model of turnout. Shachar and Nalebuff (1999) also endogenize the d term by considering a follow-the-leader model in which elites persuade voters to turn out. In a paper studying split-ticket voting and selective abstention in multiple elections, Degan and Merlo (2011) consider a model that endogenizes c to reflect the voter's psychological cost of making mistakes.

In our paper, we bring the calculus of voting model to the data without taking a particular stance on how the p, b, c, or d terms are endogenized. Specifically, our identification and estimation do not use the restriction that p is equal to the actual pivot probability, as in the rational expectations model. The p term that we recover can be broadly interpreted as the voter's perception of voting efficacy. We purposely aim to be agnostic about the different ways of modeling voter turnout so that our estimates of voter preferences and costs are robust to the specific way in which the p, b, c, or d terms are endogenized. Instead of imposing equilibrium restrictions of a particular model a priori, we let the data directly identify the p, b, and $c - d$ terms.

Relatedly, our study does not impose a priori restrictions on how the covariates enter the p, b, or $c-d$ terms, allowing, instead, the same set of covariates to affect all three terms. This is important because the way in which covariates enter the p, b, and $c - d$ terms determines the correlation structure among them, which, in turn, determines how well preferences are aggregated. In most existing studies, the sets of covariates that enter the p, b and $c-d$ terms are disjoint, precluding the possibility that preferences and costs are correlated. For example, Coate and Conlin (2004) and Coate et al. (2008) include demographic characteristics only in the b term,[10] while Shachar and Nalebuff (1999) include them in the $c - d$ term. In contrast, we let each demographic characteristic enter all three terms, allowing us to study the effects of turnout on preference aggregation.[11]

We use county-level data on voting outcomes from the 2004 U.S. presidential election to estimate the model.[12] A benefit of using actual voting data over survey data is that we can avoid serious misreporting issues often associated with survey data, such as the overreporting of turnout and reporting bias in vote choice (see, e.g., Atkeson, 1999; DellaVigna et al., 2015).[13] Our data on turnout and vote share incorporate the number of non-citizens and felons to account for the difference between the voting-eligible population and the voting-age population (McDonald and Popkin, 2001). We construct the joint distribution of demographic characteristics within each county from the 5% Public Use Microdata Sample of the Census.

We find that young, less-educated, low-income, and religious voters have high voting costs, as do African Americans, Hispanics, and other minorities. Moreover, young voters have low perception of voting efficacy, which further depresses turnout among this group. Overall, young, less-educated, and low-income voters are particularly underrepresented. In terms of preferences, minority, young, highly educated, low-income, and non-religious voters are more likely to prefer Democrats.

[10]To be more precise, Coate and Conlin (2004) and Coate et al. (2008) use demographic characteristics as covariates for the fraction of the population supporting one side.

[11]One possible exception is Degan and Merlo (2011). They consider a model based on the theories of regret in which the cost term is endogenized in a way that captures voters' preferences over candidates. They include the same set of covariates in the c and d terms. In one of their counterfactual analyses, they consider the effect of increasing voter turnout, focusing on split-ticket voting and selective abstention across presidential and congressional elections.

[12]Although we use aggregate data, we account for the issue of ecological fallacy by computing the behavior of individual voters and aggregating them at the county level.

[13]There is a set of survey-based studies that investigate the differences in preferences between voters and non-voters (see, e.g., Citrin et al., 2003; Brunell and DiNardo, 2004; Martinez and Gill, 2005; Leighley and Nagler, 2013).

Our results show that, overall, there is a positive correlation between voting cost and preference for Democrats that can be accounted for through observable characteristics. Except for two voter characteristics–years of schooling and being religious–we find that demographic characteristics that are associated with a higher cost of voting are also associated with preferring Democrats. We also find that unobservable cost shocks are positively correlated with unobservable preference shocks for Democrats. These correlations result in fewer Democratic votes relative to the preferences of the underlying population. Our estimate of turnout is significantly lower among the electorate who prefer Democrats to Republicans, at 55.7%, compared with turnout among those who prefer Republicans to Democrats, at 64.5%.[14] Moreover, we find that voters who have a strong preference for one of the parties are more likely to turn out, suggesting that preference *intensity* affects preference aggregation (see Campbell, 1999; Casella, 2005; Lalley and Weyl, 2015).

Regarding our results on the perception of voting efficacy, we find substantial across-state variation in our estimates of p_s, the state-specific coefficient in p. Furthermore, the estimates are correlated with the ex-post closeness of the election: Battleground states such as Ohio and Wisconsin tend to have high estimates of p_s, while party strongholds such as New Jersey and California have low estimates, which is consistent with the comparative statics of the pivotal voter model with rational expectations. However, the magnitude of the estimated ratio of p_s is, at most, three for any pair of states. This is in contrast to a much larger variation in the ratio implied by the pivotal voter model.[15] Our results are more consistent

[14]See DeNardo (1980) and Tucker and DeNardo (1986) for studies that report negative correlation between turnout and the Democratic vote share using aggregate data. For more recent work, see Hansford and Gomez (2010) who use rainfall as an instrument for turnout.

[15]The pivotal voter model with rational expectations predicts high variation in the ratio of pivot probabilities across states, given the winner-take-all nature of the electoral college system. Voters in only a handful of swing states have a reasonable probability of being pivotal (see, e.g., Shachar and Nalebuff, 1999).

with models of turnout in which voters' perception of efficacy is only weakly correlated with the actual pivot probabilities.

In our first counterfactual experiment, we simulate the voting outcome when all voters vote. We find that the vote share of the Democrats increases in all states. Overall, the increase in the Democrats' two-party vote share is about 3.7%. We also find that the increase in the Democratic vote share would overturn the election results in nine states, including key states such as Florida and Ohio, resulting in the Democrats winning a plurality of the electoral votes.

In our second counterfactual experiment, we compare the actual election outcome with the counterfactual outcome when we equalize the state-specific component of efficacy across states (set $p_s = p_{s'}$). Equalizing p_s across states can be interpreted as eliminating endogeneity in turnout that is driven by how state boundaries are drawn. This counterfactual is motivated by the recent development in measuring gerrymandering, in particular, the use of a metric called the "efficiency gap" (Stephanopoulos and McGhee, 2015) for determining the legality of districting plans. By eliminating state-specific endogeneity in turnout, we gauge the robustness of the efficiency gap to endogenous turnout.[16]

We find that equalizing p_s across states while keeping turnout at the actual levels changes the efficiency gap by 2.2 percentage points. We also find that the efficiency gap changes substantially when we vary the value of p_s that we use. For example, a change in the level of p_s corresponding to a change in turnout from 50% to 80% increases the efficiency gap by 14.1 percentage points. To put these numbers in perspective, 2.2 percentage points is comparable to about half of the increase in the efficiency gap for the Republican party over

[16]The literature on districting focuses on the mapping from vote share to seat share, and little attention is paid to the fact that turnout is endogenous. See, e.g., Coate and Knight (2007), and Friedman and Holden (2008). For a survey, see Nagle (2015).

the past 30 years in the U.S. state legislative elections. A change in the efficiency gap of 14.1 percentage points is larger than the proposed threshold value of 8 percentage points above which Stephanopoulos and McGhee (2015) argue that districting plans should be deemed presumptively unlawful. These results suggest that the efficiency gap is quite sensitive to endogenizing turnout. One natural alternative to comparing actual votes and seat shares, which is what the efficiency gap does as well as other measures such as partisan symmetry (e.g., Grofman and King, 2007), is to compare underlying voter preferences and seat shares. If we think about elections as a way to aggregate preferences into outcomes, evaluating the electoral system in terms of its ability to aggregate preferences seems most coherent.

2.2. Model

Anticipating the empirical application of the paper, we tailor our model to the U.S. presidential election. Let $s \in \{1, ..., S\}$ denote a U.S. state and $m \in \{1, ..., M_s\}$ denote a county in state s.

Preference of Voters. We consider a model of voting with two candidates, D and R. Each voter chooses to vote for one of the two candidates or not to vote. We let b_{nk} denote voter n's utility from having candidate $k \in \{D, R\}$ in office, p_n ($p_n > 0$) denote her perception of voting efficacy, and c_n denote her cost of voting. Given that there are only two possible outcomes (either D wins or R wins the election), the utility of voting for candidate k, U_{nk} depends only on $b_{nD} - b_{nR}$ rather than on b_{nD} and b_{nR} individually:

$$U_{nD} = p_n(b_{nD} - b_{nR}) - c_n, \tag{2.2.1}$$

$$U_{nR} = p_n(b_{nR} - b_{nD}) - c_n, \tag{2.2.2}$$

$$U_{n0} = 0,$$

where U_{n0} is the utility of not turning out, which we normalize to zero (see Appendix A for a derivation).[17] When p_n is the actual pivot probability, the behavior of the voters under our model is the same as the equilibrium play of the voters under the pivotal voter model of Palfrey and Rosenthal (1983, 1985). However, we interpret p_n broadly as the voter's subjective perception of voting efficacy, as we discuss below. The cost of voting, c_n, includes both physical and psychological costs, as well as possible benefits of fulfilling one's civic duty. Hence, c_n can be either positive or negative. When c_n is negative, the voter turns out regardless of the value of p_n and $b_{nD} - b_{nR}$.

We let the preferences of voter n in county m of state s depend on her demographic characteristics, $\mathbf{x}_n$, as follows:

$$b_{nk} = b_k(\mathbf{x}_n) + \lambda_{sk} + \xi_{mk} + \varepsilon_{nk}, \text{ for } k \in \{D, R\},$$

where λ_{sk} is a state-specific preference intercept that captures state-level heterogeneity in voter preferences. ξ_{mk} and ε_{nk} are unobserved random preference shocks at the county level and at the individual level, respectively. ξ_{mk} captures the unobserved factors that affect preferences at the county level, such as the benefits that the voters in county m receive from policies supported by candidate k. Then, the expression for the utility difference is as follows:

$$b_{nR} - b_{nD} = b(\mathbf{x}_n) + \lambda_s + \xi_m + \varepsilon_n,$$

where $b(\mathbf{x}_n) \equiv b_R(\mathbf{x}_n) - b_D(\mathbf{x}_n)$, $\lambda_s \equiv \lambda_{sR} - \lambda_{sD}$, $\xi_m \equiv \xi_{mR} - \xi_{mD}$, and $\varepsilon_n \equiv \varepsilon_{nR} - \varepsilon_{nD}$. We assume that ε_n follows the standard normal distribution.

[17]Note that expressions (2.2.1) and (2.2.2) take the familiar form of $pb - c$.

We also let voting cost c_n be a function of voter n's characteristics as

$$c_n = c_s(\mathbf{x}_n) + \eta_m,$$

where $c_s(\mathbf{x}_n)$ is the cost function and η_m is a county-level shock on the cost of voting. The cost function, $c_s(\cdot)$, is allowed to depend on s. Given that previous studies (e.g., Smith, 2001) find that neither the presence nor the closeness of congressional elections affect turnout in presidential elections, we do not incorporate other elections as a cost shifter in our model. We assume that ξ_m and η_m are both independent of $\mathbf{x}_n$, but we allow ξ_m and η_m to be correlated with each other.

We let the voting efficacy term, p_n, depend on both the demographic characteristics of voter n and the state in which she votes as follows:

$$p_n = p_s(\mathbf{x}_n) = p_s \times \tilde{p}(\mathbf{x}_n),$$

where p_s is a state specific coefficient that we estimate. It is important to let p_n depend on the state in which the voter votes because of the winner-take-all nature of the electoral votes in each state.[18] For example, in the 2004 presidential election, a vote in key states such as Ohio was predicted to matter considerably more than a vote elsewhere. Our specification also allows for the possibility that p_n depends on voters' characteristics, $\mathbf{x}_n$. Previous work has shown that voters' social and economic status affects her general sense of political efficacy (see, e.g., Karp and Banducci, 2008).

[18]In U.S. presidential elections, the winner is determined by the Electoral College. Each U.S. state is allocated a number of electoral votes, roughly in proportion to the state's population. The electoral votes of each state are awarded on a winner-takes-all basis in all states, except for Maine and Nebraska. The Presidential candidate who wins the plurality of electoral votes becomes the winner of the election.

Note that the behavior of the voters under our model is the same as the equilibrium play of the voters under the pivotal voter model if we set p_s equal to the actual pivot probability in state s and set $\tilde{p}(\mathbf{x}_n)$ equal to 1. In this sense, our specification nests the pivotal voter model as a special case. However, instead of imposing the pivotal voter model (and, hence, placing equilibrium restrictions on p_n), we estimate p_s and $\tilde{p}(\cdot)$ directly from the data. This approach allows us to interpret p_n consistently with models of turnout that endogenize p_n in various ways.

Substituting the expressions for $b_{nR} - b_{nD}$, c_n, and p_n into equations (2.2.1) and (2.2.2), the utility from choosing each of the alternatives can be expressed as follows:

$$\begin{aligned} U_{nD}(\mathbf{x}_n) &= p_s(\mathbf{x}_n)\left[-b_s(\mathbf{x}_n) - \xi_m - \varepsilon_n\right] - c_s(\mathbf{x}_n) - \eta_m, \\ U_{nR}(\mathbf{x}_n) &= p_s(\mathbf{x}_n)\left[b_s(\mathbf{x}_n) + \xi_m + \varepsilon_n\right] - c_s(\mathbf{x}_n) - \eta_m, \\ U_{n0}(\mathbf{x}_n) &= 0, \end{aligned}$$

where $b_s(\mathbf{x}_n)$ denotes $b(\mathbf{x}_n) + \lambda_s$.

A Voter's Decision. Voter n's problem is to choose the alternative $k \in \{D, R, 0\}$ that provides her with the highest utility:

$$k = \underset{\kappa \in \{D,R,0\}}{\arg\max}\ U_{n\kappa}(\mathbf{x}_n). \tag{2.2.3}$$

We can write the probability that voter n votes for candidate R as

$$
\begin{aligned}
&\Pr\left(R = \arg\max_{\kappa\in\{D,R,0\}} U_{n\kappa}\right) \\
&\quad = \Pr\left(U_{nR} > U_{nD} \text{ and } U_{nR} > 0\right) \\
&\quad = \Pr\left(\varepsilon_n > -b_s(\mathbf{x}_n) - \xi_m \text{ and } \varepsilon_n > -b_s(\mathbf{x}_n) - \xi_m + \frac{c_s(\mathbf{x}_n) + \eta_m}{p_s(\mathbf{x}_n)}\right) \\
&\quad = 1 - \Phi\left(\max\left\{-b_s(\mathbf{x}_n) - \xi_m, -b_s(\mathbf{x}_n) - \xi_m + \frac{c_s(\mathbf{x}_n) + \eta_m}{p_s(\mathbf{x}_n)}\right\}\right),
\end{aligned}
$$

where Φ is the CDF of the standard normal. We can derive a similar expression for candidate D.

Figure 1 depicts the behavior of a voter as a function of ε_n. There are two cases to consider: one in which the cost of voting is positive (Case 1) and the other in which the cost of voting is negative (Case 2). In Case 1, a voter with a strong preference for one of the candidates (which corresponds to a large positive realization or a large negative realization of ε_n) votes for her preferred candidate, while a voter who is relatively indifferent between the two candidates does not turn out. That is, a voter with high preference intensity relative to cost turns out, while a voter with low preference intensity does not. In Case 2, a voter always votes, regardless of her preference intensity, as the cost of voting is negative.

Vote Share and Voter Turnout. We can express the vote share for candidate k in county m, $v_{k,m}$, and the fraction of voters who do not turn out, $v_{0,m}$, as follows:

$$v_{R,m} \equiv \int 1 - \Phi\left(\max\left\{-b_s(\mathbf{x}_n) - \xi_m, -b_s(\mathbf{x}_n) - \xi_m + \frac{c_s(\mathbf{x}_n) + \eta_m}{p_s(\mathbf{x}_n)}\right\}\right) dF_{\mathbf{x},m}(\mathbf{x}_n), \tag{2.2.4}$$

$$v_{D,m} \equiv \int \Phi\left(\min\left\{-b_s(\mathbf{x}_n) - \xi_m, -b_s(\mathbf{x}_n) - \xi_m - \frac{c_s(\mathbf{x}_n) + \eta_m}{p_s(\mathbf{x}_n)}\right\}\right) dF_{\mathbf{x},m}(\mathbf{x}_n), \tag{2.2.5}$$

$$v_{0,m} \equiv 1 - v_{D,m} - v_{R,m} \tag{2.2.6}$$

Figure 2.1. Voter's Decision as a Function of ε_n.

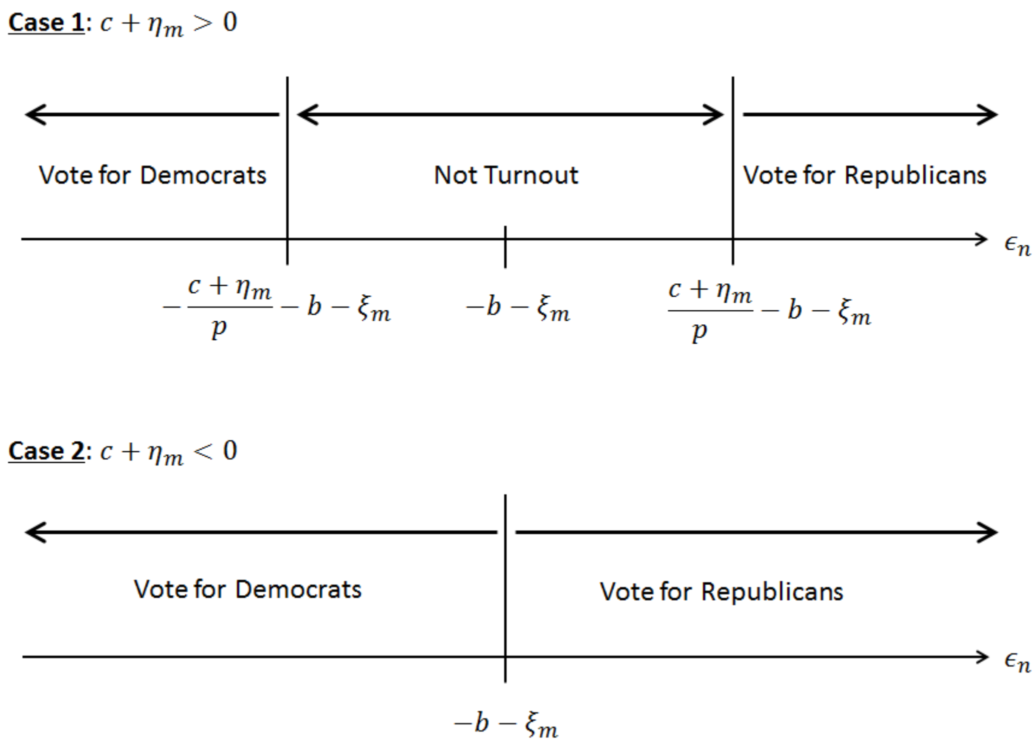

Note: The top panel corresponds to the case in which a voter has positive costs of voting. The bottom panel corresponds to the case in which a voter has negative costs of voting.

where $F_{\mathbf{x},m}$ denotes the distribution of $\mathbf{x}$ in county m. Denoting the number of eligible voters in county m by N_m and the number of counties in state s as M_s, the vote share for candidate k in state s can be expressed as $\sum_{m=1}^{M_s} N_m v_{k,m} \Big/ \sum_{m=1}^{M_s} N_m$. The candidate with the highest vote share in state s is allocated all of the electors of that state.[19] The candidate who wins the plurality of the electors becomes the overall winner of the presidential election.

Advertising and Campaign Visits. An important feature of presidential elections not explicitly modeled thus far is the campaign activities of candidates. Candidates target key states with advertisements and campaign visits during the election. These campaign activities

[19]Maine and Nebraska use a different allocation method. Hence, we drop these two states from our sample.

are endogenous and depend on the expected closeness of the race in each state (see, e.g., Strömberg, 2008; Gordon and Hartmann, 2013).

While we do not have a specific model of political campaigns, the model accounts for the effect of campaigns on voters through the state-specific preference intercept λ_s. Because we treat λ_s as parameters to be estimated, λ_s may be arbitrarily correlated with the characteristics of the state, the closeness of the race in the state, etc. Hence, our estimates of the primitives of the model are consistent even in the presence of campaign activities. We note, however, that the results of our counterfactual experiments take the level of campaigning as given.

Discussion on Voter's Information. Another factor that we do not specifically model is voter's information. One way to explicitly model information is by endogenizing the voters' information acquisition (see, e.g., Matsusaka, 1995; Degan and Merlo, 2011). In these models, the voters decide on the amount of information to acquire about the candidates by paying the cost of information acquisition. Our specification of voter preference and costs can be thought of as the indirect utility of these models to the extent that information acquisition costs are functions of voter demographics. In fact, in our estimation results, we find that income and education are associated with low voting costs, which suggests that the information acquisition cost may comprise an important part of the voting cost (as the opportunity cost for these voters tend to be higher).

Another way to model information is to consider a common value environment in which voters obtain signals about the quality of the candidates (see, e.g., Feddersen and Pesendorfer, 1996, 1997). In these models, voters' utility consists partly of the expected quality of the candidates, which is computed by conditioning on the event that the voter is pivotal. To the extent that the prior beliefs over candidate quality and the signal distribution depend on the voter's demographic characteristics, the common value component is also a function of these characteristics. Hence, our specification of the utility can also be interpreted as a reduced-form of model with a common value component.

Discussion on p. The modeling in our paper is purposely agnostic about how p is endogenized: We do not impose a particular model of p, such as rational expectations (Palfrey and Rosenthal, 1983, 1985), overconfidence (Duffy and Tavits, 2008), or cognitive dissonance (Minozzi, 2013). Similarly, our estimation approach avoids using restrictions specific to a particular way of modeling voter beliefs. The important point for our purpose is that there exists an equilibrium p that corresponds to the data-generating process regardless of the way in which p is endogenized. Our approach is to identify and estimate both the model primitives and the equilibrium p directly from the data with as little structure as possible. This empirical strategy is similar in spirit to that in the estimation of incomplete models, in which some primitives are estimated from the data without fully specifying a model. For example, Haile and Tamer (2003) recovers bidder values without fully specifying a model of the English auction, using only the restriction that the winning bid lies between the valuations of the losers and the winner. Given that their estimation procedure also avoids using restrictions specific to a particular model of the English auction, the estimates are consistent under a variety of models.

In Section **??**, we show that the key primitives of the model are identified without fully specifying how voters form p. We show that the equilibrium p is also identified directly from the data.[20] The strength of our approach is that we impose few restrictions on beliefs,

[20]More precisely, $p_s/p_{s'}$ is identified for any states s and s', and $\tilde{p}(\cdot)$ is identified up to a scalar normalization. See Section 4 for details.

and, thus, our estimates of preferences and costs are consistent under a variety of behavioral assumptions regarding how p is formed. On the other hand, this approach limits the types of counterfactual experiments that we can conduct since we do not specify a particular model regarding p.

2.3. Data

In this section, we describe our data and provide summary statistics. We combine county-level voting data and demographics data.The county-level voting data is obtained from David Leip's Atlas of U.S. Presidential Elections. This dataset is a compilation of election data from official sources such as state boards of elections. The demographics data is obtained from the U.S. Census Bureau. We construct the data on eligible voters for each county by combining the population estimates from the 2004 Annual Estimates of the Resident Population and age and citizenship information from the 2000 Census. We then adjust for the number of felons at the state level using the data from McDonald (2016). Hence, our data account for the difference between the voting age population and voting eligible population (see McDonald and Popkin, 2001).

We construct the joint distribution of voters' demographic characteristics and citizenship at the county level from the 2000 Census by combining the county-level marginal distribution of each demographic variable and the 5% Public Use Microdata Sample (see Appendix B for details). We augment the Census data with county-level information on religion using the Religious Congregations and Membership Study 2000. In particular, we define the variable Religious using adherence to either "Evangelical Denominations" or "Church of Jesus Christ of Latter-day Saints."

Our data consist of a total of 2,909 counties from forty states. Because we need a large number of counties within each state to identify the state-specific parameters, p_s and λ_s, we drop states that have fewer than 15 counties. These states are Alaska, Connecticut, District of Columbia, Delaware, Hawaii, Massachusetts, New Hampshire, Rhode Island, and Vermont. In addition, we drop Maine and Nebraska because these two states do not adopt the winner-takes-all rule to allocate electors. We also drop counties with a population below 1,000 because their vote shares and turnout rates can be extreme due to small population size.[21] Table 2.1 presents the summary statistics of the county-level vote share, turnout, and demographic characteristics. Note that a Hispanic person may be of any race according to the definition used in the Census.

In order to illustrate the degree to which turnout and expected closeness are related, Figure 2.2 plots the relationship between the (ex-post) winning margin and voter turnout at the state level. The two variables are negatively correlated, although the fitted line is relatively flat. The slope of the fitted line implies that a decrease in the (ex-post) winning margin of ten percentage points is associated with an increase in turnout of only about 1.6 percentage points. While the negative correlation may be capturing some of the forces of the rational-expectations pivotal voter model, the flatness of the slope suggests that turnout is unlikely to be fully accounted for by the pivotal voter model.

2.4. Identification

In this section, we discuss the identification of our model as the number of counties *within each state* becomes large ($M_s \to \infty$). Given that we have state-specific parameters for $p_s(\cdot)$ and $b_s(\cdot)$, we require the number of observations in each state to be large. Our discussion

[21] In addition, we drop one county, Chattahoochee, GA, as the turnout rate is extremely low (18.8%) relative to all other counties. The turnout rate for the next lowest county is 33%.

Table 2.1. Summary Statistics of Voting Outcome and Demographic Characteristics of Eligible Voters.

Voting Data					
	Obs	Mean	Std. Dev.	Min	Max
Vote Share: Democrat	2,909	0.24	0.08	0.04	0.57
Vote Share: Republican	2,909	0.37	0.09	0.07	0.70
1 - Turnout Rate	2,909	0.40	0.09	0.12	0.67

County Demographics					
	Obs	Mean	Std. Dev.	Min	Max
% Hispanic	2,909	0.05	0.12	0.00	0.97
% Black/African American	2,909	0.09	0.15	0.00	0.87
% Neither Black nor White	2,909	0.05	0.08	0.00	0.95
Mean Age	2,909	46.75	2.64	35.94	56.07
Mean Income (USD 1,000)	2,909	42.71	9.23	23.33	93.40
Mean Years of Schooling	2,909	12.86	0.60	10.84	15.18
% Religious	2,909	0.26	0.17	0.00	1.00

Note: For Age, Income, and Years of Schooling, the table reports the mean, standard deviation, minimum, and maximum of the county mean. “% Religious” is the share of the population with adherence to either “Evangelical Denomination” or “Church of Jesus Christ of Latter-day Saints.”

in this section builds on the idea initially proposed by Barkume (1976) in the context of property tax referenda for school districts.

Recall that the observed vote shares are expressed as:

$$\begin{aligned}
v_{R,m} &\equiv \int 1 - \Phi\left(\max\left\{-b_s(\mathbf{x}_n) - \xi_m, -b_s(\mathbf{x}_n) - \xi_m + \frac{c_s(\mathbf{x}_n) + \eta_m}{p_s(\mathbf{x}_n)}\right\}\right) dF_{\mathbf{x},m}(\mathbf{x}_n), \\
v_{D,m} &\equiv \int \Phi\left(\min\left\{-b_s(\mathbf{x}_n) - \xi_m, -b_s(\mathbf{x}_n) - \xi_m - \frac{c_s(\mathbf{x}_n) + \eta_m}{p_s(\mathbf{x}_n)}\right\}\right) dF_{\mathbf{x},m}(\mathbf{x}_n), \\
v_{0,m} &\equiv 1 - v_{D,m} - v_{R,m}.
\end{aligned}$$

Figure 2.2. Relationship between the Ex-Post Winning Margin and Voter Turnout.

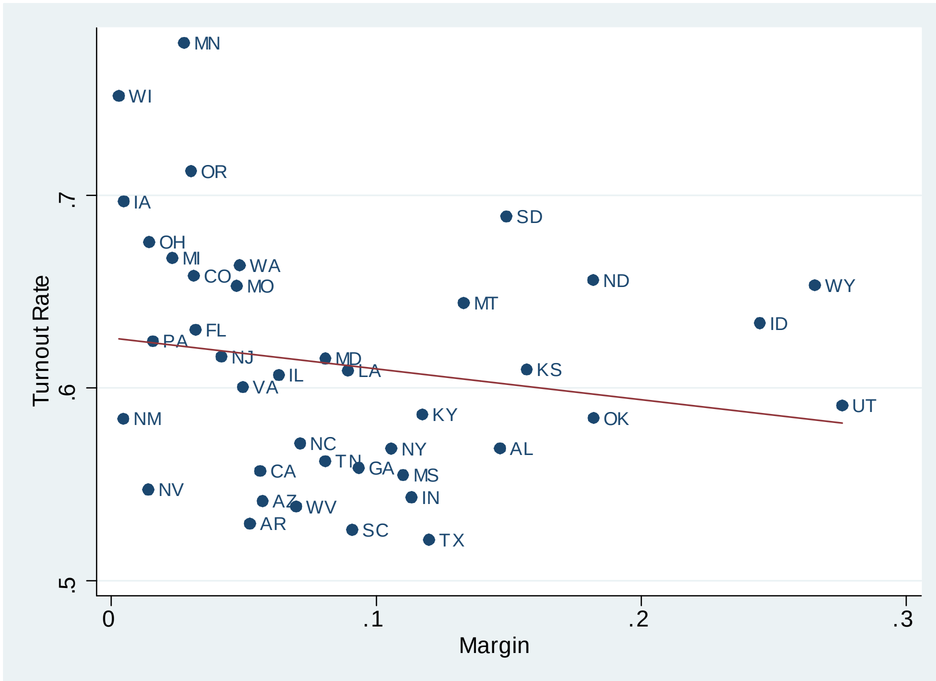

Note: The slope coefficient is -0.16 and not statistically significant.

For exposition, consider the simple case in which there is no heterogeneity in voters' observable characteristics, so that $\mathbf{x}_n = \overline{\mathbf{x}}_m$ for all n in county m.[22] In this case, the above expressions simplify as follows:

$$(2.4.1) \qquad v_{R,m} \equiv 1 - \Phi\left(\max\left\{-b_s(\overline{\mathbf{x}}_m) - \xi_m, -b_s(\overline{\mathbf{x}}_m) - \xi_m + \frac{c_s(\overline{\mathbf{x}}_m) + \eta_m}{p_s(\overline{\mathbf{x}}_m)}\right\}\right),$$

$$(2.4.2) \qquad v_{D,m} \equiv \Phi\left(\min\left\{-b_s(\overline{\mathbf{x}}_m) - \xi_m, -b_s(\overline{\mathbf{x}}_m) - \xi_m - \frac{c_s(\overline{\mathbf{x}}_m) + \eta_m}{p_s(\overline{\mathbf{x}}_m)}\right\}\right),$$

$$(2.4.3) \qquad v_{0,m} \equiv 1 - v_{D,m} - v_{R,m}.$$

[22]Note that we are well aware of the issues of ecological fallacy. In what follows, we consider a simplified setup with $\mathbf{x}_n = \overline{\mathbf{x}}_m$ for all n in county m, just for expositional purposes. In our empirical exercise, we fully address the fact that each county has a distribution of $\mathbf{x}$ by integrating the vote share for each $\mathbf{x}$ with respect to $F_{\mathbf{x},m}(\cdot)$.

We now show that the primitives of the model are identified from expressions (2.4.1), (2.4.2), and (2.4.3).

Using the fact that Φ is a strictly increasing function, we can rewrite expressions (2.4.1) and (2.4.2) as follows:

$$
\begin{aligned}
\Phi^{-1}(1-v_{R,m}) &= \max\left\{-b_s(\overline{\mathbf{x}}_m)-\xi_m,\ -b_s(\overline{\mathbf{x}}_m)-\xi_m+\frac{c_s(\overline{\mathbf{x}}_m)+\eta_m}{p_s(\overline{\mathbf{x}}_m)}\right\}, \\
\Phi^{-1}(v_{D,m}) &= \min\left\{-b_s(\overline{\mathbf{x}}_m)-\xi_m,\ -b_s(\overline{\mathbf{x}}_m)-\xi_m-\frac{c_s(\overline{\mathbf{x}}_m)+\eta_m}{p_s(\overline{\mathbf{x}}_m)}\right\}.
\end{aligned}
$$

Rearranging these two equations, we obtain the following expressions:

$$
\frac{\Phi^{-1}(1-v_{R,m})+\Phi^{-1}(v_{D,m})}{-2} = b_s(\overline{\mathbf{x}}_m)+\xi_m, \text{ and} \tag{2.4.4}
$$

$$
\frac{\Phi^{-1}(1-v_{R,m})-\Phi^{-1}(v_{D,m})}{2} = \max\left\{0, \frac{c_s(\overline{\mathbf{x}}_m)}{p_s(\overline{\mathbf{x}}_m)}+\frac{\eta_m}{p_s(\overline{\mathbf{x}}_m)}\right\}. \tag{2.4.5}
$$

Note that the left hand side of (2.4.4) closely reflects the difference in vote share, and the left hand side of (2.4.5) reflects the turnout rate. This is because, if we ignore the nonlinearity of $\Phi^{-1}(\cdot)$ and the denominator, the left hand side of (2.4.4) reduces to $1-v_{R,m}+v_{D,m}$ and the left hand side of (2.4.5) to $1-v_{R,m}-v_{D,m}$. The former is one minus the difference in vote share, and the latter is one minus voter turnout. The left hand side of expressions (2.4.4) and (2.4.5) can be directly computed using data on vote shares, $v_{D,m}$ and $v_{R,m}$.

We first consider the identification of $b_s(\cdot)$ and the distribution of ξ, $F_\xi(\cdot)$. Taking the expectation of (2.4.4) conditional on $\overline{\mathbf{x}}_m$, we have

$$
\mathbf{E}\left[\left.\frac{\Phi^{-1}(1-v_{R,m})+\Phi^{-1}(v_{D,m})}{-2}\right|\overline{\mathbf{x}}_m\right] = b_s(\overline{\mathbf{x}}_m), \tag{2.4.6}
$$

because $\mathbf{E}\left[\xi_m|\overline{\mathbf{x}}_m\right] = 0$. As the left hand side of the above expression is identified, $b_s(\cdot)$ is (nonparametrically) identified for each s (note that the asymptotics is with respect to the number of counties within each state). Given that $b_s(\cdot)$ is the utility difference between Republicans and Democrats, it is intuitive that $b_s(\cdot)$ is identified by the difference in vote share. In this model, changes in $b_s(\cdot)$ is reflected in the vote share margin which cannot be undone by changes in $c_s(\cdot)$.

Now, consider the identification of $F_\xi(\cdot)$. Given that $b_s(\cdot)$ is identified and the left hand side of (2.4.4) is observable, each realization of ξ_m can be recovered from (2.4.4). Hence, $F_\xi(\cdot)$ is also identified. Note that if $b_s(\cdot)$ is linear in $\overline{\mathbf{x}}_m$ (i.e., $b_s(\overline{\mathbf{x}}_m) = \beta\overline{\mathbf{x}}_m$), one can simply regress the left hand side of expression (2.4.4) on $\overline{\mathbf{x}}_m$ by OLS to obtain β as coefficients and ξ_m as residuals.

We now discuss the identification of $p_s(\cdot)$, $c_s(\cdot)$, and $F_\eta(\cdot)$. For simplicity, consider the case in which the second term inside the max operator of expression (2.4.5) is positive with probability 1;

$$\frac{\Phi^{-1}\left(1 - v_{R,m}\right) - \Phi^{-1}\left(v_{D,m}\right)}{2} = \frac{c_s(\overline{\mathbf{x}}_m)}{p_s(\overline{\mathbf{x}}_m)} + \frac{\eta_m}{p_s(\overline{\mathbf{x}}_m)}. \tag{2.4.7}$$

This corresponds to the case in which the turnout rate is always less than 100%. We show in Appendix C that $p_s(\cdot)$, $c_s(\cdot)$, and $F_\eta(\cdot)$ are identified without this assumption.

Taking the conditional moments of (2.4.7), we have

$$\mathbf{E}\left[\left.\frac{\Phi^{-1}\left(1 - v_{R,m}\right) - \Phi^{-1}\left(v_{D,m}\right)}{2}\right|\overline{\mathbf{x}}_m\right] = \frac{c_s(\overline{\mathbf{x}}_m)}{p_s(\overline{\mathbf{x}}_m)}, \text{ and} \tag{2.4.8}$$

$$\mathbf{Var}\left[\left.\frac{\Phi^{-1}\left(1 - v_{R,m}\right) - \Phi^{-1}\left(v_{D,m}\right)}{2}\right|\overline{\mathbf{x}}_m\right] = \frac{\sigma_\eta^2}{(p_s(\overline{\mathbf{x}}_m))^2}, \tag{2.4.9}$$

where σ_η^2 is the variance of η_m. Using (2.4.9), $p_s(\cdot)$ is identified up to a scalar constant (i.e., up to σ_η^2) because the left hand side of (2.4.9) is identified. This implies that $c_s(\overline{\mathbf{x}}_m)$ is also identified up to σ_η^2 using (2.4.8). Given that $p_s(\cdot)$ and $c_s(\cdot)$ are identified, we can recover the realization of η_m from (2.4.7), implying that $F_\eta(\cdot)$ is also identified up to σ_η^2. Intuitively, the left hand side of (2.4.8) and (2.4.9) closely reflect the mean and variance of the rate of abstention. Hence, the average cost of voting normalized by p_s (i.e., $c_s(\cdot)/p_s(\cdot)$) is identified from (2.4.8).

Importantly, while $p_s(\cdot)$ is identified only up to a scalar constant, the ratio $p_{s'}/p_{s''}$ for any states s' and s'' is identified given our specification of $p_s(\cdot)$ as $p_s(\cdot) = p_s \times \tilde{p}(\cdot)$. To see this, note that the ratio of (2.4.9) for two counties with the same demographics in states s' and s'' directly identify $p_{s'}/p_{s''}$.

The discussion has, thus far, been based on the simplified case in which all voters in county m have the same demographic characteristics–i.e., $\mathbf{x}_n = \overline{\mathbf{x}}_m$ for all n in m. As long as there is sufficient variation in $F_{\mathbf{x},m}(\mathbf{x})$, we can recover the vote shares conditional on each $\mathbf{x}$ and apply the identification discussion above.

Correlation between Unobserved Cost and Preference Shocks. Our identification makes no assumptions regarding the correlation between the unobservables ξ_m and η_m. As ξ_m and η_m enter separately in (2.4.4) and (2.4.5), $\xi_m \perp \overline{\mathbf{x}}_m$ and $\eta_m \perp \overline{\mathbf{x}}_m$ are sufficient to identify the unknown primitives on the right hand side in each equation. Hence, we do not require any restrictions on the joint distribution of ξ_m and η_m. In fact, we can nonparametrically identify the joint distribution of ξ_m and η_m from the joint distribution of the residuals in each equation. In our estimation, we specify the joint distribution of ξ_m and η_m as a bivariate Normal with correlation coefficient ρ.

Exogenous Cost Shifters. Lastly, we discuss identification when there exist instruments (e.g., rainfall) that shift the cost of voting but not the preferences of the voters. The point we wish to make is that the existence of exogenous cost shifters are neither necessary nor sufficient for identification.

To illustrate this point, consider the following discrete choice setup with instruments $\mathbf{z}_n$,

$$\begin{aligned} V_A &= u_A(\mathbf{x}_n) - c_{\mathbf{x}}(\mathbf{x}_n) - c_{\mathbf{z}}(\mathbf{z}_n), \\ V_B &= u_B(\mathbf{x}_n) - c_{\mathbf{x}}(\mathbf{x}_n) - c_{\mathbf{z}}(\mathbf{z}_n), \\ V_0 &= 0, \end{aligned}$$

where V_k denotes the mean utility of choosing $k \in \{A, B, 0\}$. Here, $u_A(\mathbf{x}_n)$ is not necessarily equal to $-u_B(\mathbf{x}_n)$, and the cost function is separated into two components, $c_{\mathbf{x}}(\mathbf{x}_n)$ and $c_{\mathbf{z}}(\mathbf{z}_n)$, where $\mathbf{z}_n$ is a vector of cost shifters excluded from $u_k(\mathbf{x}_n)$. For any arbitrary function $g(\mathbf{x}_n)$, consider an alternative model with $\tilde{u}_k(\mathbf{x}_n) = u_k(\mathbf{x}_n) + g(\mathbf{x}_n)$ ($k \in \{A, B\}$) and $\tilde{c}_{\mathbf{x}}(\mathbf{x}_n) = c_{\mathbf{x}}(\mathbf{x}_n) + g(\mathbf{x}_n)$, as follows:

$$\begin{aligned} V_A &= \tilde{u}_A(\mathbf{x}_n) - \tilde{c}_{\mathbf{x}}(\mathbf{x}_n) - c_{\mathbf{z}}(\mathbf{z}_n), \\ V_B &= \tilde{u}_B(\mathbf{x}_n) - \tilde{c}_{\mathbf{x}}(\mathbf{x}_n) - c_{\mathbf{z}}(\mathbf{z}_n), \\ V_0 &= 0. \end{aligned}$$

Because $\tilde{u}_k(\mathbf{x}_n) - \tilde{c}_{\mathbf{x}}(\mathbf{x}_n) = u_k(\mathbf{x}_n) - c_{\mathbf{x}}(\mathbf{x}_n)$, the two models are observationally equivalent, and thus, $u_k(\cdot)$ and $c_{\mathbf{x}}(\cdot)$ are not separately identified. In this model, one cannot differentiate between the case in which voters have intense preferences and large voting costs and one in which voters have weak preferences and small costs. However, in the model we take to the data, intense preferences for one of the parties get reflected in the vote share margin, which cannot be offset by increasing the voting cost. It is *not* the availability of instruments, but rather the observation that we can express $u_A(\mathbf{x}_n) = -u_B(\mathbf{x}_n)$ that identifies the primitives of the model.

2.5. Specification and Estimation

2.5.1. Specification

We now specify $b_s(\cdot)$, $c_s(\cdot)$, $p_s(\cdot)$ and the joint distribution of ξ_m and η_m for our estimation. We specify $b_s(\cdot)$, which is the utility difference from having candidates R and D in office, as a function of a state-level preference intercept, λ_s, and demographic characteristics, $\mathbf{x}_n$, consisting of age, race, income, religion, and years of schooling:

$$b_s(\mathbf{x}_n) = \lambda_s + \beta_b' \mathbf{x}_n.$$

The intercept, λ_s, is a parameter that we estimate for each state. It captures the state-level net preference shock for the Republicans that demographic characteristics do not account for. Note that the linear specification for the utility difference can be derived from a spatial voting model in which a voter has quadratic loss and her bliss point is linear in $\mathbf{x}_n$.[23]

Voting cost, $c(\cdot)$, is also specified as a linear function of $\mathbf{x}_n$ as

$$c(\mathbf{x}_n) = \beta_c[1, \mathbf{x}_n']'.$$

[23]To illustrate this point, consider a unidimensional spatial voting model in which candidate D's ideological position is 0; candidate R's position is 1; and a voter's bliss point is $\alpha_n = \beta^{bliss}\mathbf{x}_n$. Under the quadratic loss function, the utility from electing candidates D and R are $-\alpha_n^2$ and $-(1-\alpha_n)^2$, respectively, and the utility difference, $b_s(\mathbf{x}_n)$, is written as $-(1-\alpha_n)^2 + \alpha_n^2 = 2\beta^{bliss}\mathbf{x}_n - 1$. Thus, $b_s(\mathbf{x}_n)$ is linear in $\mathbf{x}_n$ in such a model.

We do not specifically model the presence of other elections, such as gubernatorial and congressional elections, because previous studies (e.g., Smith, 2001) find that neither the presence nor the closeness of other elections affects turnout in presidential elections. We also do not include weather-related variables in $c(\cdot)$ because there was insufficient variation in precipitation and temperature on the day of the 2004 presidential election to have affected turnout significantly.[24] Although $c(\cdot)$ is allowed to be state specific in the identification, we opt for a simpler specification to keep the number of parameters manageable.[25]

We specify the voter's perception of efficacy as $p_s \times \tilde{p}(\mathbf{x}_n)$, where $\tilde{p}(\cdot)$ is a function of her age, income, and years of schooling, as follows:[26]

$$\tilde{p}(\mathbf{x}_n) = \exp(\beta_p' \mathbf{x}_n).$$

We normalize $p_s = 1$ for Alabama and normalize $\tilde{p}(\cdot)$ such that $\tilde{p}(\bar{\mathbf{x}}) = 1$, where $\bar{\mathbf{x}}$ is the national average of $\mathbf{x}_n$.[27]

[24]We included weather variables in the simple model that assumes $\mathbf{x}_n = \bar{\mathbf{x}}_m$ (i.e., the demographic characteristics of voters in each county are assumed to be the same within county) and found the coefficients on the weather variables to be small and insignificant.

[25]If we include a state-specific cost intercept, this increases the number of parameters by the number of states. Given that some states have only a moderate number of counties, we found it hard to include a state-specific cost term in addition to λ_s and p_s.

[26]The set of variables included in $p_s(\mathbf{x}_n)$ is a subset of $\mathbf{x}_n$ that takes continuous values. Here, we do not include dummy variables such as race, and religion. The reason is as follows. The variation in $c(\cdot)$ changes the utility level additively, while the variation in $p_s(\cdot)$ changes it multiplicatively as $p_s \times b_s$. As dummy variables take only 0 and 1, it is difficult, in practice, to distinguish whether the effects of those variables are additive or multiplicative. Thus, estimating the model with dummy variables in both cost and efficacy is difficult, and we include only continuous variables in $p_s(\cdot)$.

[27]Note that we need two normalizations. Because we express $p_s(\mathbf{x}_n)$ as $p_s \times \tilde{p}(\mathbf{x}_n)$, we need a scalar normalization on either p_s or $\tilde{p}(\mathbf{x}_n)$. We normalize $p_s = 1$ for Alabama. We also need an additional normalization because $p_s(\cdot)$ is identified only up to the variance of η– i.e., the level of p_s is not identified in our model. Assuming that $\tilde{p}(\bar{\mathbf{x}}) = 1$ eliminates this degree of freedom.

We specify the joint distribution of county-level preference shock ξ and cost shock η as a bivariate normal, $N(0,\Sigma)$, where Σ is the variance-covariance matrix with diagonal elements equal to σ_ξ^2, σ_η^2 and off-diagonal elements $\rho\sigma_\xi\sigma_\eta$.

2.5.2. Estimation

We use the method of moments to estimate the model parameters.[28] Recall that the vote shares (as a fraction of eligible voters) and turnout in county m are given by expressions (2.2.4), (2.2.5) and (2.2.6), where $F_{\mathbf{x},m}$ is the distribution of $\mathbf{x}_n$ in county m. For a fixed vector of the model parameters, $\theta = (\beta_b, \{\lambda_s\}, \beta_c, \{p_s\}, \beta_p, \sigma_\xi, \sigma_\eta, \rho)$, we can compute the moments of expressions (2.2.4), (2.2.5) and (2.2.6) by integrating over ξ and η. Our estimation is based on matching the moments generated by the model with the corresponding sample moments.

Specifically, we define the first and second order moments implied by the model as follows:

$$\begin{aligned}
\hat{v}_{k,m}(\theta) &= \mathbf{E}_{\xi,\eta}[v_{k,m}(\xi,\eta;\theta)],\ \forall k \in \{D,R\}, \\
\hat{v}_{k,m}^{squared}(\theta) &= \mathbf{E}_{\xi,\eta}[v_{k,m}(\xi,\eta;\theta)^2],\ \forall k \in \{D,R\}, \\
\hat{v}_{m}^{cross}(\theta) &= \mathbf{E}_{\xi,\eta}[v_{D,m}(\xi,\eta;\theta)v_{R,m}(\xi,\eta;\theta)],
\end{aligned}$$

where $v_{k,m}(\xi,\eta;\theta)$ is the vote share of candidate k given a realization of (ξ,η) and parameter θ.[29,30] Denoting the observed vote share of candidate k in county m as $\mathsf{v}_{k,m}$, our objective

[28]In contrast to maximum likelihood estimation, which requires us to solve for (ξ,η) that rationalizes the observed vote share for each parameter value, the method of moments only requires integration with respect to ξ and η by simulation. The latter is substantially less costly in terms of computation.

[29]Computing $\hat{v}_{k,m}(\theta)$, $\hat{v}_{k,m}^{squared}(\theta)$, and $\hat{v}_{m}^{cross}(\theta)$ requires integration over (ξ,η). For integration, we use a quadrature with $[5\times 5]$ nodes and pruning (see Jäckel, 2005) with a total of 21 nodes.

[30]Note that we do not need to know the value of ρ for computing $v_{k,m}(\xi,\eta;\theta)$. Thus, for each $\theta\backslash\{\rho\}$ (i.e., the parameters except for ρ), we can compute $v_{k,m}(\xi,\eta;\theta)$ and solve for the implied realizations (ξ_m,η_m) that give the observed vote shares in county m. Berry (1994) guarantees that there exists a unique pair of (ξ_m,η_m) that equates $v_{k,m}(\xi,\eta;\theta)$ to the observed shares. By computing the correlation between the implied realizations of ξ_m and η_m, we can obtain the value of ρ that is consistent with the observed data.

function, $J(\theta)$, is given by

$$J(\theta)=\sum_{k=\{D,R\}}\left(\frac{J_{1,k}(\theta)}{\widehat{\mathbf{Var}}(\mathsf{v}_{k,m})}+\frac{J_{2,k}(\theta)}{\widehat{\mathbf{Var}}(\mathsf{v}_{k,m}^2)}\right)+\frac{J_3(\theta)}{\widehat{\mathbf{Var}}(\mathsf{v}_{D,m}\mathsf{v}_{R,m})},$$

where

$$\begin{aligned}
J_{1,k}(\theta) &= \frac{1}{M}\sum_{s=1}^{S}\sum_{m=1}^{M_s}(\hat{v}_{k,m}(\theta)-\mathsf{v}_{k,m})^2,\ \forall k\in\{D,R\},\\
J_{2,k}(\theta) &= \frac{1}{M}\sum_{s=1}^{S}\sum_{m=1}^{M_s}(\hat{v}_{k,m}^{squared}(\theta)-\mathsf{v}_{k,m}^2)^2,\ \forall k\in\{D,R\},\\
J_3(\theta) &= \frac{1}{M}\sum_{s=1}^{S}\sum_{m=1}^{M_s}\left(\hat{v}_m^{cross}(\theta)-\mathsf{v}_{D,m}\mathsf{v}_{R,m}\right)^2.
\end{aligned}$$

$J_{1,k}$ is the sum of the squared differences between the expectation of the predicted vote share ($\hat{v}_{k,m}(\theta)$) and the actual vote share ($\mathsf{v}_{k,m}$). $J_{2,k}$ is the sum of the squared differences between $\hat{v}_{k,m}^{squared}(\theta)$ and the squared vote share, $\mathsf{v}_{k,m}^2$. J_3 is the sum of the squared differences between the predicted and actual cross terms. M is the total number of counties across all states, $\sum_{s=1}^{S} M_s$, and $\widehat{\mathbf{Var}}(z)$ denotes the sample variance of z.

The construction of our objective function follows our identification argument closely. The first moment, $J_{1,k}$, matches the conditional expectation of the vote shares from the model with that from the data. Intuitively, $J_{1,k}$ corresponds to (2.4.6) and (2.4.8), and pins down β_b, $\{\lambda_s\}$, β_c/p_s, β_c/β_p and $p_s/p_{s'}$. The second and third moments, $J_{2,k}$ and J_3, correspond to (2.4.9). These moments pin down β_p,σ_ξ,σ_η and ρ.

We then impose this value of ρ to compute $\hat{v}_{k,m}(\theta)$, $\hat{v}_{k,m}^{squared}(\theta)$, and $\hat{v}_m^{cross}(\theta)$. Our estimation procedure can be thought of as minimizing the objective function with respect to $\theta\backslash\{\rho\}$, and the estimate of ρ is obtained by computing the correlation between the implied realizations of ξ_m and η_m given the estimate of $\theta\backslash\{\rho\}$.

2.6. Results

The set of parameters that we estimate include those that are common across all states (β_b,σ_ξ, β_c,σ_η,β_p,ρ) and those that are specific to each state ($\{\lambda_s\}$,$\{p_s\}$). Table 2.2 reports the estimates of the former set, while Figures 2.4 and 2.5 plot the parameter estimates of the latter set.

Estimates of β_b, σ_ξ, β_c, σ_η, β_p, and ρ. The first column of Table 2.2 reports the estimates of the preference parameters. We find that Age and Income enter the utility difference, $b_R - b_D$, positively, implying that old and high-income voters are more likely to prefer Republicans. Years of Schooling enters the utility difference negatively, thus more-educated voters are more likely to prefer Democrats. We also find that Hispanics, African Americans, and Other Races prefer Democrats. The Religion variable carries a positive coefficient, implying that religious voters prefer Republicans. The estimate of the constant term corresponds to the preference of the voter who has $\mathbf{x}_n$ equal to the national average and has λ_s equal to that of Alabama.

In the second column of Table 2.2, we report the estimates of the cost parameters. The estimated costs are inclusive of any benefits from fulfilling civic duty. Moreover, our cost estimates include not only physical and opportunity costs but also psychological costs, such as information acquisition costs. We find that Age, Years of Schooling, and Income enter voting cost negatively. This implies that older, more-educated, and higher-income voters have a lower cost of voting.[31] Hispanics, African Americans, and Other Races have a higher cost of voting relative to non-Hispanics and Whites. Religious voters also have a higher cost

[31] Given that high-income and more-educated voters tend to have high opportunity cost, the negative coefficients on income and education might suggest that information acquisition cost can be an important part of the voting cost.

Table 2.2. Parameter Estimates

	Preference		Cost		Efficacy	
	Estimate	SE	Estimate	SE	Estimate	SE
Age	0.0185	(0.0036)	-0.0086	(0.0091)	0.0648	(0.0157)
Years of Schooling	-0.0788	(0.0160)	-0.2470	(0.0655)	-0.1694	(0.0623)
log(income)	0.3747	(0.0496)	-0.1996	(0.0825)	-0.0534	(0.1728)
Hispanic	-0.2474	(0.0843)	0.0750	(0.0726)		
African American	-1.3392	(0.0592)	0.2799	(0.0807)		
Other Races	-0.7380	(0.0745)	0.2502	(0.0859)		
Religious	0.6616	(0.0577)	0.1442	(0.0440)		
Constant	0.1034	(0.0374)	0.4778	(0.0341)		
Sigma	-0.1978	(0.0086)	0.0097	(0.9045)		
Rho		-0.0840 (0.0521)				

Note: The table reports the parameter estimates of voters' preferences, costs, and perception of voting efficacy. The estimate of the constant terms in the first and second columns corresponds to the preference and costs of the voter who has $\mathbf{x_n}$ equal to the national average and has λ_s equal to that of Alabama. The variable log(income) is the log of income divided by 1000. Excluded categories are non-Hispanic, White, and non-Religious. Standard errors are computed by analytically deriving the asymptotic variance covariance matrix. The standard errors are reported in parentheses.

of voting compared to non-religious voters.The estimate of the constant term in the second column corresponds to the cost of a voter whose characteristics are set to the national mean.

The third column of Table 2.2 reports the estimates of the efficacy parameters. We find that Age enters the perception of efficacy positively, while Years of Schooling and Income enter negatively. This implies that older, less-educated, and lower-income citizens tend to have a higher perception of efficacy. Given that older voters have lower voting costs as well, they are more likely to be overrepresented than young voters. Regarding Years of Schooling and Income, the overall effect on participation depends on the relative magnitudes of the cost and efficacy coefficients. We discuss the net effect in the next subsection.

The last row of Table 2.2 reports the estimate of ρ, which is the correlation between unobservable shocks ξ and η. The estimate is negative (–0.084), implying that the correlation in the unobservable shocks tends to suppress turnout among voters who prefer Democrats.

Representation and Preference Aggregation. To better understand what our estimates imply about representation across demographic groups and their preferences, Figure 3 plots the estimated share of voters who prefer Democrats over Republicans (right axis) and an estimated measure of voter representation (left axis) by demographic groups. The share of voters who prefer Democrats over Republicans is simply the two-party vote share of the Democrats unconditional on turnout, computed using the preference estimates. The representation measure that we use is based on Wolfinger and Rosenstone (1980), and it is defined as follows:

$$WR(\mathbf{x}) = \frac{\text{share of group } \mathbf{x} \text{ among those who turn out}}{\text{share of group } \mathbf{x} \text{ among the overall electorate}},$$

where $\mathbf{x}$ is a demographic group. Overrepresented demographic groups have representation measures larger than one, while underrepresented groups have measures less than one.

Figure 2.3 shows that there are significant differences in the representation and preference measures among demographic groups. For example, panel (a) shows that the representation measure of 75-year-old voters is 1.65, while that of 25-year-old voters is 0.37. The fraction of 75-year-old voters who prefer Democrats is 36.4%, while that of 25-year-old voters is 59.9%.

Figure 2.3 also illustrates how demographic characteristics are related to preference aggregation. In particular, if the two curves in each panel have the same (opposite) slope, the voters who prefer Republicans (Democrats) are underrepresented. For example, in panels (a), (c), and (d), the overrepresented groups tend to prefer Republicans, while the underrepresented groups tend to prefer Democrats.

Figure 2.3. Representation and Preference by Demographic Characteristics.

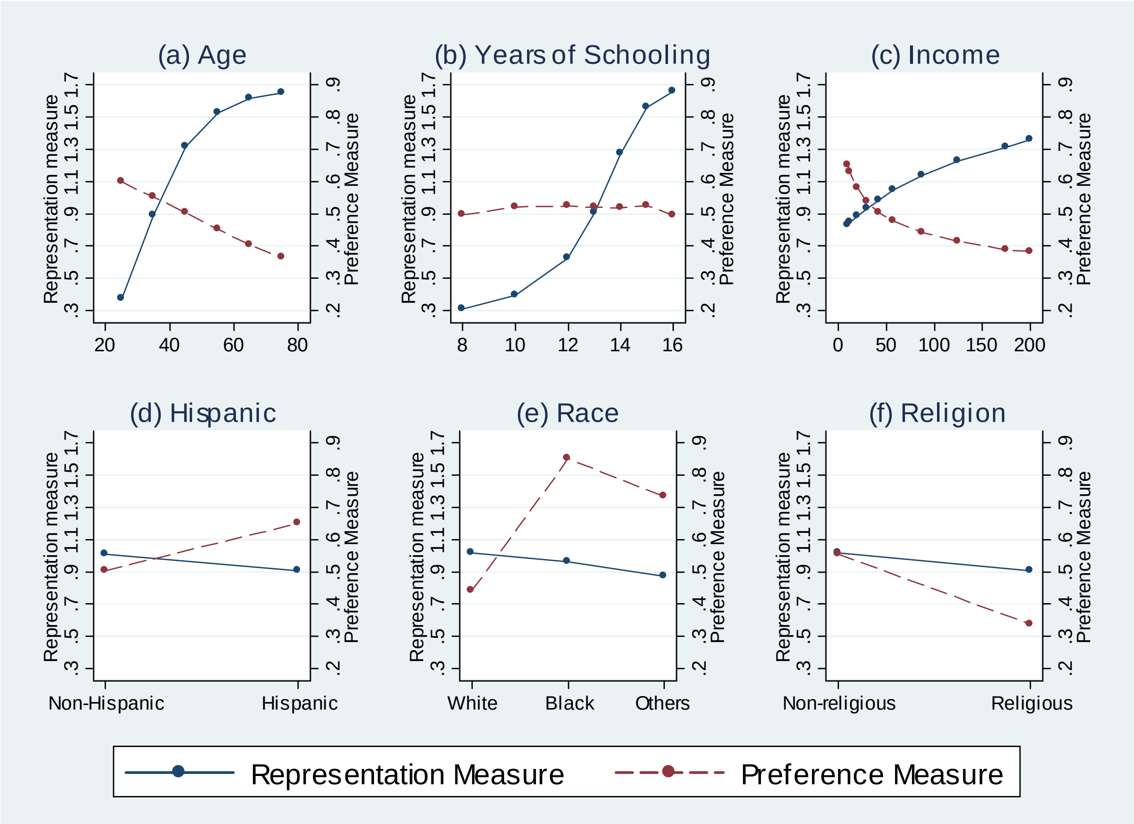

Note: The horizontal axis corresponds to the level/category of the demographic variable. The left vertical axis corresponds to the representation measure, and the right vertical axis corresponds to preference of the group in terms of the two-party vote share. The horizontal axis in Panel (c) is income in units of 1,000 USD.

Panel (a) of Figure 2.3 shows that old voters are overrepresented *and* prefer Republicans, while young voters are underrepresented *and* tend to prefer Democrats. Similarly, low-income voters and those who are Hispanic, African American, and of Other Races are underrepresented and tend to prefer Democrats. On the other hand, panel (f) shows that religious voters are underrepresented and prefer Republicans. These results show that there is a systematic selection in the preferences of those who turn out.

Figure 2.4. Estimates of the State-level Preference Intercept, λ_s

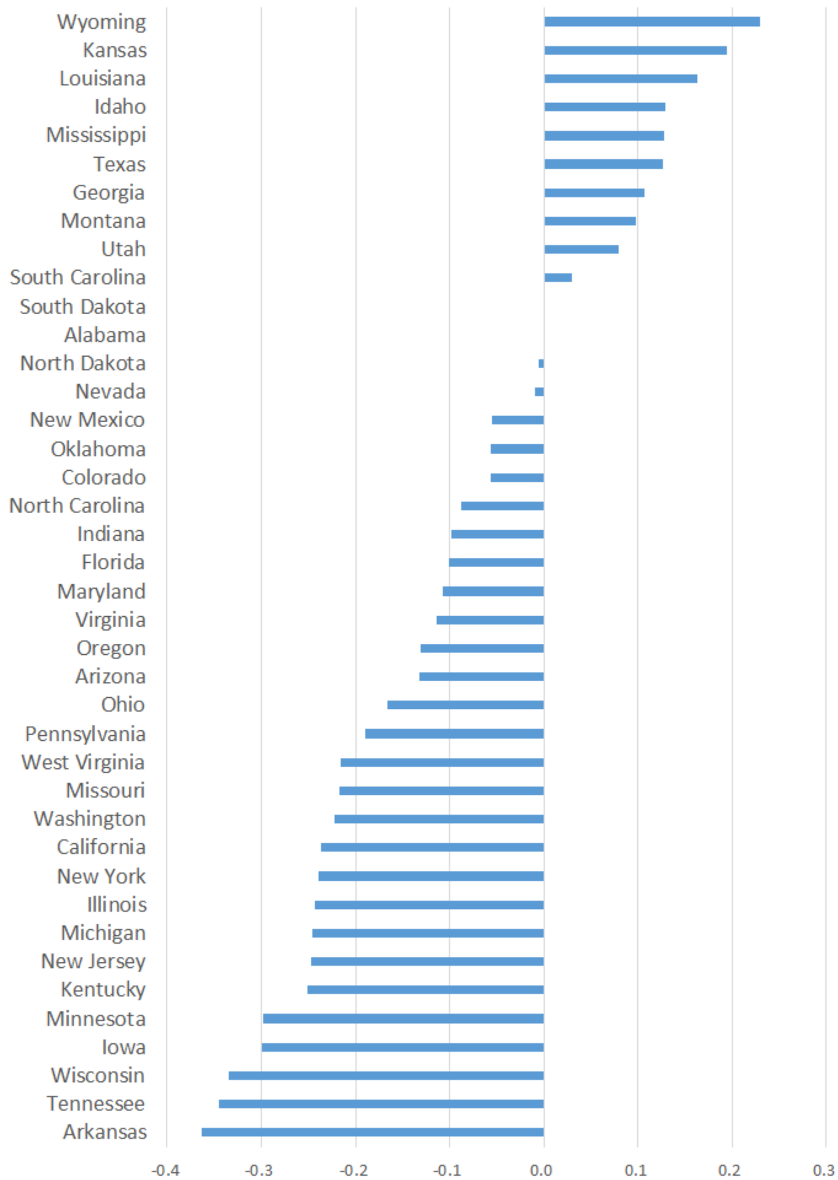

Note: $\lambda_{Alabama}$ is normalized to zero.

Estimates of State-Specific Effects, λ_s and p_s. Figure 2.4 plots our estimates of the state specific intercepts in the voter's utility relative to that of Alabama, which is normalized to 0 (Table B.1 in the Online Appendix reports the point estimates and the standard errors).

Figure 2.5. Estimates of the State-level Efficacy Coefficient, p_s.

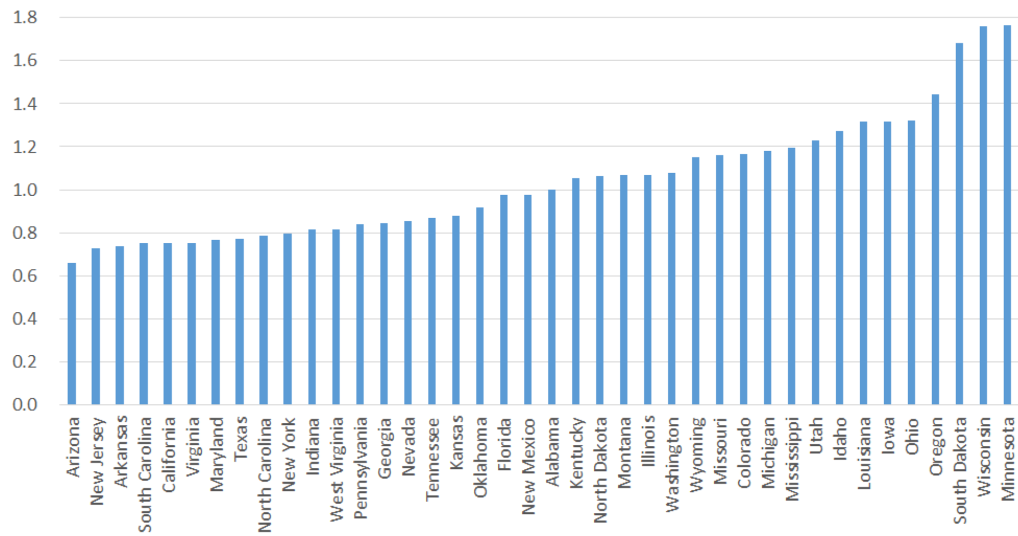

Note: Our estimates of p_s are relative to $p_{Alabama}$ where $p_{Alabama}$ is normalized to 1.

Larger values of λ_s imply that the voters in the corresponding state prefer Republicans, *net of the effect of demographic characteristics*. These state fixed effects may include the inherent preferences of voters and/or the effect of campaign activities by candidates. The figure show that states such as Wyoming, Kansas, and Louisiana have the strongest preference for Republicans, while states such as Arkansas, Tennessee, and Wisconsin have the strongest preference for Democrats. Overall, Democratic strongholds such as New York and California tend to have low estimated values of λ_s, while Republican strongholds such as Georgia and Texas tend to have high estimates.

Figure 2.5 plots the estimates of the state-specific component of the perception of voting efficacy, p_s, with normalization $p_{Alabama} = 1$ (Table B.2 in the Online Appendix reports the point estimates and standard errors). High values of p_s correspond to high perception of voting efficacy, after controlling for demographics. The perception of voting efficacy varies across states, which may partly reflect the fact that the electors are determined at the state level. We find that battleground states such as Minnesota, Wisconsin, Ohio and Iowa have some of the highest estimated values of voting efficacy. We also find that states considered party strongholds, such as California and Texas, have low estimated values. These results suggest a positive relationship between perception of voting efficacy and pivot probability.

While some of the forces of the pivotal voter model seem to be at play, the estimated values of p_s suggest that the rational expectations pivotal voter model is unlikely to explain the overall voting pattern very well. Models of voting based on rational expectations would require p_s in battleground states to be orders of magnitude greater than those in party strongholds (see, e.g., Shachar and Nalebuff, 1999). However, our estimates of p_s fall within a narrow range: the ratio of the estimated state-level efficacy parameters, $p_s/p_{s'}$, is, at most, three. Our results highlight the importance of relaxing the assumption of rational expectations on the pivot probability.

Fit. To assess the fit of our model, Figure 2.6 plots the county-level vote share, voter turnout, and vote share margin predicted from the model against the data. The predicted vote share, turnout, and vote share margin are computed by evaluating expressions (2.2.4), (2.2.5), and (2.2.6) at the estimated parameter and integrating out ξ and η. The plots line up around the 45-degree line, which suggests that the model fits the data well. In Online Appendix B, we provide further discussion of fit.

In previous work, Coate et al. (2008) discusses the difficulty of fitting the winning margin using the rational expectations pivotal voter model. In our paper, we do not impose the rational expectations assumption, and the model fits the winning margin in the data well.

Turnout and Preference Intensity. We now discuss how preference intensity is related to turnout and how this affects preference aggregation. Our discussion is closely related to the theoretical work of Campbell (1999) that shows that minorities with intense preferences can

Figure 2.6. Fit

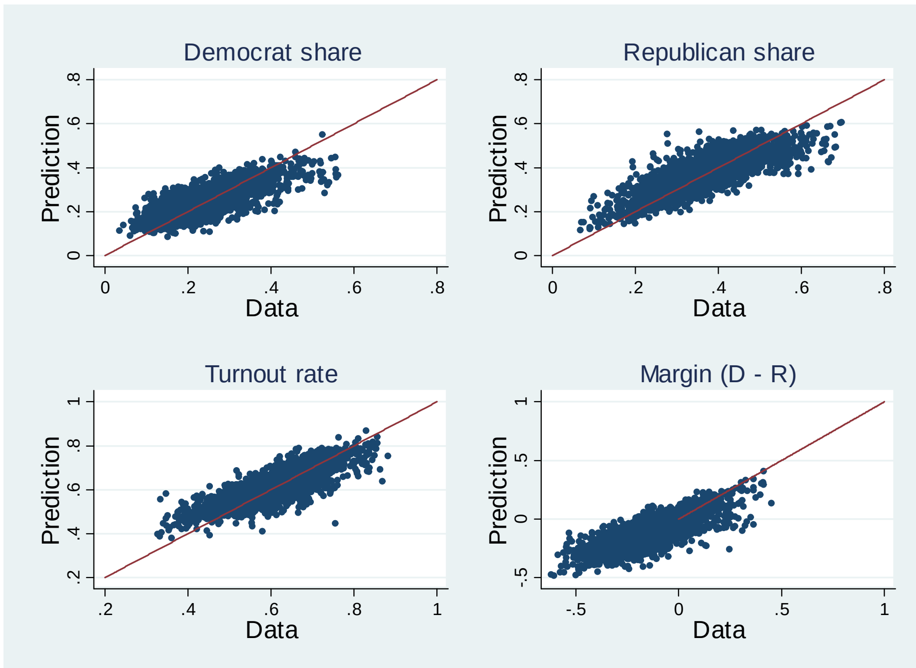

Note: The figure plots the predicted vote share, turnout, and vote share margin against the data for each county.

win elections with costly voting (see, also, Casella, 2005; Lalley and Weyl, 2015). Note that our discussion of intensity in this subsection depends on the distributional assumption of idiosyncratic preference error ε_n.

Figure 2.7 plots the histogram of $b(\mathbf{x}_n) \equiv b_s(\mathbf{x}_n) + \varepsilon_n$ for all eligible voters in the forty states included in our sample (top panel), and the proportion of those who turn out for given levels of $b(\mathbf{x}_n)$ (bottom panel). The top panel shows that the distribution of the utility difference is roughly centered around zero, and has a slightly fatter tail on the Democrats' side. The bottom panel shows that turnout is higher among Republican supporters than Democratic supporters at the same level of preference intensity. For example, voters with preferences in the bin [–2.5, –2) turns out with 70.3%, while voters with preference in the bin

[2, 2.5) turns out with 87.6%. Overall, we estimate that turnout is significantly lower among the electorate who prefer Democrats over Republicans, at 55.7%, than turnout among those who prefer Republicans over Democrats, at 64.5%.

The panel also shows that there is high turnout among voters with high preference intensity for either party. For example, voters with preferences in the bin [–2.5, –2) are almost twice as likely to turn out as those with preferences in the bin [–0.5, 0) (70.3% compared to 36.5%). This implies that voters with intense preferences effectively have "more votes" than those who are indifferent, as pointed out by Campbell (1999).

2.7. Counterfactual Experiments

We conduct two counterfactual experiments to quantify the degree to which the correlation among preferences, voting cost, and efficacy affects preference aggregation.

2.7.1. Preference Aggregation under Compulsory Voting

In our first counterfactual experiment, we consider what the election outcome would be if the preferences of all eligible voters were aggregated. The counterfactual result can be thought of as the election outcome under compulsory voting. The election outcome under this counterfactual can be computed by setting voting cost to zero. That is, individuals vote for Democrats or Republicans depending on the sign of $\widehat{b}_s(\mathbf{x}_n) + \widehat{\xi}_m + \varepsilon_n$, where $\widehat{b}_s(\cdot)$ and $\widehat{\xi}_m$ are the estimates of the net utility difference and county-level preference shock.[32] Hence, we

[32]Note that there is a unique value of $(\hat{\xi}_m, \hat{\eta}_m)$ that rationalizes the actual vote outcome given our estimates of preference, cost and perception of efficacy, as discussed in footnote 30. We use these values of $(\hat{\xi}_m, \hat{\eta}_m)$ to compute our counterfactual outcome.

Figure 2.7. Histogram of Preference Intensity and Turnout Rate by Preference Intensity

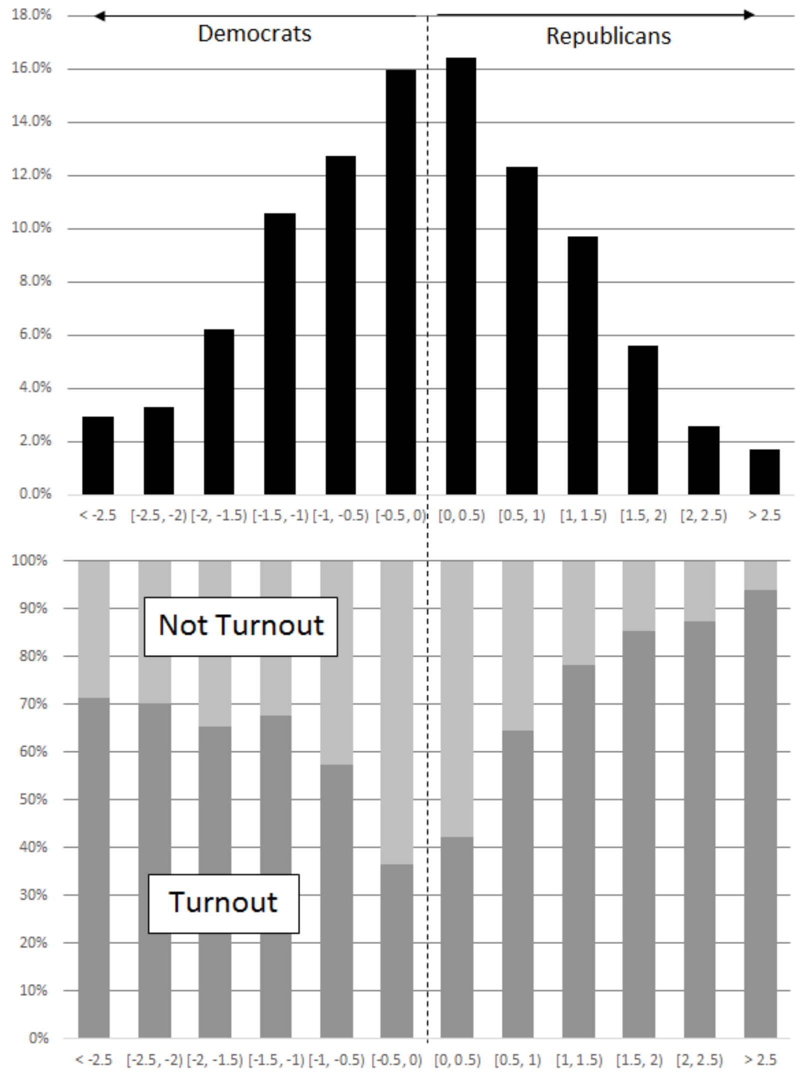

Note: The top panel plots the histogram of the utility difference, $b(\mathbf{x}_n)$. The bottom panel plots the proportion of those who turn out for given levels of preference intensity.

can express the counterfactual county-level vote shares, $\tilde{v}_{D,m}$ and $\tilde{v}_{R,m}$, as

$$\begin{aligned}\tilde{v}_{D,m} &= \int \Phi\left(-\widehat{b}_s(\mathbf{x}_n) - \widehat{\xi}_m\right) dF_{\mathbf{x},m}(\mathbf{x}_n), \\ \tilde{v}_{R,m} &= 1 - \tilde{v}_{D,m}.\end{aligned}$$

Table 2.3. Counterfactual Outcome Under Full Turnout

	Two-Party Vote Share			# of Electors	
	Democrats	Republicans	Turnout Rate	Democrats	Republicans
Actual	48.2%	51.8%	60.1%	208	278
Counterfactual	51.9%	48.1%	100.0%	310	176
	(1.1%)	(1.1%)	n.a.	(30.5)	(30.5)

Note: The table compares the actual outcome with the counterfactual outcome in which all voters turn out. The reported outcomes do not include the results for the eleven states that we drop from the sample. Standard errors are reported in parentheses.

Note that our counterfactual results are robust to equilibrium adjustments to voters' perception of efficacy because a voter's decision depends only on the sign of the utility difference irrespective of the perception of efficacy.

Table 2.3 compares the actual outcomes with the counterfactual outcomes for the forty states in our sample. The first row of Table 2.3 reports the actual vote share, turnout rate, and the number of electors for the two parties. We report our counterfactual results in the second row. We find that the Democratic two-party vote share increases from 48.2% to 51.9% in the counterfactual, reflecting our earlier finding that the preference for Democrats and voting costs are positively correlated.

In terms of electoral votes, we find that the Democrats increase the number of electoral votes by 102, from 208 to 310. Although ten states and the District of Columbia (D.C.) are not included in our sample, 310 electoral votes is larger than the threshold needed to win the election (270) even if the 10 excluded states and D.C. all vote for the Republican electors.[33] Hence, our estimates suggest that the Democrats would likely have won the 2004 presidential election if the preferences of all voters had been aggregated. The standard errors

[33] There are a total of 538 electoral votes, including the states that are excluded from our sample. A candidate needs 270 electoral votes to win.

Table 2.4. State-level Simulation Results under Full Turnout

	Two-Party Vote Share of Democrats			Turnout Rate	# of Electors for Democrats			Prob(Democrats win in counterfactual)
	Actual	Counterfactual		Actual	Actual	Counterfactual		
Alabama	37.1%	43.0%	(1.2%)	56.9%	0	0	(0.0)	0.0%
Arizona	44.7%	50.0%	(1.4%)	54.1%	0	10	(5.0)	48.8%
Arkansas	45.1%	50.2%	(1.4%)	53.0%	0	6	(3.0)	52.9%
California	55.0%	58.1%	(1.4%)	55.7%	55	55	(0.0)	100.0%
Colorado	47.6%	50.5%	(0.9%)	65.8%	0	9	(4.3)	65.3%
Florida	47.5%	51.2%	(1.1%)	63.0%	0	27	(8.6)	88.6%
Georgia	41.6%	46.7%	(1.2%)	55.9%	0	0	(0.9)	0.3%
Idaho	30.7%	36.1%	(1.2%)	63.4%	0	0	(0.0)	0.0%
Illinois	55.2%	57.9%	(1.2%)	60.7%	21	21	(0.0)	100.0%
Indiana	39.6%	45.5%	(1.3%)	54.3%	0	0	(0.0)	0.0%
Iowa	49.7%	51.9%	(0.9%)	69.7%	0	7	(0.4)	99.7%
Kansas	37.1%	42.5%	(1.2%)	61.0%	0	0	(0.0)	0.0%
Kentucky	40.0%	44.9%	(1.2%)	58.6%	0	0	(0.0)	0.0%
Louisiana	42.7%	47.3%	(1.0%)	60.9%	0	0	(0.9)	1.0%
Maryland	56.6%	59.7%	(1.1%)	61.5%	10	10	(0.0)	100.0%
Michigan	51.7%	54.4%	(1.0%)	66.7%	17	17	(0.0)	100.0%
Minnesota	51.8%	52.8%	(0.7%)	77.9%	10	10	(0.0)	100.0%
Mississippi	40.1%	45.9%	(1.2%)	55.5%	0	0	(0.0)	0.0%
Missouri	46.4%	49.5%	(0.9%)	65.3%	0	0	(4.7)	24.2%
Montana	39.7%	44.2%	(1.0%)	64.4%	0	0	(0.0)	0.0%
Nevada	48.7%	53.2%	(1.3%)	54.7%	0	5	(0.3)	99.7%
New Jersey	53.4%	56.5%	(1.2%)	61.6%	15	15	(0.0)	100.0%
New Mexico	49.6%	53.2%	(1.2%)	58.4%	0	5	(0.3)	99.7%
New York	59.3%	62.5%	(1.3%)	56.9%	31	31	(0.0)	100.0%
North Carolina	43.8%	48.3%	(1.2%)	57.1%	0	0	(3.8)	6.7%
North Dakota	36.1%	41.0%	(1.0%)	65.6%	0	0	(0.0)	0.0%
Ohio	48.9%	51.5%	(0.9%)	67.6%	0	20	(3.4)	97.0%
Oklahoma	34.4%	40.6%	(1.3%)	58.4%	0	0	(0.0)	0.0%
Oregon	52.1%	53.6%	(0.8%)	71.3%	7	7	(0.0)	100.0%
Pennsylvania	51.3%	54.0%	(1.1%)	62.4%	21	21	(0.0)	100.0%
South Carolina	41.4%	47.2%	(1.3%)	52.6%	0	0	(0.8)	1.0%
South Dakota	39.2%	43.1%	(0.8%)	68.9%	0	0	(0.0)	0.0%
Tennessee	42.8%	47.7%	(1.2%)	56.2%	0	0	(1.8)	2.7%
Texas	38.5%	45.1%	(1.5%)	52.1%	0	0	(0.0)	0.0%
Utah	26.7%	32.6%	(1.5%)	59.1%	0	0	(0.0)	0.0%
Virginia	45.9%	50.1%	(1.1%)	60.1%	0	13	(6.5)	48.5%
Washington	53.6%	55.6%	(1.0%)	66.4%	11	11	(0.0)	100.0%
West Virginia	43.5%	49.2%	(1.4%)	53.8%	0	0	(2.2)	25.3%
Wisconsin	50.2%	51.7%	(0.7%)	75.2%	10	10	(0.0)	100.0%
Wyoming	29.7%	35.2%	(1.1%)	65.3%	0	0	(0.0)	0.0%

Note: The shaded rows correspond to the states in which the winning party in the counterfactual differs from that in the actual data. The total number of electors is 538 and the number of electors for the states included in our data is 486. 270 electors are needed to win the election. Standard errors are reported in parenthesis.

in our parameter estimates translate to an 87.7% confidence level that the number of electors for the Democrats exceeds 270. Note that this number is a lower bound on the confidence level that the Democrats win the election because it assumes that all states excluded from our sample vote for the Republicans. If we assume, instead, that all states excluded from our sample vote in the same way as they did in the actual election, the confidence level that the Democrats win is 96.0%.

Table 2.4 presents the state-level breakdown of the counterfactual results for the forty states in our sample. We find that the two-party vote share of the Democrats increases in the counterfactual in all states, and that the results are overturned in nine states (shaded in the table) in the counterfactual. The table also shows that there is considerable heterogeneity in the magnitude of the change across states. For example, in Texas, we find that the change in the two-party vote share for the Democrats is more than 5 percentage points (from 38.5% to 45.1%), while, in Minnesota, the change is only 1.0 percentage point. An important variable that explains the heterogeneity is the actual turnout. Figure 2.8 plots the state-level change in the two-party vote share against turnout, and shows that the change tends to be small in states with high voter turnout, while it tends to be large in states with low voter turnout.

Figure 2.8. Changes in the Democrat's Vote Share by State

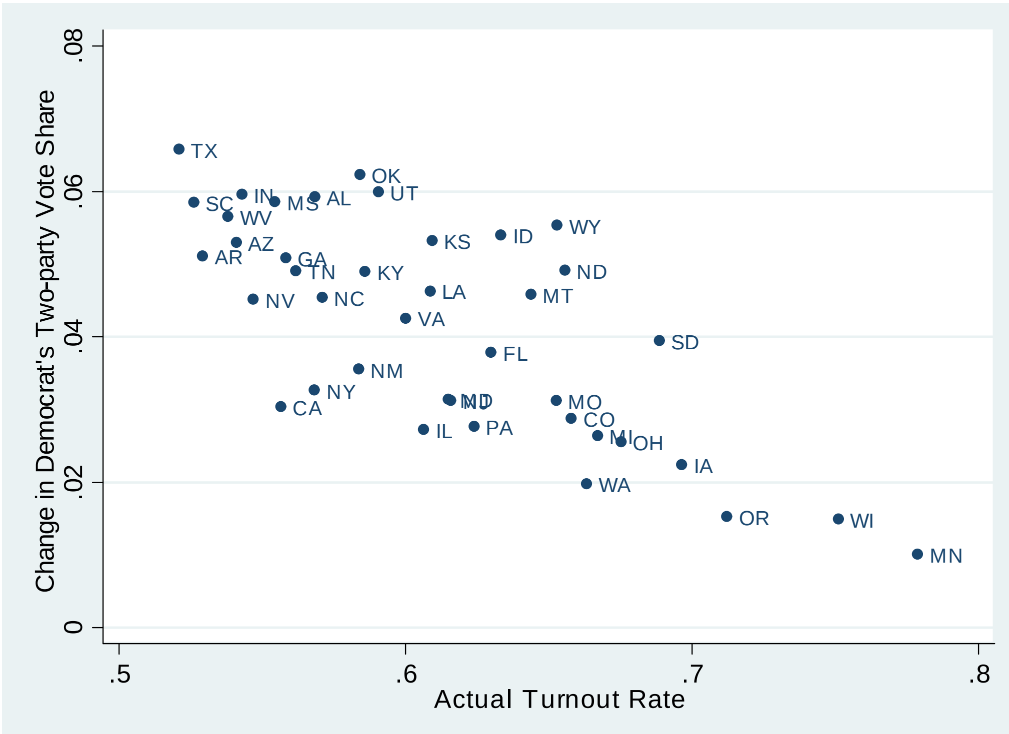

Note: The figure plots the changes in the Democrat's two-party vote share between the actual and the counterfactual against the level of actual turnout rate.

2.7.2. Efficiency Gap and Endogenous Turnout

In our second counterfactual experiment, we study the implications of endogenous turnout for using the efficiency gap as a measure of gerrymandering. Our counterfactual is motivated by a recent U.S. Supreme Court case involving districting for Wisconsin's state legislature (*Gill v. Whitford*) in which the plaintiffs introduces a metric called the "efficiency gap" to measure the extent of gerrymandering.

The efficiency gap, proposed by Stephanopoulos and McGhee (2015) is a measure of how well vote shares map into seat shares, and it is defined as the difference in the wasted votes between the parties divided by the total votes. The top half of Table 2.5 is a numerical example from Stephanopoulos and McGhee (2015) that illustrates how the efficiency gap is computed. The wasted vote for party k in district d is either $v_{kd} - (1/2)(v_{kd} + v_{-kd})$ or v_{kd}, depending on whether v_{kd} is greater than v_{-kd}, where v_{kd} is the votes obtained by party k in district d and v_{-kd} is the votes obtained by the other party. The efficiency gap is simply the sum of the difference in the wasted votes across districts, divided by the total votes. Stephanopoulos and McGhee (2015) argue that an efficiency gap exceeding 8% should be presumptively unlawful. Table 2.5 is an example of a districting plan that favors party A (party A wins districts 1 through 8 with a state-wide vote share of 55%), and indeed, this is captured by the fact that this has a high efficiency gap of 20%.

While the efficiency gap captures the imbalance in the way votes are translated to seat shares, the validity of the efficiency gap as a measure of gerrymandering requires that turnout is exogenous. To see this, consider again the example in Table 2.5, but suppose now that half of the voters are low cost types, who always vote, and the other half are high cost types, who vote only when the partisanship of the district is relatively balanced. The bottom part of Table 2.5 illustrates a redistricting plan in which the shares of A and B supporters are kept the same as in the original Stephanopoulos and McGhee (2015)'s example, but voters now have heterogeneous costs. Suppose that the high cost voters turn out only in districts 4 to 8, which are the districts with a balanced partisanship. This would lower the efficiency gap from 20% to 4.87% while keeping the seat shares unchanged (i.e., party A wins districts 1 through 8).

Table 2.5. Calculation of the Efficiency Gap

District	1	2	3	4	5	6	7	8	9	10	Total	
	Stephanopolous and McGhee (2015)'s example											
Total Population	100	100	100	100	100	100	100	100	100	100	1000	
A Supporter	70	70	70	54	54	54	54	54	35	35	550	
B Supporter	30	30	30	46	46	46	46	46	65	65	450	
A Wasted Vote	20	20	20	4	4	4	4	4	35	35	150	Efficiency Gap
B Wasted Vote	30	30	30	46	46	46	46	46	15	15	350	20.00%
	Endogenous Turnout											
Low Cost A Supporter	70	70	70	0	0	0	0	0	33	32	275	
High Cost A Supporter	0	0	0	54	54	54	54	54	2	3	275	
Low Cost B Supproter	0	0	0	46	43	23	23	23	34	33	225	
High Cost B Supporter	30	30	30	0	3	23	23	23	31	32	225	
A Vote	70	70	70	54	54	54	54	54	33	32	545	
B Vote	0	0	0	46	46	46	46	46	34	33	297	
Turnout	70	70	70	100	100	100	100	100	67	65	842	
A Wasted Vote	35	35	35	4	4	4	4	4	33	32	190	Efficiency Gap
B Wasted Vote	0	0	0	46	46	46	46	46	0.5	0.5	231	4.87%

Note: This example is taken from Figure 1 of Stephanopolus and McGhee (2015). The top half is the original example. The bottom half is our modification with two cost types and endogenous turnout.

This example illustrates the potential issue with using the efficiency gap as a measure of gerrymandering. The original example of Stephanopoulos and McGhee (2015) and the modified example are just as biased for party A in terms of mapping the overall preferences of the voters to seat shares, and yet they result in very different efficiency gap measures. Moreover, districting planners can take advantage of the heterogeneity in voting costs to draw plans that give one party a disproportionate advantage while keeping the efficiency gap low. Whether the efficiency gap is a good measure of gerrymandering depends on the extent to which turnout levels can be manipulated through districting plans.

In order to empirically evaluate the robustness of the efficiency gap to endogenous turnout, we compute the efficiency gap for the 2004 U.S. Presidential election when we equalize the state-specific component of efficacy across states (set $p_s = p_{s'}$). To the extent that variation in p_s across states reflects how state boundaries affect turnout, equalizing p_s across states can be interpreted as eliminating endogeneity in turnout that is state specific. Although the intended use of the efficiency gap measure is mainly for Congressional and state legislative elections, it is possible to compute the efficiency gap for presidential elections as well. In our context, we compute the wasted vote in each state for the two parties and then sum the difference across all of the states.

Table 2.6 reports the results. The first row corresponds to the efficiency gap computed using the actual data. Rows 2 through 7 corresponds to the results when we equalize p across states. We consider six different levels of p to target aggregate turnout levels between 40% and 80% in 10% increments and the actual aggregate turnout level (60.1%). Comparing the first row to the 5th row (60.1%), we find that the efficiency gap changes by 2.2%. This implies that equalizing p_s across states in a way that does not change the overall turnout affects the efficiency gap by 2.2%. The table also shows that the efficiency gap changes by about 14% when turnout is exogenously increased from 50% to 80%. These changes in the efficiency measure seem significant especially in light of the fact that Stephanopoulos and McGhee (2015) argue that districting plans exceeding the threshold value of 8% should be deemed presumptively unlawful.

Table 2.6. Election Outcomes and Efficiency Gap when p_s is Equalized across States

	Turnout	Vote Share		Two-party Vote Share	# of Winning States		Efficiency Gap
		D	R	for D	D	R	
Actual	60.1%	29.0%	31.1%	48.2%	11	29	-0.47%
Counterfactual	40%	18.1%	21.9%	45.3%	6	34	-9.35%
	50%	23.4%	26.6%	46.8%	7	33	-7.21%
	60%	28.9%	31.1%	48.2%	10	30	-2.61%
	60.1%	29.0%	31.1%	48.2%	10	30	-2.64%
	70%	34.9%	35.1%	49.8%	13	27	-3.38%
	80%	40.9%	39.1%	51.1%	16	24	6.89%

Note: Efficiency gap is defined for Democrat candidate. Positive (negative) efficiency gap means that the districting plan gives Democrats (Republicans) an advantage.

The sensitivity of the efficiency gap to endogeneity of turnout points to a more general problem with comparing actual votes and seat shares as a measure of gerrymandering. While our discussion has so far focused on the efficiency gap, any measure based on a comparison between actual votes and seat shares is subject to the same concerns. This includes the concept of partisan symmetry, which a majority of the Supreme court justices expressed support in the case *League of United Latin American Citizens v Perry*. The extant arguments seem to take turnout as fixed and exogenous. Our results illustrate the importance of considering the implications of endogenous turnout when thinking about how to measure gerrymandering. One natural alternative is to compare the difference between the actual underlying voter preferences and seat shares. If we think about elections as a way to aggregate preferences into outcomes, evaluating the electoral system in terms of its ability to aggregate preferences seems most coherent. The methods developed in this paper can be used for that purpose.

References

Abito, Jose Miguel, "Agency Costs in Environmental Regulation: Evidence from Regulated Electric Utilities," *Working Paper*, 2014.

Aguirregabiria, Victor and Pedro Mira, "Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models," *Econometrica*, 2002, *70* (4), 1519–1543.

——— **and** ———, "Sequential estimation of dynamic discrete games," *Econometrica*, 2007, *75* (1), 1–53.

——— **and** ———, "Dynamic discrete choice structural models: A survey," *Journal of Econometrics*, 2010, *156* (1), 38–67.

Akerlof, George A and William T Dickens, "The Economic Consequences of Cognitive Dissonance," *American Economic Review*, 1982, pp. 307–319.

Atkeson, Lonna Rae, ""Sure, I Voted for the Winner!" Overreport of the Primary Vote for the Party Nominee in the National Election Studies," *Political Behavior*, 1999, *21* (3), 197–215.

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

Barkume, Anthony J, "Identification of Preference for Election Outcomes from Aggregate Voting Data," *Public Choice*, 1976, *27* (1), 41–58.

Bendor, Jonathan, Daniel Diermeier, and Michael Ting, "A Behavioral Model of

Turnout," *American Political Science Review*, 2003, *97* (02), 261–280.

Berry, Steven T, "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 1994, *25* (2), 242–262.

Breen, William J, Laurie Simon Hodrick, and Robert A Korajczyk, "Predicting equity liquidity," *Management Science*, 2002, *48* (4), 470–483.

Brunell, Thomas L and John DiNardo, "A Propensity Score Reweighting Approach to Estimating the Partisan Effects of Full Turnout in American Presidential Elections," *Political Analysis*, 2004, *12* (1), 28–45.

Brunnermeier, Markus K and Jonathan A Parker, "Optimal Expectations," *American Economic Review*, 2005, *95* (4).

Campbell, Colin M, "Large Electorates and Decisive Minorities," *Journal of Political Economy*, 1999, *107* (6), 1199–1217.

Cantillon, Estelle and Aurelie Slechten, "Price formation in the European carbon market: the role of firm participation and market structure," *Working Paper*, 2015.

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L Palmer, "Sulfur dioxide control by electric utilities: What are the gains from trade?," *Journal of political Economy*, 2000, *108* (6), 1292–1326.

Cascio, Elizabeth U and Ebonya Washington, "Valuing the Vote: The Redistribution of Voting Rights and State Funds Following the Voting Rights Act of 1965," *Quarterly Journal of Economics*, 2013.

Casella, Alessandra, "Storable Votes," *Games and Economic Behavior*, 2005, *51* (2), 391–419.

Chan, H Ron, B Andrew Chupp, Maureen L Cropper, and Nicholas Z Muller, "The Impact of Trading on the Costs and Benefits of the Acid Rain Program," Technical

Report, National Bureau of Economic Research 2015.

Chan, Ron, "How Large are the Cost Savings from Emissions Trading? An Evaluation of the U.S. Acid Rain Program," *Working Paper*, 2015.

Citrin, Jack, Eric Schickler, and John Sides, "What if Everyone Voted? Simulating the Impact of Increased Turnout in Senate Elections," *American Journal of Political Science*, 2003, *47* (1), 75–90.

Coase, Ronald H, *The problem of social cost*, Springer, 1960.

Coate, Stephen and Brian Knight, "Socially optimal districting: a theoretical and empirical exploration," *The Quarterly Journal of Economics*, 2007, *122* (4), 1409–1471.

______ **and Michael Conlin**, "A Group Rule-Utilitarian Approach to Voter Turnout: Theory and Evidence," *American Economic Review*, 2004, pp. 1476–1504.

______, ______, **and Andrea Moro**, "The Performance of Pivotal-Voter Models in Small-Scale Elections: Evidence from Texas Liquor Referenda," *Journal of Public Economics*, 2008, *92* (3), 582–596.

Collard-Wexler, Allan, "Demand Fluctuations in the Ready-Mix Concrete Industry," *Econometrica*, 2013, *81* (3), 1003–1037.

Cullen, Joseph A, "Dynamic response to environmental regulation in the electricity industry," *Working Paper*, 2015.

______ **and Stanley S Reynolds**, "Market Dynamics and Investment in the Electricity Sector," 2017.

Dardati, Evangelina, "Pollution Permit Systems and Firm Dynamics: How does the Allocation Scheme Matter?," *Working Paper*, 2014.

Dávila, Eduardo and Cecilia Parlatore, "Trading Cost and Informational Efficiency," 2017.

Degan, Arianna and Antonio Merlo, "A Structural Model of Turnout and Voting in Multiple Elections," *Journal of the European Economic Association*, 2011, *9* (2), 209–245.

DellaVigna, Stefano, John A List, Ulrike Malmendier, and Gautam Rao, "Voting to Tell Others," *Working Paper*, 2015.

DeNardo, James, "Turnout and the Vote: The Joke's on the Democrats," *American Political Science Review*, 1980, *74* (02), 406–420.

Dhillon, Amrita and Susana Peralta, "Economic Theories of Voter Turnout," *Economic Journal*, 2002, *112* (480), 332–352.

Downs, Anthony, *An Economic Theory of Democracy*, New York, 1957.

Duffy, John and Margit Tavits, "Beliefs and Voting Decisions: A Test of the Pivotal Voter Model," *American Journal of Political Science*, 2008, *52* (3), 603–618.

EIA, "Electric Power Annual 2012," Technical Report 2012.

Ellerman, A Denny and Juan-Pablo Montero, "The efficiency and robustness of allowance banking in the US Acid Rain Program," *The Energy Journal*, 2007, pp. 47–71.

______, Paul L Joskow, Richard Schmalensee, Juan-Pablo Montero, and Elizabeth M Bailey, *Markets for clean air: The US Acid Rain Program*, Cambridge University Press, 2000.

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

Esponda, Ignacio and Demian Pouzo, "Conditional Retrospective Voting in Large Elections," *American Economic Journal: Microeconomics, Forthcoming*, 2016.

Fabra, Natalia and Mar Reguant, "Pass-Through of Emissions Costs in Electricity Markets," *American Economic Review*, 2014, *104* (9), 2872–99.

Feddersen, Timothy, "Rational Choice Theory and the Paradox of Not Voting," *Journal* *of Economic Perspectives*, 2004, pp. 99–112.

——— **and Alvaro Sandroni**, "A Theory of Participation in Elections," *American Economic Review*, 2006, pp. 1271–1282.

——— **and Wolfgang Pesendorfer**, "Voting behavior and information aggregation in elections with private information," *Econometrica*, 1997, pp. 1029–1058.

Feddersen, Timothy J and Wolfgang Pesendorfer, "The swing voter's curse," *American economic review*, 1996, pp. 408–424.

Fowlie, Meredith, "Allocating emissions permits in cap-and-trade programs: Theory and evidence," *Working Paper*, 2010.

———, "Emissions trading, electricity restructing, and investment in pollution abatement," *The American Economic Review*, 2010, pp. 837–869.

——— **and Jeffrey M Perloff**, "Distributing pollution rights in cap-and-trade programs: are outcomes independent of allocation?," *Review of Economics and Statistics*, 2013, *95* (5), 1640–1652.

——— **and Nicholas Muller**, "Market-based emissions regulation when damages vary across sources: What are the gains from differentiation?," Technical Report, Working Paper 2013.

———**, Mar Reguant, and Stephen P Ryan**, "Market-based emissions regulation and industry dynamics," *forthcoming in Journal of Political Economy*, 2014.

———**, Stephen P Holland, and Erin T Mansur**, "What do emissions markets deliver and to whom? Evidence from Southern California's NOx trading program," *The American economic review*, 2012, *102* (2), 965–993.

Friedman, John N and Richard T Holden, "Optimal gerrymandering: sometimes pack, but never crack," *American Economic Review*, 2008, *98* (1), 113–44.

Fujiwara, Thomas, "Voting Technology, Political Responsiveness, and Infant Health: Evidence from Brazil," *Econometrica*, 2015, *83* (2), 423–464.

Gârleanu, Nicolae and Lasse Heje Pedersen, "Dynamic trading with predictable returns and transaction costs," *The Journal of Finance*, 2013, *68* (6), 2309–2340.

Gordon, Brett R and Wesley R Hartmann, "Advertising Effects in Presidential Elections," *Marketing Science*, 2013, *32* (1), 19–35.

Grofman, Bernard and Gary King, "The future of partisan symmetry as a judicial test for partisan gerrymandering after LULAC v. Perry," *Election Law Journal*, 2007, *6* (1), 2–35.

Haile, Philip A and Elie Tamer, "Inference with an Incomplete Model of English Auctions," *Journal of Political Economy*, 2003, *111* (1), 1–51.

Hansford, Thomas G and Brad T Gomez, "Estimating the Electoral Effects of Voter Turnout," *American Political Science Review*, 2010, *104* (02), 268–288.

Harsanyi, John C, "Rule Utilitarianism, Rights, Obligations and the Theory of Rational Behavior," *Theory and Decision*, 1980, *12* (2), 115–133.

Helfand, Gloria E, Michael R Moore, and Yimin Liu, "Testing for dynamic efficiency of the Sulfur Dioxide Allowance Market," 2006.

Holland, Stephen P, Erin T Mansur, Nicholas Z Muller, and Andrew J Yates, "Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors," *American Economic Review*, 2016, *106* (12), 3700–3729.

Hopenhayn, H. A., "Industry Equilibrium Dynamics: A General Competitive Theory," March 1990.

Hopenhayn, Hugo A, "Entry, exit, and firm dynamics in long run equilibrium," *Econometrica: Journal of the Econometric Society*, 1992, pp. 1127–1150.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Husted, Thomas A and Lawrence W Kenny, "The Effect of the Expansion of the Voting Franchise on the Size of Government," *Journal of Political Economy*, 1997, pp. 54–82.

Jäckel, Peter, "A Note on Multivariate Gauss-Hermite Quadrature," Technical Report 2005. Sourced from: www.pjaeckel.webspace. virginmedia.com/ANoteOnMultivariateGaussHermiteQuadrature.pdf.

Jaraitė-Kažukauskė, Jūratė and Andrius Kažukauskas, "Do transaction costs influence firm trading behaviour in the european emissions trading system?," *Environmental and Resource Economics*, 2015, *62* (3), 583–613.

Joskow, Paul L and Richard Schmalensee, "The Political Economy of Market-Based Environmental Policy: the US Acid Rain Program 1," *The journal of law and economics*, 1998, *41* (1), 37–84.

______, ______, and Elizabeth M Bailey, "The market for sulfur dioxide emissions," *American Economic Review*, 1998, pp. 669–685.

Jovanovic, Boyan, "Selection and the Evolution of Industry," *Econometrica: Journal of the Econometric Society*, 1982, pp. 649–670.

Kalouptsidi, Myrto, "Time to build and fluctuations in bulk shipping," *The American Economic Review*, 2014, *104* (2), 564–608.

Kanazawa, Satoshi, "A Possible Solution to the Paradox of Voter Turnout," *Journal of Politics*, 1998, *60* (04), 974–995.

Karp, Jeffrey A and Susan A Banducci, "Political Efficacy and Participation in Twenty-Seven Democracies: How Electoral Systems Shape Political Behaviour," *British Journal* *of Political Science*, 2008, *38* (02), 311–334.

Kellogg, Ryan, "The effect of uncertainty on investment: evidence from Texas oil drilling," *The American Economic Review*, 2014, *104* (6), 1698–1734.

Keohane, Nathaniel, "Cost savings from allowance trading in the 1990 Clean Air Act: Estimates from a choice-based model," *Moving to markets in environmental regulation: Lessons from twenty years of experience*, 2006, pp. 194–229.

Kolstad, Jonathan T and Frank A Wolak, "Using Environmental Emissions Permit Prices to Raise Electricity Prices: Evidence from the California Electricity Market," *Working Paper*, 2008.

Kroth, Verena, Valentino Larcinese, and Joachim Wehner, "A Better Life For All? Democratization and Electrification in Post-Apartheid South Africa," Technical Report 2013.

Krusell, Per and Anthony A Smith Jr, "Income and wealth heterogeneity in the macroeconomy," *Journal of political Economy*, 1998, *106* (5), 867–896.

Lalley, P Steven and E Glen Weyl, "Quadratic Voting," *Working Paper*, 2015.

Ledyard, John O, "The Pure Theory of Large Two-Candidate Elections," *Public Choice*, 1984, *44* (1), 7–41.

Lee, Donghoon and Kenneth I Wolpin, "Intersectoral labor mobility and the growth of the service sector," *Econometrica*, 2006, *74* (1), 1–46.

Leighley, Jan E and Jonathan Nagler, *Who Votes Now?: Demographics, Issues, Inequality, and Turnout in the United States*, Princeton University Press, 2013.

Levine, David K and Thomas R Palfrey, "The Paradox of Voter Participation? A Laboratory Study," *American Political Science Review*, 2007, *101* (01), 143–158.

Lijphart, Arend, "Unequal Participation: Democracy's Unresolved Dilemma Presidential

Address," *American Political Science Review*, 1997, *91* (01), 1–14.

Lillo, Fabrizio, J Doyne Farmer, and Rosario N Mantegna, "Econophysics: Master curve for price-impact function," *Nature*, 2003, *421* (6919), 129–130.

Liski, Matti and Juan-Pablo Montero, "Market Power in an Exhaustible Resource Market: The Case of Storable Pollution Permits," *The Economic Journal*, 2011, *121* (551), 116–144.

Martinez, Michael D and Jeff Gill, "The Effects of Turnout on Partisan Outcomes in US Presidential Elections 1960–2000," *Journal of Politics*, 2005, *67* (4), 1248–1274.

Matsusaka, John G, "Explaining voter turnout patterns: An information theory," *Public choice*, 1995, *84* (1), 91–117.

McDonald, Michael P., "1980-2014 November General Election," 2016. accessed May 26th, 2016.

McDonald, Michael P and Samuel L Popkin, "The Myth of the Vanishing Voter," *American Political Science Review*, 2001, *95* (4), 963–974.

Merlo, A and Aureo de Paula, "Identification and Estimation of Preference Distributions When Voters Are Ideological," *Working Paper*, 2016.

Merlo, Antonio, "Whither Political Economy? Theories, Facts and Issues," in Whitney Newey Richard Blundell and Torsten Persson, eds., *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society*, Vol. I, Cambridge: Cambridge University Press, 2006, pp. 381–421.

Miller, Grant, "Women's Suffrage, Political Responsiveness, and Child Survival in American History," *The Quarterly Journal of Economics*, 2008, *123* (3), 1287.

Minozzi, William, "Endogenous Beliefs in Models of Politics," *American Journal of Political Science*, 2013, *57* (3), 566–581.

Muller, Nicholas Z and Robert Mendelsohn, "Efficient pollution regulation: getting the prices right," *The American Economic Review*, 2009, *99* (5), 1714–1739.

Myatt, David, "A Rational Choice Theory of Voter Turnout," Technical Report 2012.

Myerson, Roger B and Robert J Weber, "A Theory of Voting Equilibria.," *American Political Science Review*, 1993, *87* (01), 102–114.

Nagle, John F, "Measures of partisan bias for legislating fair elections," *Election Law Journal*, 2015, *14* (4), 346–360.

Palfrey, Thomas R and Howard Rosenthal, "A Strategic Calculus of Voting," *Public Choice*, 1983, *41* (1), 7–53.

——— **and** ———, "Voter Participation and Strategic Uncertainty," *American Political Science Review*, 1985, *79* (01), 62–78.

Reguant, Mar and A Denny Ellerman, "Grandfathering and the endowment effect: An Assessment in the context of the Spanish National Allocation Plan," *Working Paper*, 2008.

Richards-Shubik, Seth, "Peer effects in sexual initiation: Separating demand and supply mechanisms," *Quantitative Economics*, 2015, *6* (3), 663–702.

Riker, William H and Peter C Ordeshook, "A Theory of the Calculus of Voting," *American Political Science Review*, 1968, *62* (01), 25–42.

Robert, Engle, Ferstenberg Robert, and Russell Jeffrey, "Measuring and modeling execution cost and risk," *The Journal of Portfolio Management*, 2012, *38* (2), 14–28.

Rubin, Jonathan D, "A model of intertemporal emission trading, banking, and borrowing," *Journal of Environmental Economics and Management*, 1996, *31* (3), 269–286.

Rust, John, "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica: Journal of the Econometric Society*, 1987, pp. 999–1033.

———, "Structural estimation of Markov decision processes," *Handbook of econometrics*, 1994, *4*, 3081–3143.

Ryan, Stephen P, "The costs of environmental regulation in a concentrated industry," *Econometrica*, 2012, *80* (3), 1019–1061.

Schennach, Susanne M, "The economics of pollution permit banking in the context of Title IV of the 1990 Clean Air Act Amendments," *Journal of Environmental Economics and Management*, 2000, *40* (3), 189–210.

Schmalensee, Richard and Robert N Stavins, "The SO2 Allowance Trading System: The Ironic History of a Grand Policy Experiment," *The Journal of Economic Perspectives*, 2013, *27* (1), 103–121.

Shachar, Ron and Barry Nalebuff, "Follow the Leader: Theory and Evidence on Political Participation," *American Economic Review*, 1999, pp. 525–547.

Smith, Mark A, "The Contingent Effects of Ballot Initiatives and Candidate Races on Turnout," *American Journal of Political Science*, 2001, pp. 700–706.

Stavins, Robert N, "Transaction costs and tradeable permits," *Journal of environmental economics and management*, 1995, *29* (2), 133–148.

Stephanopoulos, Nicholas O and Eric M McGhee, "Partisan gerrymandering and the efficiency gap," *The University of Chicago Law Review*, 2015, pp. 831–900.

Strömberg, David, "How the Electoral College Influences Campaigns and Policy: the Probability of Being Florida," *American Economic Review*, 2008, *98* (3), 769–807.

Tucker, Harvey J and James DeNardo, "Does Heavy Turnout Help Democrats in Presidential Elections?," *American Political Science Review*, 1986, *80* (04), 1291–1304.

U.S. Environmental Protection Agency, "The National Allowance Data Base Version 2.1 Technical Support Document," 1993.

———, "Technical Documentation for Phase II Allowance Allocations," Technical Report, EPA-430-R-93-002 1993.

Wolfinger, Raymond E and Steven J Rosenstone, *Who votes?*, Vol. 22, Yale University Press, 1980.

APPENDIX A

Appendix for Chapter 1

A.1. Computational Details on Solving the Structural Model

Appendix A.1 explains the details on the computational procedure of solving the structural model.

A.1.1. Individual Optimization

I first explain the computational procedure for solving an individual problem. For notational simplicity, I omit the script i for a particular firm. Because the model has a finite period, it can be solved by backward induction.

(1) Phase II (2003 to 2000): I solve the optimization problem from 2003 to 2000. Note that I use $CV_{T+1}(h_{T+1}, R^2)$ as a continuation value in the terminal period 2003. By solving in a backward way, I obtain the policy function $\hat{x}_t(h_t, I_t, R^2)$ for emissions e_t, net purchase b_t, and banking h_{t+1}, and the expected value function in 2000 $EV_{2000}(h_{2000}, I_{2000}, R^2)$.

(2) Investment decision for Phase II: I define the continuation value at the timing of making the investment decision for Phase II by $W_{2000}(h_{2000}, I_{2000}, R^2)$. The decision problem is given by

$$W_{2000}(h_{2000}, I_{2000}, R^1) \equiv \max_{R^2} \quad EV_{2000}(h_{2000}, I_{2000}, R^2) - \Gamma(R^2, R^1).$$

$$\text{s.t.} \quad R^2 \leq R^1$$

By solving this problem, I obtain the investment policy function $R^{2*}(h_{2000}, I_{2000}, R^1)$.

(3) Phase I (1999 to 1995): I repeat the same procedure as step 1. Note that the continuation value in the problem at $t = 1999$ is given by $W_{2000}(h_{2000}, I_{2000}, R^1)$.

(4) Investment for Phase I: The problem is given by

$$\max_{R^1} EV_{1995}(0, 0, R_{P1}) - \Gamma(R^1, R^0).$$
$$\text{s.t.} R^1 \leq R^0$$

Note that $h_{1995} = 0$ and $I_{1995} = 0$ in 1995.

A.1.2. Computation of a Dynamic Competitive Equilibrium

The computational procedure for finding an equilibrium is parallel to the estimation procedure that I introduced in section 1.5.

(1) Fix a candidate of permit prices $\mathbf{P} = \{P_t\}_{t=1995}^{2003}$.

(2) Solve the individual problem by backward induction and obtain the policy function $\hat{x}_{it}(h_{it}, I_{it}, R_{it})$ for emissions e_t, net purchase b_t, and banking h_{t+1}, participation probability $P_{it}(h_{it}, R_{it})$, and the investment decisions $R_i^1(h_{i1995}, I_{i1995})$ and , $R_i^2(h_{i,2000}, I_{i,2000}, R_i^1)$.

(3) Consider the timing of market participation. Denote the year of participation by $s \in \{\emptyset, 1995, \cdots, 2003\}$. $s = \emptyset$ means that a firm does not trade at all in the period.

(4) For each path of participation timing, I simulate the optimal decisions using the policy functions.

(5) Calculate the probability that each path of participation timing is realized.

(6) The simulated optimal decisions are given as

$$\hat{x}_{it} = \sum_{s \in \{\emptyset, 1995, \cdots, 2003\}} Prob_i^{enter}(s) \hat{x}_{it}(s).$$

(7) Check the market clearing condition as

$$\sum_i \hat{b}_{it}(\mathbf{P}) + \bar{B}_t^{fringe}(P_t) = 0 \ \forall t = 1995, \cdots, 2003.$$

(8) Repeat steps 1-7 until the market clearing conditions are satisfied.

In practice, I stop the iteration when the following condition is satisfied:

$$\max_{t=1995,\cdots,2003} \left| \sum_i \hat{b}_{it}(\mathbf{P}) + \bar{B}_t^{fringe}(P_t) \right| < 1000.$$

This criterion is sufficiently tight so that the absolute value of the price change is in the order of magnitude of 1e-1.

To update the price in the above procedure, I construct the following rule that exploits the market clearing conditions and the optimality conditions. Denote a current candidate of an equilibrium price by $\mathbf{P}^l = \{P_t^l\}_{t=1995}^{2003}$. The next candidate of price P_t^{l+1} is given by solving the equation

$$\sum_i \sum_s P_{i,enter}(s) \cdot TC'^{(-1)} \left(\hat{\lambda}_{it}(\mathbf{P}^l, s) - P_t^{l+1} \right) + \bar{B}_t^{fringe}(P_t^{l+1}) = 0,$$

where $\hat{\lambda}_{it}(\mathbf{P}^l, s)$ is the prediction of the shadow values under the current candidate of prices $\mathbf{P}^l$.

A.2. Special Cases of Structural Model

Appendix A.2 introduces the special cases of the structural model I introduced in the paper. Those cases are used in counterfactual simulations.

A.2.1. Case without Permit Banking and with Transaction Costs

I explain the case in which firms are not allowed to bank emissions permits. Once I shut down permit banking, permit holding h_{it} is no longer a state variable in the model. However, the dynamic consideration still plays a role due to abatement investment and participation decisions.

I first consider individual optimization problems. Consider the case in which a firm is a trader. The problem is given by

$$
\begin{aligned}
V_{it}^{1}(I_{it}, R_{it}) = \quad & \max_{\{q_{jt}\}_j, b_t} \quad \pi_{it}\left(\{q_{jt}\}_j\right) - \left(P_t b_{it} + TC(b_{it})\right) + \beta V_{i,t+1}(1, R_{i,t+1}) \\
& \text{s.t.} \quad e_{it}\left(\{q_{jt}, \rho_{jt}\}_j\right) = a_{it} + b_{it}.
\end{aligned}
$$

Note that the choice of $\{q_{jt}\}_j$ and b_t does not affect the continuation value. The optimality conditions of the problem are given by equation (1.3.6) and (1.3.7).

Next, consider the case in which a firm is a non-trader:

$$
\begin{aligned}
V_{it}^{0}(I_{it}, R_{it}) = \quad & \max_{\{q_{jt}\}_j, b_t} \quad \pi_{it}\left(\{q_{jt}\}_j\right) + \beta V_{i,t+1}(0, R_{i,t+1}) \\
& \text{s.t.} \quad e_{it}\left(\{q_{jt}, \rho_{jt}\}_j\right) \leq a_{it}.
\end{aligned}
$$

In this case, a firm may not use all the permits, due to the capacity constraints of production. The emissions level is given by

$$e^*_{it} = \min\{a_{it}, e^{max}_{it}\},$$

where e^{max}_{it} is the emissions level when a firm is facing zero shadow costs of permits.

Other components, including the participation and the investment decisions are the same as in the baseline case (i.e., the case that includes both permit banking and transaction costs).

A.2.2. Shutting Down Transaction Costs without Permit Banking

This section explains the case in which I shut down both transaction costs and permit banking. In this case, I do not need to consider the participation decision.

Given an emissions rate R_{it}, the individual problem in period t is given by

$$\begin{aligned} \Pi_{it}(R_{it}) = \quad & \max_{\{q_{jt}\}_{j\in J_{it}}, b_{it}} \quad \pi_{it}\left(\{q_{jt}\}_j\right) - P_t b_{it} \\ & \text{s.t.} \qquad e_{it}\left(\{q_{jt}, \rho_{jt}\}_j\right) \leq a_{it} + b_{it}. \end{aligned}$$

The FOC is by $\partial \pi_{it}/\partial q_{jt} = P_t \ \forall j$. This gives the optimal choice for $e_{it}(R_{it}, P_t)$ and $b_{it}(R_{it}, P_t)$.

The investment decision for Phase II is then given as

$$\begin{aligned} W_{i,2000}(R^1_i) = \quad & \max\nolimits_{R^2_i} \quad \sum_{t=2000}^{2003} \beta^{t-2000}\Pi_{it}(R^1_i) - \Gamma(R^2_i, R^1_i) \\ & s.t. \qquad R^2_i \leq R^1_i \end{aligned}$$

and the problem for Phase I is

$$\begin{aligned} \max_{R_i^1} \quad & \sum_{t=1995}^{1999} \beta^{t-1995}\Pi_{it}(R_i^1) + \beta^{2000-1995}W_{i,2000}(R_i^1) - \Gamma(R_i^1, R_i^0). \\ s.t. \quad & R_i^1 \leq R_i^0. \end{aligned}$$

To close the model, consider an equilibrium of the permit market. The permit price should satisfy the market clearing conditions:

$$\sum_i b_{it}^*(R_{it}^*(P), P_t) + \bar{B}_t^{fringe}(P_t) = 0 \ \forall t = 1995, \cdots, 2003.$$

A.2.3. Shutting Down Transaction Costs with Permit Banking

I now consider the case with permit banking. In the absence of transaction costs, Rubin (1996) has shown that the equilibrium path of permit prices grows at the rate of β^{-1} as long as the aggregate banking is positive, which implies that

$$\begin{aligned} P_{t+1} = \beta^{-1}P_t. \\ \Longleftrightarrow P_t = & \beta^{-(t-1)}P_{1995}. \end{aligned}$$

The optimal decision on emissions, given the emissions rate, is determined by $\partial\pi_{it}/\partial q_{jt} = P_t \ \forall j$, which is the same as the one in appendix A.2.2. As I discussed in section 1.3.8.1, individual decisions on net purchase and banking are not determined from the model, because the current shadow value $\lambda_t = P_t$ is equal to the discounted marginal value of banking $\beta\lambda_{t+1} = \beta P_{t+1} = P_t$. In other words, banking and trading decisions are arbitrary as long as a firm can produce the level of emissions determined by the optimality condition.

Now I consider investment decisions. The continuation value at the beginning of Phase II is given by

$$
\begin{aligned}
V_{i,2000}(h_{i,2000}, R_i^2) &= \sum_{t=2000}^{2003} \beta^{t-2000} \left[\pi_{it}\left(\{q_{jt}\}_j, R_i^2\right) - P_t b_{it}\right] + \beta^{2003-2000} CV(h_{i,T+1}) \\
&= \sum_{t=2000}^{2003} \beta^{t-2000} \left[\pi_{it}\left(\{q_{jt}\}_j, R_i^2\right) - P_t \cdot (e_{it} - a_{it})\right] \\
&\quad + \beta^{2003-2000} \left\{CV(h_{i,T+1}) - P_T h_{i,T+1}\right\} \\
&\quad + \sum_{t=2000}^{2003} \beta^{t-2000} P_t h_{it} + \sum_{t=2000}^{2002} \beta^{t-2000} P_t h_{it+1} \\
&= \sum_{t=2000}^{2003} \beta^{t-2000} \left[\pi_{it}\left(\{q_{jt}\}_j, R_i^2\right) - P_t \cdot (e_{it} - a_{it})\right] \\
&\quad + \beta^{2003-2000} \left\{CV(h_{i,T+1}) - P_T h_{i,T+1}\right\} + P_{2000} h_{i,2000},
\end{aligned}
$$

where the last equality uses the equilibrium relationship $\beta P_{t+1} = P_t$. The investment problem is

$$
\begin{aligned}
W_{i,2000}(h_{i,2000}, R_i^1) = \quad & \max_{R_i^2} \quad V_{2000}(h_{i,2000}, R_i^2) - \Gamma(R_i^2, R^1). \\
& s.t. \quad R_i^2 \leq R_i^1.
\end{aligned}
$$

Note that $h_{i,2000}$ does not affect the optimal investment level of R_i^2.

The continuation value at the beginning of Phase I is given as

$$\begin{aligned} V_{1995}(h_{i,1995}, R_i^1) &= \sum_{t=1995}^{1999} \beta^{t-1995} \left[\pi_{it}\left(\{q_{jt}\}_j, R_i^1\right) - P_t(e_{it} - a_{it})\right] \\ &\quad + \beta^{1999-1995}\left(\beta W_{2000}(h_{i,2000}, R_i^1) - P_{1999} h_{i,2000}\right). \end{aligned}$$

The investment problem is similar to the one in Phase II.

Finally, I consider the market clearing condition. By aggregating the transition equation of permit holding (1.3.2) over individual firms and time, we have

$$\sum_{t=1995}^{2003} E_t(P_t) + H_{T+1} = \sum_{t=1995}^{2003} A_t + \sum_{t=1995}^{2003} B_t, \tag{A.2.1}$$

where $E_t = \sum_i e_{it}(P_t)$, and other uppercase variables are similarly defined. The market clearing condition in each period is

$$B_t + \bar{B}_t^{fringe}(P_t) = 0.$$

By putting this condition into equation (A.2.1), we have

$$\sum_{t=1995}^{2003} E_t\left(\beta^{-(t-1)} P_{1995}\right) + H_{T+1}(\beta^{-(T-1)} P_{1995}) = \sum_{t=1995}^{2003} A_t + \sum_{t=1995}^{2003} -\bar{B}_t^{fringe}\left(\beta^{-(t-1)} P_{1995}\right).$$

The equilibrium price P_1 is determined by this equation, and thus the whole path of the equilibrium price.

A.3. Derivations and Proof

A.3.1. Derivation of $\partial EV_t(h_t, I_t)/\partial h_t$

I omit an index i for a particular firm for expositional purposes. I focus on the derivation of $\frac{\partial EV_t(h_t,0)}{\partial h_t}$. Recall that

$$EV_t(h_t, 0) \;\; = \;\; \int \max\left\{V_t^0(h_t), V_t^1(h_t) - F_t - \epsilon\right\} dG(\epsilon).$$

By the chain rule, we have

$$\frac{dEV_t(h_t,0)}{dh_t} = \frac{\partial EV_t}{\partial V_t^0}\frac{dV_t^0}{dh_t} + \frac{\partial EV_t}{\partial V_t^1}\frac{dV_t^1}{dh_t}.$$

First, we derive $\frac{\partial EV_t}{\partial V_t^k}$ for $k = 0, 1$. This is an application of the Williams-Daly-Zachary theorem (see Theorem 3.1 in Rust, 1994):

$$\frac{\partial EV_t(h_t)}{\partial h_t} = \mathbb{P}(h_t)\left\{P_t + T'(b_t^{trade})\right\} + (1 - \mathbb{P}(h_t))\,\pi'(e_t^{not}).$$

By using the interchange of integration and differentiation, we have the following (I omit h_t for expositional purposes in the following derivation):

$$\begin{aligned}
\frac{\partial EV_t}{\partial V_t^1} &= \frac{\partial}{\partial V_t^1}\int \max\left\{V_t^1 - F_t - \epsilon, V_t^0\right\} dG(\epsilon) \\
&= \frac{\partial}{\partial V_t^1}\int_{\Upsilon^1}(V_t^1 - F_t - \epsilon)dG(\epsilon) + \frac{\partial}{\partial V_t^1}\int_{\Upsilon^0} V_t^0 dG(\epsilon) \\
&= \int_{\Upsilon^1}\frac{\partial}{\partial V_t^1}(V_t^{trade} - F_t - \epsilon)dG(\epsilon) + \int_{\Upsilon^0}\frac{\partial}{\partial V_t^1}V_t^0 dG(\epsilon) \\
&= \int_{\Upsilon^1} dG(\epsilon) \\
&= \mathbb{P}_t(h_t),
\end{aligned}$$

where Υ^1 is the set of ϵ such that a firm chooses to participate, i.e., $\Upsilon^1 \equiv \{\epsilon : V_t^1 - F_t - \epsilon > V_t^0\}$, and Υ^0 is similarly defined. Note that we can apply the similar derivation to obtain $\frac{\partial EV_t}{\partial V_t^0} = 1 - \mathbb{P}(h_t)$.

Next, we calculate $\frac{\partial V_t^k}{\partial h_t}$for $k = 0, 1$. The derivation is a direct application of the envelope theorem (or the Benveniste-Scheinkman formula):

$$\frac{\partial V_t^k}{\partial h_t} = \lambda_t^k,$$

where λ_{it}^k is the Lagrange multipliers in the corresponding optimization problems. Thus, we obtain

$$\frac{dEV_t(h_t, 0)}{dh_t} = \mathbb{P}_t(h_t)\lambda_t^1 + (1 - \mathbb{P}_t(h_t))\,\lambda_t^0.$$

A.3.2. Proof of Comparative Statics

This subsection shows the comparative statics in section 1.3.8. I omit an index of firm i for expositional purposes. I also re-write the profit from the electricity market as a function of emissions volume e_{it} for the purpose of exposition. The optimization problem for the trader is now given by

$$\begin{aligned}
\max_{e_t, b_t, h_{t+1}} \quad & \pi_t(e_t) - (P_t b_t + TC(b_t)) + \beta EV_{t+1}(h_{t+1}, 1) \\
\text{s.t.} \quad & e_t + h_{t+1} = a_t + h_t + b_t, \\
& h_{t+1} \geq 0,
\end{aligned}$$

and the problem for the non-trader is similarly defined. For simplicity, I assume that the non-borrowing constraint is not binding: $\mu_t = 0$ throughout the proof.

First, I focus on the case in which a firm does not participate in trading. I show that e_t and h_{t+1} are increasing in h_t. The optimality condition in this case is given by

$$\pi_t'(e_t) = \beta \frac{\partial EV_{t+1}(\overbrace{a_t + h_t - e_t}^{=h_{t+1}}, 0)}{\partial h_{t+1}}.$$

Using an implicit function theorem, we have

$$\begin{aligned}
\frac{de_t}{dh_t} &= -\frac{-\beta EV_{t+1}''(a_t + h_t - e_t, 0)}{\pi''(e_t) + \beta EV_{t+1}''(a_t + h_t - e_t, 0)} \\
&= \frac{\beta EV_{t+1}''(a_t + h_t - e_t, 0)}{\pi''(e_t) + \beta EV_{t+1}''(a_t + h_t - e_t, 0)} \\
&\in (0, 1),
\end{aligned}$$

where $EV_{t+1}''(\cdot, 0) = \partial^2 EV_{t+1}(h_{t+1}, 0)/\partial h_{t+1}^2$. Note that both $EV_{t+1}(\cdot)$ and $\pi(\cdot)$ are concave functions, so that their second derivatives are non-positive. Because $h_{t+1} = (a_t + h_t) - e_t$,

$$\frac{\partial h_{t+1}}{\partial h_t} = 1 - \frac{\partial e_t}{\partial h_t} > 0.$$

Next, I consider the case in which a firm participates in trading. The optimality conditions, given by equations (1.3.6) and (1.3.8), are

$$\begin{aligned}
\pi'(e_t) - P_t - TC'(e_t + h_{t+1} - a_t - h_t) &= 0 \\
\pi'(e_t) - \beta EV_{t+1}'(h_{t+1}, 1) &= 0.
\end{aligned}$$

Taking the total derivative of these equations with respect to h_t, we have

$$\begin{aligned}
(\pi'' - TC'')\frac{\partial e_t}{\partial h_t} + (-TC'')\frac{\partial h_{t+1}}{\partial h_t} + TC'' &= 0 \\
\pi''\frac{\partial e_t}{\partial h_t} - \beta EV''_{t+1}\frac{\partial h_{t+1}}{\partial h_t} &= 0.
\end{aligned}$$

Solving these equations gives me

$$\begin{aligned}
\frac{\partial e_t}{\partial h_t} &= \frac{-TC''}{\pi'' - TC'' - TC''\frac{\pi''}{\beta EV''_{t+1}}} > 0 \\
\frac{\partial h_{t+1}}{\partial h_t} &= \frac{\pi''}{\beta EV''_{t+1}}\frac{\partial e_t}{\partial h_t} > 0.
\end{aligned}$$

Thus, h_t increases both e_t and h_{t+1}. Finally, h_t decreases b_t because

$$\begin{aligned}
\frac{\partial b_t}{\partial h_t} &= \frac{\partial e_t}{\partial h_t} + \frac{\partial h_{t+1}}{\partial h_t} - 1 \\
&= \frac{-TC''}{\pi'' - TC'' - TC''\frac{\pi''}{\beta EV''_{t+1}}} + \frac{-TC''\frac{\pi''}{\beta EV''_{t+1}}}{\pi'' - TC'' - TC''\frac{\pi''}{\beta EV''_{t+1}}} - 1 \\
&= \frac{-TC'' - TC''\frac{\pi''}{\beta EV''_{t+1}}}{\pi'' - TC'' - TC''\frac{\pi''}{\beta EV''_{t+1}}} - 1 \\
&= \frac{-\pi''}{\pi'' - TC'' - TC''\frac{\pi''}{\beta EV''_{t+1}}} < 0.
\end{aligned}$$

APPENDIX B

Appendix for Chapter 2

B.1. Derivation of the Calculus of Voting Model

In this Appendix, we follow Riker and Ordeshook (1968) and present a derivation of expressions (1) and (2). We classify the situation of a voter into the following five mutually exclusive events:

E_0: votes for D and R are tied without her vote;

E_{D1}: D has exactly one more vote than R without her vote;

E_{R1}: R has exactly one more vote than D without her vote;

E_{D2}: D has two or more votes than R without her vote;

E_{R2}: R has two or more votes than D without her vote.

Let q_l denote the probability of E_l for $l \in \{0, D1, R1, D2, R2\}$. Let π be the probability that D wins the election in case of a tie. Then, the utility of the voter for voting for candidates D and R, as well as not voting, are written as

$$\begin{aligned}
U_D &= q_0 b_D + q_{D1} b_D + q_{R1}(\pi b_D + (1-\pi) b_R) + q_{D2} b_D + q_{R2} b_R - c, \\
U_R &= q_0 b_R + q_{D1}(\pi b_D + (1-\pi) b_R) + q_{R1} b_R + q_{D2} b_D + q_{R2} b_R - c, \\
U_0 &= q_0(\pi b_D + (1-\pi) b_R) + (q_{D1} + q_{D2}) b_D + (q_{R1} + q_{R2}) b_R.
\end{aligned}$$

By taking the difference between voting for D and not voting, we have

$$\begin{aligned} U_D - U_0 &= q_0 b_D + q_{D1} b_D + q_{R1}(\pi b_D + (1-\pi) b_R) + q_{D2} b_D + q_{R2} b_R - c \\ &\quad - q_0(\pi b_D + (1-\pi) b_R) - (q_{D1} + q_{D2}) b_D - (q_{R1} + q_{R2}) b_R \\ &= \left(q_0(1-\pi) + q_{R1}\pi\right) b_D - \left(q_0(1-\pi) + q_{R1}\pi\right) b_R - c \\ &= \left(q_0(1-\pi) + q_{R1}\pi\right)(b_D - b_R) - c. \end{aligned}$$

Similarly, we have

$$\begin{aligned} U_R - U_0 &= q_0 b_R + q_{R1} b_R + q_{D1}(\pi b_D + (1-\pi) b_R) + q_{R2} b_R + q_{D2} b_D - c \\ &\quad - q_0(\pi b_D + (1-\pi) b_R) - (q_{R1} + q_{R2}) b_R - (q_{D1} + q_{D2}) b_D \\ &= \left(q_{D1}(1-\pi) + q_0\pi\right) b_R + \left(-q_{D1}(1-\pi) - q_0\pi\right) b_D \\ &= \left(q_{D1}(1-\pi) + q_0\pi\right)(b_R - b_D) - c. \end{aligned}$$

Assuming that $q_0 = q_{D1} = q_{R1} \equiv p$ in a large election (see page 103 of Myerson and Weber, 1993, for a justification), we can rewrite U_D, U_R, and U_0 as

$$\begin{aligned} U_D &= p(b_D - b_R) - c \\ U_R &= p(b_R - b_D) - c \\ U_0 &= 0. \end{aligned}$$

B.2. Data Construction

In this Appendix, we explain how we construct the joint distribution of demographic characteristics and citizenship status at the county level. We first use the 5% Public Use

Microdata Sample of the 2000 U.S. Census (hereafter PUMS), which is an individual-level dataset, to estimate the covariance matrix between the demographic variables and citizenship information within each public use microdata area (PUMA). In particular, we estimate the joint distribution of the discrete demographic characteristics (Race, Hispanic, Citizenship) by counting the frequency of occurrence. We then estimate a covariance matrix for the continuous demographic variables (Age, Income, Years of Schooling) for each bin. Because the PUMA and counties do not necessarily coincide, we estimate covariance matrices for each PUMA and then use the correspondence chart provided in the PUMS website to obtain estimates at the county level.

In the second step, we construct the joint distribution of demographic characteristics by combining the covariance matrix estimated in the first step and the marginal distributions of each of the demographic variables at the county level obtained from Census Summary File 1 through File 3. We discretize continuous variables into coarse bins. We discretize age into three bins: (1)18-34 years old; (2) 35-59 years old; and (3) above 60 years old; income into 6 bins: (1) \$0-\$25,000; (2) \$25,000-\$50,000; (3) \$50,000-\$75,000; (4) \$75,000-\$100,000; (5) \$100,000-\$150,000; and (6) above \$150,000; and years of schooling into 5 bins: (1) Less than 9th grade; (2) 9th-12th grade with no diploma; (3) high school graduate; (4) some college with no degree or associate degree; and (5) bachelor's degree or higher. Thus, there are 540 bins in total. The joint distribution of demographic characteristics that we create gives us a probability mass over each of the 540 bins for each county.

Finally, we augment the census data with religion data obtained from Religious Congregations and Membership Study 2000. These data contain information on the share of the population with adherence to either "Evangelical Denominations" or "Church of Jesus Christ of Latter-day Saints" at the county level. Because the Census does not collect information on religion, we do not know the correlation between the religion variable and the demographic characteristics in the Census. Thus, we assume independence of the religion variable and other demographic variables. As a result, there are 1,080 bins in our demographics distribution.

B.3. Identification of $c(\cdot)$, $p(\cdot)$, and F_η in the general case

In this Appendix, we show that $c(\cdot)$, $p(\cdot)$, and F_η are identified even when the max operator in equation (2.4.5) binds with positive probability. Note that our argument in the main text considered only the case in which the max operator never binds. Recall that

$$\underbrace{\frac{\Phi^{-1}\left(1-v_{R,m}\right)-\Phi^{-1}\left(v_{D,m}\right)}{2}}_{\equiv Y_m}=\max\left\{0,\frac{c(\bar{\mathbf{x}}_m)}{p(\bar{\mathbf{x}}_m)}+\frac{\eta_m}{p(\bar{\mathbf{x}}_m)}\right\},\ \eta_m\perp\bar{\mathbf{x}}_m. \tag{B.3.1}$$

In this Appendix, we work with the normalization that the value of $p(\cdot)$ at some $\bar{\mathbf{x}}_m=\mathbf{x}_0$ as $p(\mathbf{x}_0)=1$. This amounts to a particular normalization of variance of η. Note that the distribution of Y_m (the left hand side of equation (B.3.1)) conditional on $\bar{\mathbf{x}}_m=\mathbf{x}_0$ is a truncated distribution with mass at zero. Figure B.1 illustrates this when the mass at zero is less than 50%, and F_η is symmetric and single-peaked at zero.

First, we present our identification discussion for the case that F_η is symmetric and single-peaked at zero. As Figure B.1 illustrates, the median of Y_m conditional on $\bar{\mathbf{x}}_m=\mathbf{x}_0$ directly identifies $c(\mathbf{x}_0)$ under these assumptions. Also, the density of η, f_η, is identified above the point of truncation. Formally, $f_\eta(F_\eta^{-1}(t))$ is identified for any $t>t(\mathbf{x}_0)$, where

$$t(\mathbf{x}_0)=\Pr\left(Y_m=0|\,\mathbf{x}_0\right).$$

Hence, $f_\eta(0)$ is identified from the height of the density of Y_m at the median.

Figure B.1. The distribution of Y_m conditional on $\mathbf{x} = \mathbf{x}_0$ and $\mathbf{x} = \mathbf{x}_1$.

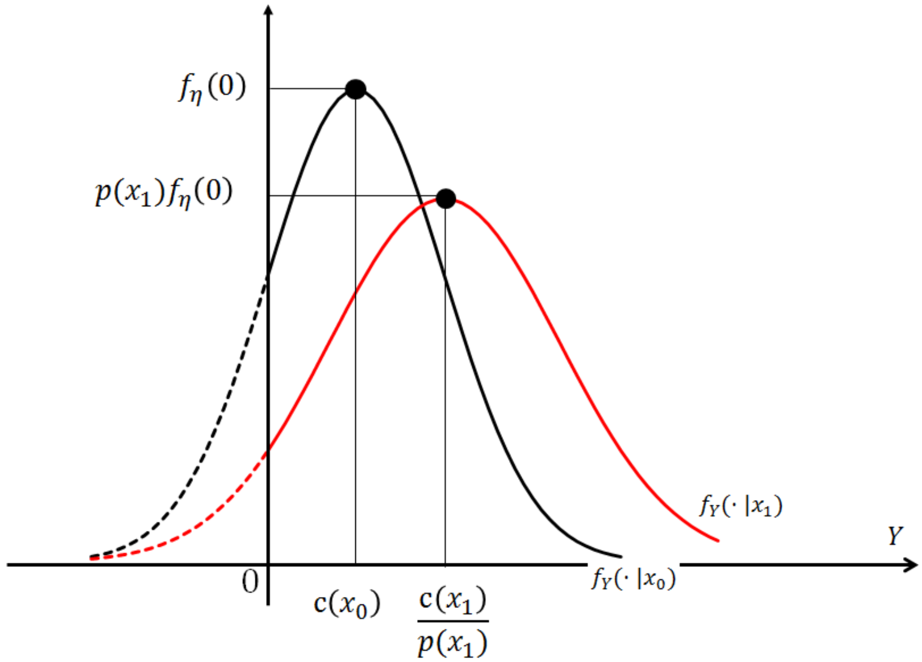

Note: The distribution of Y_m conditional on $\mathbf{x} = \mathbf{x}_0$ and $\mathbf{x} = \mathbf{x}_1$ is the one when the distribution of η is symmetric and single-peaked, and $t(\mathbf{x}_0)$, $t(\mathbf{x}_1) < 0.5$, where $t(\mathbf{x})$ is the probability that Y_m is equal to zero conditional on $\mathbf{x}$. The distribution of Y_m is truncated at zero. The conditional median of Y_m identifies $c(\mathbf{x}_0)$ and $c(\mathbf{x}_1)/p(\mathbf{x}_1)$, and the height of the density at the conditional median identifies $f_\eta(0)$ and $p(\mathbf{x}_1)f_\eta(0)$.

Now, consider $\mathbf{x}_1 \neq \mathbf{x}_0$. Assume, again, that $t(\mathbf{x}_1) < 0.5$. Then, $c(\mathbf{x}_1)/p(\mathbf{x}_1)$ is identified from the conditional median of Y_m, and $p(\mathbf{x}_1)f_\eta(0)$ is identified by the height of the conditional density of Y_m at the median. Given that $f_\eta(0)$ is identified, $c(\mathbf{x}_1)$ and $p(\mathbf{x}_1)$ are both identified. Moreover, F_η is identified over its full support if there exists sufficient variation in $\mathbf{x}$, i.e., $\inf_{\mathbf{x}} t(\mathbf{x}) = 0$.

We now consider the case in which F_η is not restricted to being symmetric and single-peaked and $t(\mathbf{x}_0)$ may be less than 0.5. The distribution of Y_m is identified above $t(\mathbf{x}_0)$, as before. Now, consider $\mathbf{x}_1 \neq \mathbf{x}_0$. Similar to before, we identify $p(\mathbf{x}_1)f_\eta(F_\eta^{-1}(\tau))$ for τ

above $t(\mathbf{x}_1)$.[1] If we let τ be any number larger than $\max\{t(\mathbf{x}_0), t(\mathbf{x}_1)\}$, both $f_\eta(F_\eta^{-1}(\tau))$ and $p(\mathbf{x}_1)f_\eta(F_\eta^{-1}(\tau))$ are identified. Hence, $p(\mathbf{x}_1)$ is identified. Similarly, $p(\cdot)$ is identified for all $\mathbf{x}$.

We now consider identification of $c(\cdot)$. We present two alternative assumptions on F_η and show that $c(\cdot)$ can be identified under either assumption. First, assume that the median of η is zero, $Med(\eta) = 0$, and that there exists $\mathbf{x} = \mathbf{x}_2$ such that $t(\mathbf{x}_2) < 1/2$. The latter assumption means that more than half of the counties have turnout less than 100% when $\mathbf{x} = \mathbf{x}_2$. Then, the median of Y_m conditional on $\mathbf{x}_2$ identifies $c(\mathbf{x}_2)/p(\mathbf{x}_2)$. Now, consider any $\mathbf{x}_1 \neq \mathbf{x}_2$ and let τ be any number larger than $\max\{t(\mathbf{x}_2), t(\mathbf{x}_1)\}$. Let z_1 and z_2 be the τ quantile of Y_m conditional on $\mathbf{x}_1$ and $\mathbf{x}_2$, respectively. z_1 and z_2 are clearly identified. Then,

$$
\begin{aligned}
& p(\mathbf{x}_1)\left[F^{-1}_{\eta/p(\mathbf{x}_1)}(\tau) - F^{-1}_{\eta/p(\mathbf{x}_1)}(1/2)\right] = p(\mathbf{x}_2)\left[F^{-1}_{\eta/p(\mathbf{x}_2)}(\tau) - F^{-1}_{\eta/p(\mathbf{x}_2)}(1/2)\right] \\
& \Leftrightarrow p(\mathbf{x}_1)\left[z_1 - \frac{c(\mathbf{x}_1)}{p(\mathbf{x}_1)}\right] = p(\mathbf{x}_2)\left[z_2 - \frac{c(\mathbf{x}_2)}{p(\mathbf{x}_2)}\right] \\
& \Leftrightarrow \frac{c(\mathbf{x}_1)}{p(\mathbf{x}_1)} = z_1 - \frac{p(\mathbf{x}_1)}{p(\mathbf{x}_2)}\left(z_2 - \frac{c(\mathbf{x}_2)}{p(\mathbf{x}_2)}\right). \qquad \text{(B.3.2)}
\end{aligned}
$$

Given that all of the terms on the right hand side of (B.3.2) are identified, $c(\mathbf{x}_1)/p(\mathbf{x}_1)$ is identified.

Alternatively, assume that $\mathbf{E}(\eta) = 0$ and $\inf_{\mathbf{x}} t(\mathbf{x}) = 0$. We now show that $c(\cdot)$ is identified under these alternative assumptions. Intuitively, this latter assumption means that there exist values of $\mathbf{x}$ for which the max operator is never binding. In this case, we can fully recover the distribution of $F_\eta(\cdot)$. Then, we can identify the distribution of $c(\mathbf{x})/p(\mathbf{x}) + \eta_m/p(\mathbf{x})$ for any $\mathbf{x}$. Hence, we identify $c(\cdot)/p(\cdot)$.

[1] Note that we identify $f_{\eta/p(\mathbf{x}_1)}\left(F^{-1}_{\eta/p(\mathbf{x}_1)}(t)\right)$, where $f_{\eta/p(\mathbf{x}_1)}(\cdot)$ and $F^{-1}_{\eta/p(\mathbf{x}_1)}(\cdot)$ are the density of $\eta/p(\mathbf{x}_1)$ and the inverse distribution of $\eta/p(\mathbf{x}_1)$, respectively. Note that $f_{\eta/p(\mathbf{x}_1)}\left(F^{-1}_{\eta/p(\mathbf{x}_1)}(t)\right) = p(\mathbf{x}_1)f_\eta(F_\eta^{-1}(t))$.

Table B.1. Estimates of State Preference Fixed Effects Relative to $\lambda_{Alabama}$

	Estimate	SE		Estimate	SE
Alabama	0 (Normalized)		Nevada	-0.009	(0.059)
Arizona	-0.132	(0.050)	New Jersey	-0.247	(0.041)
Arkansas	-0.363	(0.034)	New Mexico	-0.055	(0.069)
California	-0.236	(0.060)	New York	-0.239	(0.036)
Colorado	-0.057	(0.055)	North Carolina	-0.088	(0.033)
Florida	-0.101	(0.039)	North Dakota	-0.006	(0.040)
Georgia	0.107	(0.028)	Ohio	-0.166	(0.035)
Idaho	0.129	(0.044)	Oklahoma	-0.057	(0.041)
Illinois	-0.243	(0.030)	Oregon	-0.130	(0.054)
Indiana	-0.099	(0.031)	Pennsylvania	-0.190	(0.037)
Iowa	-0.300	(0.032)	South Carolina	0.030	(0.031)
Kansas	0.194	(0.038)	South Dakota	0.000	(0.039)
Kentucky	-0.251	(0.032)	Tennessee	-0.345	(0.034)
Louisiana	0.163	(0.030)	Texas	0.127	(0.034)
Maryland	-0.107	(0.053)	Utah	0.079	(0.062)
Michigan	-0.246	(0.031)	Virginia	-0.114	(0.035)
Minnesota	-0.297	(0.032)	Washington	-0.223	(0.055)
Mississippi	0.128	(0.033)	West Virginia	-0.216	(0.037)
Missouri	-0.217	(0.029)	Wisconsin	-0.335	(0.033)
Montana	0.098	(0.044)	Wyoming	0.230	(0.057)

Note: Standard errors are reported in parentheses. Higher values imply a stronger preference for Democrats.

B.4. Additional Tables

We report the estimates of state-specific effects on preference and efficacy in Tables B.1 and B.2, which we use to plot Figures 2.4 and 2.5.

B.5. Fit

In this Appendix, we report further on the fit of the model. Figure B.2 plots the distribution of Democratic and Republican vote shares in the data and in the model prediction.

Table B.2. Estimates of State-level Fixed Effects of Voting Efficacy

	Estimate	SE		Estimate	SE
Alabama	1 (Normalized)		Nevada	0.854	(0.075)
Arizona	0.660	(0.109)	New Jersey	0.725	(0.054)
Arkansas	0.739	(0.043)	New Mexico	0.976	(0.083)
California	0.751	(0.053)	New York	0.795	(0.052)
Colorado	1.165	(0.073)	North Carolina	0.787	(0.040)
Florida	0.975	(0.054)	North Dakota	1.062	(0.056)
Georgia	0.842	(0.039)	Ohio	1.323	(0.057)
Idaho	1.274	(0.081)	Oklahoma	0.919	(0.041)
Illinois	1.067	(0.041)	Oregon	1.441	(0.086)
Indiana	0.814	(0.042)	Pennsylvania	0.838	(0.046)
Iowa	1.319	(0.055)	South Carolina	0.751	(0.045)
Kansas	0.877	(0.043)	South Dakota	1.680	(0.105)
Kentucky	1.052	(0.037)	Tennessee	0.867	(0.041)
Louisiana	1.316	(0.079)	Texas	0.769	(0.041)
Maryland	0.767	(0.050)	Utah	1.229	(0.087)
Michigan	1.180	(0.048)	Virginia	0.754	(0.043)
Minnesota	1.765	(0.103)	Washington	1.080	(0.051)
Mississippi	1.193	(0.067)	West Virginia	0.814	(0.043)
Missouri	1.159	(0.040)	Wisconsin	1.762	(0.108)
Montana	1.066	(0.054)	Wyoming	1.151	(0.060)

Note: Standard errors are reported in parentheses. Alabama is set to 1 for normalization.

The figure shows that the model fits the data well for all ranges of the Democratic and Republican vote shares. Lastly, we compute the χ^2 statistic for the goodness-of-fit test. The χ^2 test statistics for Democratic and Republican vote shares are 14.07 and 45.29, respectively. The former do not reject the null that these two distributions are the same at 5% level, while the latter rejects it.

Figure B.2. Model Fit

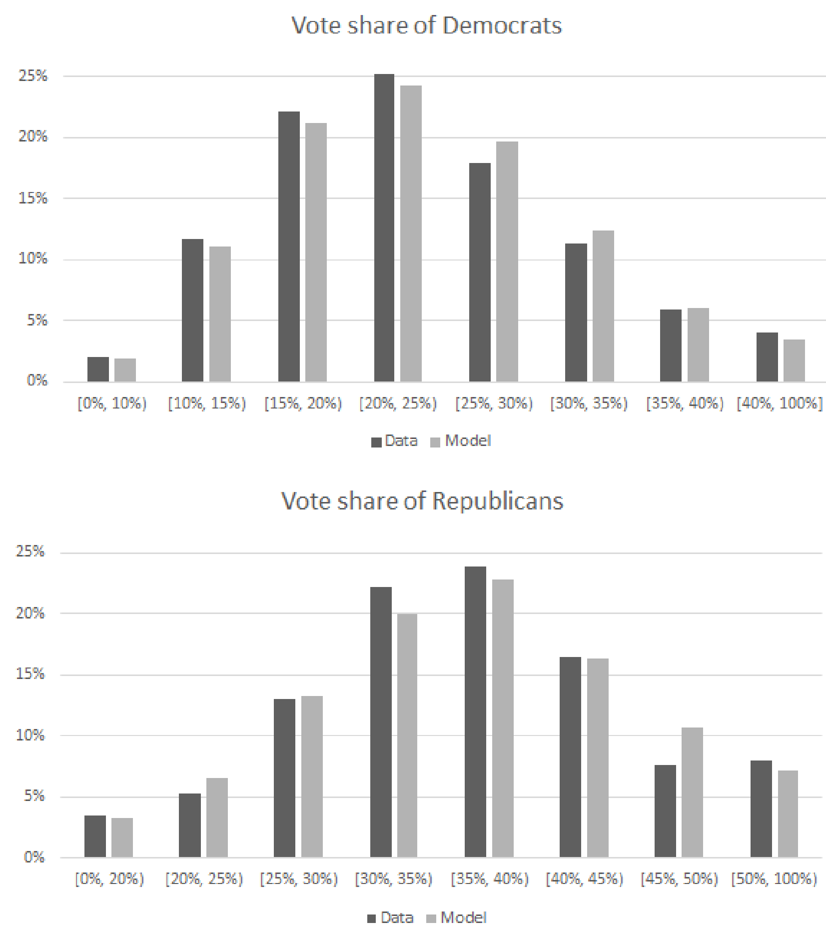

Note: The top panel plots the distributions of Democratic vote share in the data and in the model prediction. The bottom panel plots the distributions for the Republican vote share.