

Harvard Environmental Economics Program

DEVELOPING INNOVATIVE ANSWERS TO TODAY'S COMPLEX ENVIRONMENTAL CHALLENGES

May 2018 Discussion Paper 18-76

Slow Focus: Belief Evolution in the U.S. Acid Rain Program

Cuicui Chen Ph.D. 2018, Harvard University, Public Policy



heep@harvard.edu http://heep.hks.harvard.edu

Slow Focus: Belief Evolution in the U.S. Acid Rain Program

Cuicui Chen

Harvard University

The Harvard Environmental Economics Program

The Harvard Environmental Economics Program (HEEP) develops innovative answers to today's complex environmental issues, by providing a venue to bring together faculty and graduate students from across Harvard University engaged in research, teaching, and outreach in environmental, energy, and natural resource economics and related public policy. The program sponsors research projects, convenes workshops, and supports graduate education to further understanding of critical issues in environmental, natural resource, and energy economics and policy around the world. For more information, see the Program's website: http://heep.hks.harvard.edu.

Acknowledgements

The Enel Endowment for Environmental Economics at Harvard University provides major support for HEEP. The Endowment was established in February 2007 by a generous capital gift from Enel SpA, a progressive Italian corporation involved in energy production worldwide. HEEP receives additional support from the affiliated Enel Foundation.

HEEP also receives support from BP, Chevron Services Company, and Shell. HEEP enjoys an institutional home in and support from the Mossavar-Rahmani Center for Business and Government at the Harvard Kennedy School. HEEP collaborates closely with the Harvard University Center for the Environment (HUCE). The Center has provided generous material support, and a number of HUCE's Environmental Fellows and Visiting Scholars have made intellectual contributions to HEEP.

HEEP and the closely-affiliated Harvard Project on Climate Agreements are grateful for additional support from the Harvard University Climate Change Solutions Fund; the Harvard Global Institute; the Belfer Center for Science and International Affairs and the Ash Center for Democratic Governance and Innovation—both located at the Harvard Kennedy School; Christopher P. Kaneb (Harvard AB 1990); and the International Emissions Trading Association (IETA).

Previous sponsors include: the Alfred P. Sloan Foundation, the AVINA Foundation, Bank of America, Castleton Commodities International LLC, ClimateWorks Foundation, Doris Duke Charitable Foundation, Duke Energy Corporation, the James M. and Cathleen D. Stone Foundation, the National Science Foundation, the U.S. Department of Energy, and the U.S. Environmental Protection Agency.

Citation Information

Chen, Cuicui. "Slow Focus: Belief Evolution in the U.S. Acid Rain Program." Discussion Paper 2018-76. Cambridge, Mass.: Harvard Environmental Economics Program, May 2018.

The views expressed in the Harvard Environmental Economics Program Discussion Paper Series are those of the author(s) and do not necessarily reflect those of the Harvard Kennedy School or of Harvard University. Discussion Papers have not undergone formal review and approval. Such papers are included in this series to elicit feedback and to encourage debate on important public policy challenges. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

Slow Focus: Belief Evolution in the U.S. Acid Rain Program

Cuicui Chen*

Harvard University

May 2, 2018

Abstract

I study firm behavior in new markets by examining coal-dependent private electric utilities' beliefs about the sulfur dioxide allowance price following the implementation of the U.S. Acid Rain Program. The program is the first large-scale cap-and-trade program, exposing the electric utility industry to a wholly novel market for pollution allowances. I estimate firms' beliefs about the allowance price from 1995 to 2003 using a firm-level dynamic model of allowance trades, coal quality, and emission reduction investment. I find that firms initially underestimate the role of market fundamentals as a driver of allowance prices, but over time their beliefs appear to converge toward the stochastic process of allowance prices. Such beliefs in the first five years of the program cost firms around 10% of their profits. Beliefs also change the relative efficiency of cap-and-trade programs and emission taxes.

^{*}PhD Candidate in Public Policy, Harvard University. Email: cuicuichen@fas.harvard.edu. I am extremely grateful to my advisors — Ariel Pakes, Elie Tamer, Joe Aldy, Richard Zeckhauser and Rob Stavins — for their continual support and warm encouragement. I have also benefitted immensely from discussions with Rob Bray, Myrto Kalouptsidi, Robin Lee, and Forest Reinhardt. Special thanks go to Dallas Burtraw, Denny Ellerman, Kenon Smith, and the industry interviewees for patiently answering my endless questions about institutional details, and to Paul Edmon, Mike Ethier, and the AMPL company for generously providing crucial computational support. I also thank participants at the Environmental Lunch, the Industrial Organization Workshop, and Seminar in Environmental Economics and Policy at Harvard for their helpful comments. Financial support from the Environment and Natural Resources Program at Harvard Kennedy School and the Harvard Institute for Quantitative Social Science is gratefully acknowledged.

1 Introduction

Progress is impossible without change, and those who cannot change their minds cannot change anything.

> George Bernard Shaw Everybody's Political What's What, 1945

The 1990 amendments to the Clean Air Act give electric power companies the right to pay to pollute, much the way sinners bought indulgences in the Middle Ages. Yet although the trading of sulfur emissions between utilities is now permitted, there is not much activity in this market, and prices of marketable permits are far below initial projections.

Peter Passell for *The New York Times* Paying to pollute: a free market solution that's yet to be tested, 1996

New economic environments abound. Technology creates new markets, regulation defines new boundaries, and industry dynamics shape new competitive landscapes. As firms start out in an environment they lack experience in, they may not form reasonable expectations about future market conditions. While adaptability has been much advocated for in the business literature, how firms actually behave in a new environment is little understood in economics.

This paper investigates the beliefs of private electric utilities about the sulfur dioxide (SO_2) allowance price under the U.S. Acid Rain Program from 1995 to 2003. As the first large-scale cap-and-trade program, the Acid Rain Program thrust the electric utility industry into an unprecedented market-based environmental regulation. In planning to comply with the Acid Rain Program, utilities need to form beliefs about the future SO_2 allowance price. Given the novelty of this environment, do utilities hold beliefs that are at least as good as what they could have inferred from available data? Do beliefs improve over time as utilities gain experience in the allowance market? What do the beliefs imply for firm profits and social welfare? Answers to those questions are important for choosing, designing and evaluating cap-and-trade programs, now the predominant policy for addressing climate change, as well as various forms of other quantity-based instruments¹. The lessons are particularly relevant for countries, such as China, that expose traditional industries to market-based environmental regulation for the first time.

I use dynamic structural estimation to back out those beliefs about the future allowance price from firm-level allowance trades, coal quality choices, and emission reduction investment. The intuition is simple: the higher a utility expects tomorrow's allowance price to be, the more aggressively it saves allowances, by buying more (or selling fewer) allowances,

¹Cap-and-trade programs belong to the broader class of quantity-based instruments, used widely in environmental regulations. Notable examples of other quantity-based instruments include the Renewable Identification Number credits for biodiesel and other renewable fuels, and the Corporate Average Fuel Economy (CAFE) credits.

reducing the sulfur content of coal, and making more emission reduction investment. In estimating the beliefs, I make three contributions. First, I build the first empirical dynamic model of firm behavior in cap-and-trade programs. Dynamics is necessary for understanding firm behavior in cap-and-trade programs, because all compliance decisions have dynamic implications via the saving of allowances.² Second, I provide the first empirical application of two novel dynamic programming acceleration methods, the Relative Value Function Iterations method (Bray, 2017b) and the Endogenous Value Function Iterations method (Bray, 2017a). They vastly reduce the computation burden. Third, this is the first time that extensive SO₂ allowance trading data have been matched to electric operations data for academic research.³

Existing evidence suggests that the initial beliefs about the future allowance price were quite unreasonable. Figure 1 shows that by 1995, all major predictions of the 1995-1999 average allowance price (summarized by Hahn and May (1994)) exceeded the actual prices, most of which by a large margin. The expected 1995, 1997, and 1999 allowance price surveyed by Fieldston Company (1993) were no better. It is not convincing that those disparities, being consistent and vast, are merely due to the difference between expectation and realization. Neither do they arise simply from regulatory capture; the rules had been set years ago. A plausible culprit, however, is that utilities may not have made good use of available information; Montero and Ellerman (1998) attribute 60% of the over-prediction to the failure to incorporate the decline in the low-sulfur coal price that had been going on since 1980. The inefficient use of information is likely because of utilities' inexperience with the unprecedented allowance market. My private communication with utility executives, regulators, and allowance brokers confirms the utilities' initial difficulty with the allowance market; utilities were wondering about "how is there going to be a market," "what are we going to do," and many small firms appeared "not up for the idea of market and trading."

To investigate whether the beliefs have improved since 1995 as more information became available,⁴ I first estimate the beliefs from a dynamic model of cap-and-trade compliance. A coal-dependent private electric utility chooses the net purchase of allowances, the sulfur content of coal, and capital investment to allow fuel-switching, based on electricity demand, allowance price, and allowance stock, subject to the requirement that it have enough allowances to cover the SO₂ emissions. The model takes into account the allowance transaction cost, the distortion of incentives by cost-of-service regulation, the vintage structure of allowances, and the two-phase compliance structure of the Acid Rain Program. I estimate this three-dimensional continuous-state dynamic model, where the allowance price beliefs are the perceived law of motion for the allowance price state variable, using a nested-fixedpoint-style algorithm with a maximum simulated likelihood estimator. The allowance price belief parameters are identified from inter- and intra-firm variations in compliance behavior

²Fowlie et al. (2016) develops a dynamic model of the cement industry under various environmental policies, including a cap-and-trade program; they focus on entry and exit decisions and assume away allowance trading and allowance price uncertainty. Other empirical models of firm behavior in cap-and-trade programs are static, including Fowlie (2010), Chan et al. (2017), and Cicala (2015).

³The Environmental Protection Agency hired a contractor to work on the match for several years, and Ellerman et al. (2000) includes academic research that uses the first three years of the matched data.

⁴That said, when the initial beliefs in Figure 1 were formed, there had been almost two years of allowance price observations by 1995, because the allowance market had started operating in 1993.



Figure 1: Predicted 1995-1999 average allowance price by various institutions (Hahn and May, 1994), predicted 1995, 1997, and 1999 allowance prices by survey respondents (Fieldston Company, 1993), and the actual allowance price from the beginning of the allowance market to 2004.

induced by variations in the expected marginal dynamic value of allowances, which in turn are induced by exogenous variations in the allowance price and electricity demand.

Having estimated the beliefs, I then compare them with two alternative beliefs widely used in structural work in industrial organization: full-information rational expectations, and adaptive-learning beliefs. The former comparison asks whether utilities know the "overall" stochastic process of allowance prices as does an econometrician who estimates the process using a full series of data. Some argue that full-information rational expectations require too much statistical knowledge from economic agents, and propose to put economic agents and econometricians on comparable footing (Hansen, 2007). Thus, the latter comparison is based on the "adaptive" stochastic process of allowance prices, which is estimated iteratively, as time rolls forward, using those data available only up to each rolling year.⁵

I find that firms underestimate the role of market fundamentals as a driver of the allowance price; they pay too much attention to what is happening inside the allowance market (the historical allowance prices) and too little to the conditions in related markets (such as electricity and coal markets) that drive the allowance price. Consequently, they predict a flatter allowance price trajectory than full-information rational expectations; the latter pre-

⁵The data include market prices of allowances as reported by trade journals and brokerage firms. The Environmental Protection Agency also reports the clearing prices of its small-scale annual allowance auctions, which I do not use for two reasons. First, the clearing prices are downward biased due to a design flaw of those auctions (Cason, 1993, 1995; Cason and Plott, 1996). Second, to the extent that clearing prices are useful, they should have already been somewhat incorporated in the market price of allowances.

dicts the rises and falls in the actual allowance price trajectory much better.⁶ Smaller firms, and firms facing less competitive pressure, exhibit flatter beliefs. Over the years, those beliefs appear to converge toward the adaptive stochastic process of the allowance price. This trend is consistent with the shift in the management practice in electric utilities during that period,⁷ where compliance decisions are made less by engineers but more by people with experience in markets and trading; utilities seem to have gradually grasped the philosophy of market-based environmental regulation that compliance is about profit maximization, not cost minimization (Reinhardt, 2000).⁸

The beliefs in the first five years of the Acid Rain Program cost private electric utilities an average dynamic payoff equivalent to around 10% of their profits. Under cost-of-service regulation, forgone payoffs to electric utilities are the lower bound on the forgone savings to ratepayers. Therefore, policies that improve the belief formation process of firms would be financially beneficial to ratepayers. Such policies include making available comprehensive market information in a timely and transparent manner, holding workshops and conferences to facilitate communications among utilities, brokers, and regulators, introducing competition to the electricity market, and using price ceilings and floors to constrain the volatility of the allowance price.

Beliefs about future market conditions make cap-and-trade programs more efficient than emission taxes when they align the marginal cost of emission reduction with the marginal benefit of doing so. The literature has missed this efficiency determinant because it has focused static models, where beliefs about future market conditions are irrelevant. Indeed, efficiency requires that each firm reduce their emissions up to the point where its marginal cost of emission reduction equals the marginal benefit it creates. Under a static model, neither cap-and-trade nor tax achieves efficiency, as they equalize the marginal abatement cost across firms, rather than tailor it to each firm's marginal benefit. It has thus been proposed that to achieve efficiency, government intervention is necessary, for example by setting up trading ratios (Muller and Mendelsohn, 2009). However, once we move beyond statics to dynamics, where beliefs are a key driver of behavior in cap-and-trade programs but much less relevant in taxes,⁹ beliefs become a decentralized channel that affects the relative efficiency of those two policy instruments. For example, given my empirical finding that bigger utilities tend to forecast better by not over-predicting allowance prices as much as do smaller utilities, they would be less aggressive in reducing emissions. The resulting lower marginal cost aligns with their lower marginal benefit of emission reduction, as they

⁶Some may wonder why the allowance price did not turn out to match the expectations of those firms per rational expectations. Note that the coal-dependent private electric utilities studied in this paper are only some of the participants in the allowance market; other important participants include coal-independent or public utilities, non-utility energy firms, and brokerage firms.

⁷Some have suggested the possibility that public utilities commissions, overseeing the utilities under costof-service regulation, may also have "learned" to better regulate the utilities in the presence the Acid Rain Program (Bailey, 1998). Most of that "learning", if existent, should have already been conducted before 1995, the beginning of my period of analysis, as the utilities submitted compliance plans to public utilities commissions, and many researchers discussed (e.g.,Rose et al. (1992, 1993)), and regulators implemented, necessary changes to cost-of-service regulation (summarized by Lile and Burtraw (1998)).

⁸For future research, it would be interesting to collect organizational data to formally investigate the exact evolution and effect of this shift in the management practice.

⁹A tax, once in place, is politically costly to change.

tend to locate in less populated areas. This alignment pushes cap-and-trade programs to be more efficient than taxes.

Literature. This paper contributes to the burgeoning literature on firm behavior in new environments, where standard assumptions of firm behavior may not hold. Goldfarb and Xiao (2011) study firms' entry behavior shortly after the passage of the 1996 Telecommunications Act; they find that firms' ability to conjecture opponents' entry decisions improves over time. Covert (2014) studies firms' input choices in the shale gas production using the new technology of hydraulic fracturing; he finds that firms do not use all available information for decision making, and the decisions improve over time only slowly and incompletely. Hortaçsu et al. (2016) study firms' bidding behavior in the deregulated wholesale electricity market in Texas; they find that larger firms bid closer to a Nash equilibrium strategy than do small firms. Doraszelski et al. (2018) study firms' bidding behavior in the newly created frequency response market in the UK; they find that the bids do not resemble Nash equilibrium play initially but stabilize to the latter over time. Thus, the existing research focuses on competitive environments with static decisions. In contrast, this paper focuses on a single-agent environment with dynamic decisions.

Several papers have documented inefficiencies in the Acid Rain Program in its early years. As mentioned above, Montero and Ellerman (1998) find that most of the upward bias in the early predictions of allowance prices come from expectation errors. Carlson et al. (2000) find that a significant proportion of the potential gains from trade is not realized in the early years of the Acid Rain Program. Chan et al. (2017) find that a substantial number of coal units did not choose the least-cost solution to achieve the emission rate they achieved. This paper complements those papers by offering an explanation that formalizes the widely-held view that it took a while for the SO₂ allowance market to take off.

The rest of this paper proceeds as follows. Section 2 describes the institutional details of the Acid Rain Program, compliance options, the allowance price trajectory, and cost-of-service regulation under which private electric utilities operate. Section 3 explains data sources, sample construction, and suggestive evidence of biased beliefs about the future allowance price. Section 4 presents the model, followed by the estimation strategy and results in Section 5. Section 6 presents counterfactual experiment results on the consequences of biased beliefs, and compares cap-and-trade programs with emissions taxes from a dynamic perspective. Section 7 concludes.

2 Institutional Background

The U.S. Acid Rain Program. Legislated in the 1990 Clean Air Act Amendments and administered by the U.S Environmental Protection Agency (EPA), the Acid Rain Program was designed to cut acid rain by reducing SO_2 emissions from electric generating plants to about half their 1980 level. It established the SO_2 allowance market, the first large-scale emission allowance market.

Phase I (1995-1999) of the Acid Rain Program covers 263 largest, dirtiest units, almost all coal-fired, and Phase II (2000-) covers all fossil-fuel-based units exceeding 25MW generating

capacity. Around 180 units from Phase II voluntarily opted into Phase I. I do not differentiate between those voluntary units and the original 263 units, and call them "Phase I units". I use "Phase II units" to refer to the rest of the units, those that have compliance requirement only since 2000; note, however, that Phase I units have compliance obligations in Phase II as well. Figures 2 shows the geographic locations of plants with Phase I units and those with only Phase II units.



Figure 2: Plants with Phase I units (left) and those with only Phase II units (right) as of 2000. Data source: EPA's Air Market Program Data.

Each SO₂ allowance represents one ton of SO₂. The number of allowances allocated to a unit each year equals the product of its average 1985-87 heat input and a target emission rate. The target emission rate is 2.5 lb SO₂ per MMBtu of heat input for Phase I and around 1.2 lb/MMBtu for Phase II. To comply, each unit needs to hold enough allowances in its allowance account by the end of each compliance year to cover its emissions. EPA deducts from that account the number of allowances equal to the emissions, measured by the Continuous Emissions Monitoring (CEM) device installed at flue-gas stacks. Allowances can be traded and banked for future use.

Compliance options. Almost all SO_2 emissions from electric utilities come from units that burn bituminous coal. Utilities can reduce SO_2 emissions by using lower-sulfur bituminous coal (the intensive margin), switching to the ultra-low-sulfur, sub-bituminous coal (the extensive margin), or retrofitting the unit(s) with a flue-gas desulfurization equipment, or a "scrubber".

The sulfur content of bituminous coal ranges from 1.5 to 5 lb/MMBtu. It is mined widely around the Illinois Basin and the Appalachia areas, and is considered high-quality coal because of its high heat content and low moisture content. Bituminous coal has the most developed spot market, and many firms choose lower-sulfur, Central Appalachian bituminous coal to reduce SO_2 emissions (Ellerman et al., 2000). Little capital investment is necessary.

Sub-bituminous coal has very low sulfur content ranging from 0.5 to 1.5 lb/MMBtu. It is predominantly mined in the Powder River Basin in Wyoming. The mine-mouth price of subbituminous coal is generally lower than that of bituminous coal, but the long distance between the Western coal mines and the Eastern coal buyers had traditionally added significant cost to sub-bituminous coal and made it unattractive in the East. Since the deregulation of the railroad industry in the 1980s, however, the rail transport rate has been declining and so has the delivered sub-bituminous coal price. Since sub-bituminous coal has low heat content and high ash content, switching bituminous to sub-bituminous coal requires capital investment that increases the size of coal handling and storage facilities and upgrades the electrostatic precipitator to accommodate high ash throughput. Such investment costs an average of \$45 per kW installed capacity (Ellerman et al., 2000).

A more expensive capital investment is a scrubber. It can remove over 85% of the SO₂ emissions in the flue gas and works best with a sulfur content higher than 5 lb/MMBtu. It requires a large capital cost of around \$250 per kW installed capacity, and three years to install. THe two capital investments, scrubbing and switching to sub-bituminous coal, are effectively mutually-exclusive, because scrubbing works with cheap high-sulfur bituminous coal and dramatically reduces the SO₂ emissions.

Utilities typically planned abatement investments such that they would take effect by the start of the compliance period, for two reasons. First, utilities were required to submit their compliance plans to the public utilities commissions and EPA well before the start of the compliance period. They would propose scrubbing and fuel-switching investment, if any, in those plans. Second, coal contracts, existing or new, are typically structured around the starting or the ending years of the compliance periods. Therefore, proposing scrubbing or switching in the middle of the compliance periods would be neither natural nor practical. Indeed, all scrubbers built for compliance purposes started operating from late 1994 to late 1995, and virtually all fuel-switching investment for compliance purposes took effect around 1995 and around 2000. Scrubbing was irrelevant for firms with only Phase II units, because by the definition of Phase II units, they were already much cleaner than Phase I units and would not find scrubbing worthwhile. Accordingly, I will present in the next section a model of the utilities' behavior between 1995 and 2003, in which the fuel-switching investment is a one-time discrete choice in 1999. I take as given the pre-1995 scrubbing and fuel-switching investments; they reflected firms' beliefs about the allowance price when the allowance price observations were insufficient to generate a meaningful comparison benchmark.

The flexibility of the cap-and-trade program allows firms to buy allowances without cutting back emissions. A firm with excess allowances can sell them for revenue. The EPA does not impose any restriction on who can trade, when, and how. Allowances are labeled with vintages, and in a given compliance year, only the allowances of the current and the prior vintages can be used for compliance (that is, the EPA deducts only the current- and the prior-vintage allowances). This precludes borrowing from the future. For example, if a unit emitted 4000 tons of SO₂ in 1997, the unit account needs to holds at least 4000 allowances of vintage 1997 or prior by the emission deduction deadline. The EPA also holds a small allowance auction in every March to auction off 2.8% of the annually allocated allowances and returns revenue to compliance units.

The allowance price. Trade journals and brokerage firms publish monthly allowance prices. Figure 3 shows the market price of current- and prior-vintage allowances reported by three major sources from 1993-2005. The allowance price began at almost \$200 before Phase I started, but it quickly plunged to as low as \$60 during the first year of the program. It then oscillated for the next decade. I focus on the allowance prices before the end of 2003; the post-2004 prices varied wildly due to unanticipated regulatory uncertainty, which

is beyond the scope of this paper.¹⁰



Figure 3: Monthly market price of current- and prior-vintage allowances, in current US dollars, 1993-2004. Data are provided by Denny Ellerman.

Cost-of-service regulation of private electric utilities. Private electric utilities, responsible for the vast majority of electricity generation before the restructuring of the electricity sector, are subject to cost-of-service regulation by public utilities commissions. Those utilities have obligations to meet all the electricity demand at a pre-determined electricity rate, in return for the monopolistic position in its service area. The electricity rate is set at rate hearings and informal negotiations, such that the utility is reimbursed for prudently-incurred operating costs and earns no more than a fair rate of return on its capital; unanticipated costs, which are not reflected in the electricity rate, are passed on to ratepayers upon the public utilities commission's approval (Joskow, 1972). The electricity restructuring happened around 2000, involving deregulation of the wholesale electricity market and divestiture of generation assets from private electric utilities. Today, around half of the private electric utilities still remain under cost-of-service regulation.

The Acid Rain Program creates new costs and revenues for private utilities. Lile and Burtraw (1998) describe in detail the regulatory treatment of those costs and revenues by each state's public utilities commission. Approved scrubbing or fuel-switching investment would be added to the rate base, earning profits for the utility at the pre-specified rate of

¹⁰The allowance price started its steep ascent in 2004, skyrocketed to \$1600 in 2005, before it gradually declined to almost zero by 2010. A series of unanticipated policy proposals and court actions not directly related to the Acid Rain Program disrupted the allowance market. See Schmalensee and Stavins (2013) for details.

return. Prudently-incurred expenditure on lower-sulfur coal, of acquiring allowances, and the revenue from selling allowances would be passed on to ratepayers.

3 Data

I use detailed data on the U.S. electricity production and compliance with the Acid Rain Program. I compile monthly and annual data from 1995 to 2003 at the unit, plant, and firm levels from publicly available sources. This is the first time that the allowance trading data spanning both Phase I and Phase II have been matched to the electricity production data for academic research.

The main sources of electricity production data are multiple surveys administered by the Energy Information Administration (EIA) and the Federal Energy Regulatory Commission (FERC). Data on plant divestiture are from Cicala (2015), which are also based on EIA data. The Acid Rain Program compliance data are from the Air Markets Program Data (AMPD) system at the Environmental Protection Agency. Monthly current- and prior-vintage allowance prices are from Denny Ellerman, who had collected it from trade journals and brokerage firms. Monthly future-vintage prices are from the online archive of Cantor Fitzgerald / BGC Environmental Brokerage Services, the biggest allowance broker. Appendix A describes data collection and processing details.

The allowance transfers reported to the AMPD contain measurement errors. Firms are required to report to the AMPD only the allowances they intend to use for contemporaneous compliance. The data may therefore miss those allowances that firms have on hand but do not use for compliance right away, or those from transactions to be settled in the future.

Sample construction. The period of analysis is from 1995 to 2003. I focus on private electric utilities that operate un-scrubbed bituminous-coal units subject to at least one phase of the Acid Rain Program and rely little on other SO_2 -emitting fuels. The sample includes 42 firms, accounting for over half of the allowance trading volumes during the period of analysis.

Table 1 summarizes the sample construction process. First, from all electric utilities, I select those with coal units subject to at least one phase of the Acid Rain Program. This is because otherwise the utilities would have trivial compliance burden; coal is responsible for almost all SO_2 emissions from electric utilities.¹¹ Those included utilities, however, can have gas units in addition to coal units. Indeed, gas operations interfered little with coal operations during the period of analysis; gas emits zero SO_2 , and as will be described later in this section, gas prices were too high to affect the electricity generation from coal.

Second, from all electric utilities with coal units subject to at least one phase of the Acid Rain Program, I select those that are privately-owned, which are under cost-of-service regulation. Private electric utilities were responsible for the vast majority of electricity generation prior to the electricity restructuring. Studying the beliefs after the electricity restructuring is for future research, which requires a competitive, instead of single-agent, model.

¹¹In 1985, about 70% of the U.S. SO_2 emissions came from electric utilities, of which 96% were emitted by coal-fired generation units (U.S. Environmental Protection Agency, 1994).

Included are utilities that	Excluded utilities have	# of utilities as of 1995	% trade volume from 1993-2003
have coal units subject to ARP	trivial regulatory burden	161	100
+ are under cost-of-service regulation	different decision- making environment	94	71
+ emit $<10\%$ from non-coal-non-gas source	complicated strategy space	83	68
+ have un-scrubbed bituminous units	easy compliance problem	42	53

 Table 1: Sample construction.

Third, I exclude those utilities with non-coal-non-gas units, such as oil and petroleum coke, that constitute more than 10% of the utility capacity and SO_2 emissions. This is because their compliance problem would then involve choices of non-coal-non-gas fuels, complicating the strategy space. Finally, I exclude those utilities that do not operate un-scrubbed bituminous-coal units. This is because utilities with all coal units scrubbed or burning sub-bituminous coal have trivial compliance burden.¹²

Table 2 reports the sample composition and the emission abatement investment. In the "composition of coal units" panel, the number of private utilities with a given type of unit composition is larger in 1995 than 2000, because of the electricity restructuring that happened between 1998 and 2002. Of those with compliance obligations since 1995, 13 chose to scrub at least one unit, and 11 chose to switch at least one unit to sub-bituminous coal. Of those with compliance obligations since 2000, none chose scrubbing, and five chose switching. Those that neither scrub or switch fuel reduced the sulfur content of bituminous

 $^{^{12}}$ As a result of the sample construction, the analysis in this paper complements the conjecture in Schmalensee and Stavins (2013) about the likely causes of inefficiencies in the Acid Rain Program. Three of the likely causes from Schmalensee and Stavins (2013) are related to the bias toward scrubbers as a compliance option: the encouragement of scrubber installation by bonus allowances under the Acid Rain Program; pre-existing SO₂ regulation such as the New Source Review, which essentially mandated scrubber installation for new plants constructed after 1970; and constraints from state regulation, mostly to protect local high-sulfur coal. However, since my sample starts in 1995, and utilities needed to make scrubber investment two to three years before 1995 (as discussed in the previous section), I take the scrubber investment as given and examine the remaining, un-scrubbed generation capacity. The fourth likely cause from Schmalensee and Stavins (2013) is the policy uncertainty about further SO₂ regulation, discussed in the previous section. I have mitigated its effect by cutting the sample off at the end of 2003, because the allowance market began responding to that policy uncertainty only after 2004. Lastly,Schmalensee and Stavins (2013) point out the lack of information about marginal abatement costs in early years. By 1995, however, the allowance market had operated for about two years, and utilities could use those data to form some beliefs about future market conditions.

Composition of coal units	Number of private utilities			
	as of 1995	as of 2000		
Phase I units only	7	5		
Both Phase I and II units	19	11		
Phase II units only		16		
Total	26	32		
Utility abatement investment	t Number of private utilities			
	comply since 1995	comply since 2000		
Some scrubbing	13	0		
Some switching	10	5		
No investment	3	11		

Table 2: Sample composition and abatement investment.

coal or bought allowances to comply.

Evolution of allowance prices. Table 3 reports various estimates of the stochastic process of the current- and prior-vintage allowance price. Those estimates would typically serve as inputs to the estimation of a dynamic model. The point of this paper, however, is to compare those estimates with the effective beliefs that firms hold. In other words, I do not take the former as inputs to dynamic estimation; instead I back out the latter from data.

Columns (1) and (2) compare allowance price models without and with time trends, when predicting the one-year-ahead allowance price before 2000, using pre-2000 data. The regression equations are:

$$P_t = a + cP_{t-1} + \epsilon_t, \qquad t \le 2000$$

for Column (1) and

$$P_t = a + b(t-1) + (c + d(t-1))P_{t-1} + \epsilon_t, \qquad t \le 2000$$

for Column (2). Adding pre-2000 trends substantially improves the fit. Columns (3) uses the specification of Column (2) to predict the one-year-ahead allowance price after 2000, using all data. Column (4) allows the time trend to continue into Phase II; this worsens the fit.

Figures 4 and 5 plot the observed one-year-ahead allowance price against the predictions from Columns (1) and (2), and from Columns (3) and (4). Consistent with Table 3, the figures show that trends appear important for Phase I allowance price prediction, but not for Phase II. This informs the specification of belief in the dynamic model. Appendix B provides additional data patterns that inform the model in Section 4.

Suggestive evidence of biased beliefs. Figures 6 and 7 plot the distributions of sulfur contents and allowance trades against the allowance price trajectory in Phase I. The allowance price started a sharp increase in 1996. As shown in Figure 4, this increase is predictable based on available data. If firms had reasonable expectations of this increase,

	(1) 93-99	(2) 93-99	(3) 93-02	(4) 93-02
Current price	0.179^{*} (0.0869)	0.871^{***} (0.0796)	0.569^{***} (0.0635)	0.490^{***} (0.0690)
Current price \times phase I year		-0.463^{***} (0.0518)	-0.258^{***} (0.0231)	
Phase I year		81.16^{***} (9.313)	43.21^{***} (3.760)	
Current price \times year				-0.204^{***} (0.0210)
Year				33.53^{***} (3.408)
Constant	$101.6^{***} \\ (14.01)$	-15.39 (12.50)	34.03^{**} (10.48)	$\begin{array}{c} 48.85^{***} \\ (11.46) \end{array}$
N Adj. R-sq	78 0.029	$78 \\ 0.557$	114 0.404	114 0.329

Table 3: Predicting one-year-ahead current- and prior-vintage allowanceprices, in 1995 January dollar.

Robust standard errors are in parenthesis. *p < 0.05, **p < 0.01, ***p < 0.001. The allowance price is deflated using the Urban Consumer Price Index.



Figure 4: Observed one-year-ahead pre-2000 allowance price, and the predictions with v.s. without trends, based on 1993-1999 data, in 1995 January dollar.

Figure 5: Observed one-year-ahead post-2000 allowance price, and the predictions with trends in Phase I but with *v.s.* without trends in Phase II, based on 1993-2002 data, in 1995 January dollar.

they would cut back on SO_2 emissions and buy allowances. By doing so, they would benefit both statically, because the allowance price was low, and dynamically, because the allowance price was going up. Figures 6 and 7 show otherwise.

Figure 6: Distribution of sulfur content in non-scrubbed bituminous coal units during compliance, 1995-1999.

4 Model

This section presents a single-agent dynamic model of a coal-dependent private electric utility subject to the Acid Rain Program. Before the program begins, SO₂ allowances are allocated to the firm. In each year, the firm observes its allowance stock, the allowance price, and the heat input required to meet the electricity demand. It chooses the sulfur content of the bituminous coal and trades allowances. It also faces an investment opportunity in 1999 to switch some plants to sub-bituminous coal. The firm's SO₂ emissions are the product of the sulfur content and the heat input. Emissions and allowance trades affect the next-period allowance stock deterministically. The allowance price and the heat input, exogenous to the firm, evolve stochastically. As is typically assumed, the firm's belief about the heat input coincides with its stochastic process. The novelty is that the firm's belief about the future allowance price is flexibly specified.

4.1 Allowance Price Belief

When specifying the allowance price belief, it is important to allow for non-stationarity, particularly in Phase I. Indeed, Phase I is the first five years of the unprecedented allowance market, when the price formation process may still be evolving. Thus, I assume that the firm believes in year $t \in \{1995, 1996, \ldots, 2000\}$ that the allowance price in year s > t evolves

Figure 7: Distribution of net purchase of current-vintage allowances by utilities during compliance, 1995-1999.

according to the following process:

$$P_s = b_1 + b_2 \times \min\{s - 1, 2000\} + (b_3 + b_4 \times \min\{s - 1, 2000\})P_{s-1} + \epsilon_s, \tag{1}$$

where ϵ_s is an i.i.d. normal error with mean zero and standard deviation b_5 . Again, the values for the parameters $(b_1, b_2, b_3, b_4, b_5)$, which capture the firm's belief about the allowance price, will come from structural dynamic estimation in Section 5, rather than from the estimation of the stochastic process in Section 3.

Thus, the firm forecasts future allowance prices based on historical allowance prices via the slope coefficients, and the conditions of related markets (*e.g.*, coal and electricity markets) via the intercepts. Importantly, the firm believes that the allowance price process is not stationary until Phase II; Equation (1) has time trends in both the slope and the intercept only if the forecast year $s \leq 2000$. I impose stationarity on the Phase II forecasts to capture the perception that the allowance price formation process will settle down by Phase II. This also facilitates the dynamic programming of the infinite-horizon Phase II problem, as it requires stationary state transitions.

The parameters of interest are $(b_1, b_2, b_3, b_4, b_5)$. In Section 5, I will present both the pooled estimates, using all data between 1995 and 1999, and the year-specific estimates, using the data in each year separately. The former corresponds to rational expectations and the latter adaptive-learning beliefs. For the latter, I adopt the standard assumption for the adaptive estimation that the firm is myopic about the possible future belief changes; it does not internalize those possible changes when it makes its decisions today.

In year $t \in \{2000, \ldots, 2003\}$ in Phase II, the firm believes that the allowance price in year s > t simply follows the stationary process:

$$P_s = b_1' + b_3' P_{s-1} + \epsilon_s \tag{2}$$

where ϵ_s is an i.i.d. normal error with mean zero and standard deviation b'_5 . Thus, in Phase II, the firm expects the current allowance price, P_{s-1} , to effectively aggregate available market information; and the market fundamentals, $b_1 + \epsilon_s$, to contribute to the future allowance price in a constant manner on average. The parameters of interest are (b'_1, b'_3, b'_5) . As before, I will present both the pooled estimates, using all data between 2000 and 2003, and the year-specific estimates.

4.2 The Phase II Problem

In each $t \in \{2000, 2001, \ldots, \infty\}$, firm *i* observes its allowance stock W_{it} , the allowance price P_t , and its heat input H_{it} . It then chooses the net allowance purchase a_{it} and the sulfur content x_{it} . It incurs net allowance expenditure $A(a_{it}; P_t)$ and coal expenditure $C_i(x_{it}; H_{it})$.

I assume that the firm's static payoff depends on those expenditures via "internalization functions" ϕ_A and ϕ_C :

$$\pi_i(a_{it}, x_{it}; P_t, H_{it}) = \phi_A[A(a_{it}; P_t)] + \phi_C[C_i(x_{it}; H_{it})], \tag{3}$$

where the internalization functions translate the monetary value of those expenditures into what the firm deems payoff. As an example, a firm that is not subject to cost-of-service regulation and produces output D would have $\phi_A(A) = -A$ and $\phi_C(C) = -C + P(D)D$, where P is the inverse demand function. That firm would internalize 100% of the allowance and coal expenditures is because the revenue is independent of them. However, the firms in this paper are subject to cost-of-service regulation, and their revenue depends on their cost. As described in Section 2, the public utilities commission passes the costs that those firms prudently incurs to ratepayers. As a result, those firms may not internalize every dollar spent on allowances and coal.

The internalization functions flexibly captures the distortionary effect of cost-of-service regulation on the firm's internalized payoff. For example, if the firm pays A dollars for the allowances, and the public utilities commission approves pass-through of allowance expenditures with a probability of 80%, then the firm expects to receive 0.8A - A = -0.2A as the static payoff from allowance trading. Then, $\phi_A = -0.2A$. Alternatively, if the firm pays a high C for coal due to an unexpectedly high coal price, then even if the public utilities commission almost always approves pass-through of coal expenditures, the firm may offer to absorb 0.1C to build good will with the regulator. The firm receives 0.9C - C = -0.1C as the static payoff from burning coal. Then, $\phi_C = -0.1C + PD$, where P is the part of the electricity rate that generates returns on existing capital investment. For the rest of the analysis, I omit PD from the payoff because both P and D are generally out of the firm's control and therefore do not affect the firm's choices of allowance trades and sulfur contents.

The ability to save allowances for future use creates dynamics. In each year, the firm starts with its stock of allowances, W_{it} , receives free allowances $alloc_i$, trades a_{it} allowances, and submits a number of allowances equal to the SO₂ emissions, $x_{it}H_{it}$. Thus, the evolution of the allowance stock is deterministic and endogenous:

$$W_{i,t+1} = W_{it} + alloc_i + a_{it} - x_{it}H_{it}.$$

The evolution of the allowance price P follows the allowance price beliefs in the previous section. The evolution of heat input H_i is Markovian and exogenous.

The Bellman's Equation for firm i's Phase II problem is:

$$V_{i}(W_{i}, P, H_{i}) = \max_{\substack{x_{i} \in X\\a_{i} \ge x_{i}H_{i} - W_{i} - alloc_{i}}} \{\phi_{A}[A(a_{i}; P)] + \phi_{C}[C_{i}(x_{i}; H_{i})] + \beta \int V_{i}(W_{i} + alloc_{i} + a_{i} - x_{i}H_{i}, P', H'_{i})dF_{P}(P'|P)dF_{H_{i}}(H'_{i}|H_{i})\},$$

where F_P is the allowance price belief of Equation (2), and F_{H_i} is the transition of heat input.

The first constraint, $x_i \in X$, requires that the sulfur content be chosen from the physical range of sulfur contents of bituminous coal. The second constraint, the "compliance constraint", $a_i \ge x_i H_i - W_i - alloc_i$, requires that enough allowances be held to cover the SO₂ emissions each year; after rearranging it becomes $W'_i \equiv W_i + alloc_i + a_i - x_i H_i \ge 0$.

The first-order condition with regard to the sulfur content, if interior, is:

$$\frac{\partial \phi_C(C^*)}{\partial C} \frac{\partial C_i(x_i^*; H_i)}{\partial x_i} / H_i - \beta \mathbb{E} \frac{\partial V((W_i')^*, P', H_i')}{\partial W_i'} - \mu^* = 0,$$
(4)

and the first-order condition with regard to the net allowance purchase is:

$$\frac{\partial \phi_A(A^*)}{\partial A} \frac{\partial A(a_i^*; P)}{\partial a_i} + \beta \mathbb{E} \frac{\partial V((W_i')^*, P', H_i')}{\partial W_i'} + \mu^* = 0,$$
(5)

where μ^* is the Lagrange multiplier associated with the compliance constraint.

The first-order conditions explain optimal behavior intuitively.¹³ Equation (4) shows that the cost of marginally lower-sulfur coal is justified by two benefits: the discounted expected marginal value of allowances, and the shadow price of the compliance constraint. Using a lower sulfur content both avoids drawing from the allowance stock and relaxes the compliance constraint. Equation (5) shows that the cost of marginally more allowances is justified in the same way; buying more allowances both adds to the allowance stock and relaxes the compliance constraint.

Equations (4) and (5) together imply that, at the optimum, the marginal cost of lowering the sulfur content in coal equals the marginal cost of acquiring additional allowances. This is regardless of whether the compliance constraint binds or not. Indeed, sulfur content reductions and allowances are perfect substitutes. In Section 5, I use this equality to estimate the payoff parameters without solving the dynamic programming problem.

The belief over the future allowance price affects behavior via the discounted expected marginal value of allowances. For example, the higher the expected future allowance price, the higher the discounted expected marginal value of allowances, leading to a lower sulfur content and a larger net allowance purchase.

¹³Appendix C proves that the value function $V(W_i, P, H_i)$ is increasing and concave in the allowance stock, as long as the internalization function, ϕ_A , is decreasing in the allowance expenditure, and the allowance expenditure function, A, is increasing and convex in the net allowance purchase. Thus, those first-order conditions are sufficient for optimality.

Allowing for future-vintage allowance trading. The model so far has focused on the current- and prior-vintage allowances. It is straightforward to allow for the next-vintage allowances. The net allowance expenditure becomes $A(a_{it}, b_{it}; P_t, \psi(P_t))$, where b_{it} is the next-vintage net allowance purchase, and ψ maps the current-vintage allowance price to the next-vintage.

The Bellman's Equation for firm i's Phase II problem that includes next-vintage allowance trading is:

$$V_{i}(W_{i}, P, H_{i}) = \max_{\substack{x_{i} \in X\\a_{i} \ge x_{i}H_{i} - W_{i} - alloc_{i}\\a_{i}^{nv} \ge -alloc_{i}}} \{\phi_{A}[A(a_{i}, a_{i}^{nv}; P)] + \phi_{C}[C_{i}(x_{i}; H_{i})] + \beta \int V_{i}(W_{i} + alloc_{i} + a_{i} + a_{i}^{nv} - x_{i}H_{i}, P', H'_{i})dF_{P}(P'|P)dF_{H_{i}}(H'_{i}|H_{i})\}$$

The additional constraint $a_i^{nv} \ge -alloc_i$ says that the firm cannot sell more next-vintage allowances than its allocation. The compliance constraint remains intact, but the allowance stock next year now includes a_i^{nv} . Indeed, the t+1 vintage allowances purchased in t cannot be used for compliance until year t+1.

Appendix B has shown that almost all transactions between the period of analysis concern the current-year, the prior-year, and the next-year vintage. Incorporating allowance trades of further vintages would incur huge computation burden, as it would require additional state variables, one for the allowance stock of each vintage. In any case, the prices of allowances of further vintages are not always available.

4.3 Fuel-Switching Investment

I model fuel-switching investment as a discrete-choice problem. Fuel-switching investment typically involves all bituminous units within a plant, and the set of plants is discrete. Since both data and intuition suggest that switching younger plants makes a better investment, the choice set includes switching no plant, switching the youngest plant, switching the two youngest, up to switching the four youngest. I cap the size of the choice set at five to avoid computing an impractical number of Phase II dynamic programming problems. In fact, the majority of firms in the data have no more than four coal plants, and the nine exceptions are located so far away from sub-bituminous coal sources that switching more than five coal plants would be unrealistic.¹⁴ Let firm *i*'s choice set be J_i , in which each choice *j* is characterized by the capacity to be switched, k_j .

The timing for the 1999 investment problem is as follows. First, the firm observes the states (W_i, P, H_i) and chooses the next-vintage net allowance purchase; if the firm has Phase I units, it also chooses the sulfur content and the current-vintage net allowance purchase. Second, private shocks to the cost of each investment option are realized. Third, the firm chooses an investment option. The investment takes effect in 2000.

¹⁴The nine firms with more than five coal plants are Carolina Power and Light, Detroit Edison, Duke Power, PSI Energy, South Carolina Electric and Gas, Virginia Electric and Power, Southern Company, and American Electric Power. Only Southern Company switched fuel, and it switched only one plant.

To compute the value of each investment option, the firm uses the following information: 1) the heat input rate of sub-bituminous units, $\bar{\alpha}_i$, that translates capacity to heat input; 2) the sulfur content of sub-bituminous coal, \bar{x}_i , that translates heat input to emissions; and 3) the sub-bituminous coal price, \bar{p}_i . The heat input rate of sub-bituminous units is assumed constant (over time) because, conditional on switching, the sub-bituminous coal typically has lower marginal cost than bituminous coal, and therefore the switched capacity will be dispatched first. The sulfur content of sub-bituminous coal is assumed constant, because most sub-bituminous coal contracts are long-term with pre-specified sulfur contents, and sub-bituminous coal has a narrow range of sulfur contents to start with. The sub-bituminous coal price is assumed constant because of long-term contracts.

The value of investment option j to firm i after it takes effect is:

$$\begin{aligned} V_{i}^{j}(W_{i}, P, H_{i}) &= \max_{\substack{x_{i} \in X \\ a_{i} \geq x_{i} \max\{H_{i} - \bar{\alpha}_{i}k_{j}, 0\} + \bar{x}_{i} \min\{\bar{\alpha}_{i}k_{j}, H_{i}\} - W_{i} - alloc_{i}}} \{\phi_{A}[A(a_{i}, a_{i}^{nv}; P)] + \phi_{C}[C_{i}(x_{i}; H_{i})] \\ &+ \beta \int V_{i}(W_{i} + a_{i} + a_{i}^{nv} + alloc_{i} - x_{i} \max\{H_{i} - \bar{\alpha}_{i}k_{j}, 0\} - \bar{x}_{i} \min\{\bar{\alpha}_{i}k_{j}, H_{i}\}, P', H'_{i}) \\ &\times dF_{P}(P'|P)F_{H_{i}}((H_{i})'|H_{i})\}, \end{aligned}$$

where $x_i \max\{H_i - \bar{\alpha}_i k_j, 0\}$ is the emission from bituminous coal units, and $\bar{x}_i \min\{\bar{\alpha}_i k_j, H_i\}$ from sub-bituminous.

If we assume that the investment cost shocks are i.i.d. logit errors, the expected value of the investment opportunity to firm i in 1999 before the cost shocks realize is:

$$\log(\sum_{j=1}^{|J_i|} \exp(-\phi_K(k_j c^k) + \beta \mathbb{E}(V_i^j(W_{i,2000}, P_{2000}, H_{i,2000}) | W_{i,1999}, P_{1999}, H_{i,1999}))),$$

up to a constant, where c^k is the unit capital cost of fuel switching and ϕ^K is the internalization function for capital expenditure. The constant is omitted because it does not affect behavior. This value as a function of the 1999 states will be the terminal value function for the finite-horizon Phase I problem, to be introduced below.

4.4 Finite-Horizon Phase I

The Bellman's Equation for the 1999 problem of firm i is:

$$V_{i}^{1999}(W_{i,1999}, P_{1999}, H_{i,1999}) = \max_{\substack{x_{i} \in X \\ a_{i} \ge x_{i}H_{i,1999} - W_{i,1999} - alloc_{i,1999} \\ a_{i}^{nv} \ge -alloc_{i,2000}}} \{\phi_{A}[A(a_{i}, a_{i}^{nv}; P_{1999})] \\ + \phi_{C}[C_{i}(x_{i}; H_{i,1999})] + \log[\sum_{j=1}^{|J_{i}|} \exp(-\phi_{K}(k_{j}c^{k}) \\ + \beta V_{i}^{j}(W_{i,1999} + alloc_{i,1999} + a_{i} + a_{i}^{nv} + x_{i}H_{i,1999}, P_{2000}, \tilde{H}_{i,2000}) \\ \times dF_{P}(P_{2000}|P_{1999})dF_{\tilde{H}_{i}}(\tilde{H}_{i,2000}|g_{i}(H_{i,1999})]\},$$

$$(6)$$

where F_P is the Phase I allowance price transition in Equation (1) with s = 2000, and $F_{\tilde{H}_i}$ is the transition of the heat input to Phase II units; the function g_i converts the heat input to Phase I units to that to Phase II ones.¹⁵

The Bellman's Equation for the year $t \in \{1995, \ldots, 1998\}$ problem is:

$$\begin{aligned} V_i^t(W_{i,t}, P_t, H_{i,t}) &= \max_{\substack{x_i \in X\\a_i \ge x_i H_{i,t} - W_{i,t} - alloc_{i,t}\\a_i^{nv} \ge -alloc_{i,t+1}}} \{\phi_A[A(a_i, a_i^{nv}; P_t)] + \phi_C[C_i(x_i; H_{i,t})] \\ &+ \beta \int V_i^{t+1}(W_{i,t} + alloc_{i,t+1} + a_i^{nv} - x_i H_{i,t}, P_{t+1}, H_{i,t+1}) \\ &\times dF_P(P_{t+1}|P_t)F_H(H_{i,t+1}|H_{i,t}), \end{aligned}$$

where F_P is the Phase I allowance price transition in Equation (1) with s = t + 1, and F_{H_i} is the transition of heat input to Phase I units.

Modifications for firms that start compliance in 2000. If the firm is subject to only Phase II, the 1995-1998 problems are irrelevant. In its 1999 problem, only the next-vintage allowance trade and the fuel-switching investment remain as choices, and g_i is the identity map.

4.5 Identification

The parameters to be estimated are:

- 1. coal expenditure parameters in the functions C_i , for all i;
- 2. allowance expenditure parameter in the function A;
- 3. heat input transition parameters in F_{H_i} and $F_{\tilde{H}_i}$, for all *i*;
- 4. internalization parameters in the functions ϕ_A, ϕ_C ;
- 5. internalization parameter in the function ϕ_K ; and
- 6. belief parameters $(b_1, b_2, b_3, b_4, b_5)$.

Parameters 1-4 are identified without solving Bellman's Equations. The coal expenditure parameters represent how the bituminous coal price depends on the sulfur content. They are identified from within-plant variations in the price charged for its monthly shipments of bituminous coal with varying sulfur contents. The identification assumption is that the sulfur content is orthogonal to the unobserved price component, after controlling for observed coal characteristics such as ash content, BTU content, and source county.

The allowance expenditure parameter and the internalization parameters in ϕ_A, ϕ_C are identified based on the equality between the marginal cost of lower sulfur content and the

¹⁵The use of g_i avoids tracking the heat input to Phase II units as an additional state variable.

marginal cost of allowances from Equations (4) and (5). The allowance expenditure parameter concerns how the marginal cost rises with the transaction volume.¹⁶ As will be discussed in Section 5, this equality is subject to the measurement error in the reported allowance trades. I correct for the measurement error to obtain unbiased estimates of those parameters.

The heat input transition parameters are identified under standard identification assumptions of first-order autoregressive models.

Identifying Parameters 5 and 6 requires solving the dynamic problem. Exogenous variations in the allowance price and the heat inputs lead to variations in the expected marginal value of allowances both within and across firms. Parameters 5 and 6 are separately identifiable because they have different implications for optimal behavior. Parameter 5, or the capital internalization parameter, affects the investment choice, and the allowance and sulfur content choices, in opposite directions. Parameters 6, or the belief parameters, push all choices in the same direction. For example, higher capital internalization implies less switching and more reliance on allowances and lower-sulfur bituminous coal, while a higher allowance price to be expected tomorrow leads to more switching, more allowance purchase, and lower sulfur content.

5 Estimation

The estimation strategy proceeds in two steps. First, I estimate the parameters that do not rely on the solution to the dynamic problem. Those parameters are: 1) the coal expenditure parameters for each firm, estimated using fixed-effect regressions; 2) the heat input transition parameters for each firm, estimated using time-series regressions; and 3) the allowance expenditure parameter; and 4) the internalization parameters for allowance and coal expenditures. The equality derived from the two first-order conditions of the solution to the dynamic model informs the estimation of parameters 3) and 4), which uses an ordinary least squares regression with measurement error correction.

The second step is a nested-fixed-point-type algorithm to estimate the belief parameters and the capital internalization parameter. This step requires the solution to the dynamic problem.¹⁷ The inner loop is the first empirical application of the Relative Value Function Iteration and Endogenous Value Function Iteration methods (Bray, 2017b,a), which vastly accelerate dynamic programming. The outer loop uses a maximum simulated likelihood estimator, searching for the parameter values that best rationalize the observed behavior. The simulation numerically integrates out the measurement errors in the allowance stock state variable induced by the measurement errors in the net allowance purchase.

¹⁶The assumption of an increasing marginal cost is necessary to prevent allowance trades from going to infinity, which is never observed in the data. Its interpretation can be the financial transaction cost, the firm's budget constraint, the constraint from the public utilities commissions, etc.

¹⁷I do not use the BBL estimation approach (Bajari et al., 2007), because Phase I has a finite horizon, and my data do not permit reliable estimation of year-specific policy functions.

5.1 Coal Expenditure Parameters

The delivered bituminous coal price of shipment k in month-year t to plant j is:

$$p_{jkt}^{bit} = \gamma_0^{bit} + \gamma_1^{bit} x_{jkt} + \gamma_2^{bit} x_{jkt}^2 + \gamma_Z^{bit} Z_{jkt} + dummies + \epsilon_{jkt}^{bit}$$
(7)

where x is the sulfur content, Z is a flexible control variable vector, and *dummies* is a dummy variable vector. The control variables include the heat content, the ash content, the distance between the plant and the source county, and their interactions. The quadratic specification of the sulfur content is consistent with Kolstad and Turnovsky (1998).

The parameters of interest are γ_1^{bit} and γ_2^{bit} . They measure how the delivered coal price changes with the sulfur content of coal. Table 4 reports the estimates of γ_1^{bit} and γ_2^{bit} at various specifications. I use the estimates from Column (4) for the rest of the analysis.

The coal expenditure of firm i in the Phase I problem is:

$$C_i(x_{i,t}; H_{i,t}) = H_{i,t}(\gamma_1^{bit} x_{i,t} + \gamma_2^{bit} x_{i,t}^2)$$

where x is the sulfur content, and H is the heat input. I omit from the coal expenditure the non-sulfur components of the coal price function (7); they do not affect the firm's choices in Phase I.

The coal expenditure of firm i with capacity k_j switched to sub-bituminous coal in the Phase II problem is:

$$C_i(x_{i,t}; H_{i,t}) = \max\{H_i - \bar{\alpha}_i k_j, 0\} (\gamma_1^{bit} x_{i,t} + \gamma_2^{bit} x_{i,t}^2 + \bar{\gamma}_i^{bit}) + \min\{H_i, \bar{\alpha}_i k_j\} \bar{\gamma}_i^{sub}$$

where $\bar{\alpha}_i k_j$ is the heat input to the sub-bituminous units switched to, in which $\bar{\alpha}_i$ is the heat input rate of those units. Unlike in the Phase I problem, the coal expenditure in the Phase II problem includes the non-sulfur components in the bituminous coal price, summarized by $\bar{\gamma}_i^{bit}$. The non-sulfur price components affect the value of switching capacity k_j , which in turn affects the terminal value functions for Phase I. I construct the non-sulfur bituminous coal price $\bar{\gamma}_i^{bit}$ by taking the heat-input-weighted average of $(p_{jkt}^{bit} - \gamma_1^{bit} x_{jkt} - \gamma_2^{bit} x_{jkt}^2)$ over the bituminous coal shipments received by plants operated by firm *i* during Phase II.

The coal expenditure in the Phase II problem also includes sub-bituminous coal expenditure. Let $\bar{\gamma}_i^{sub}$ denote the sub-bituminous coal price. For firm *i* that receives sub-bituminous coal shipments in Phase II, I construct $\bar{\gamma}_i^{sub}$ by taking the heat-input-weighted average of p_{jkt}^{sub} over the sub-bituminous coal shipments received by plants operated by firm *i* during Phase II. For firm *i* that does not receive sub-bituminous coal shipments in Phase II, I construct the counterfactual $\bar{\gamma}_i^{sub}$ as follows. First, I conduct a sub-bituminous coal price regression in the style of Equation (7), using data from plants that receive sub-bituminous coal shipments. Second, I construct the predictors for plants that do not receive such shipments. The sulfur content would be the industry-average sulfur contents used by sub-bituminous coal price during; the other predictors would be the average values of the sub-bituminous coal produced by the coal county that yields the lowest predicted coal price for firm *i*.

	(1)	(2)	(3)	(4)
Sulfur content	-6.477**	-8.350***	-7.305***	-7.693***
(lb/MMBtu)	(2.477)	(2.249)	(1.572)	(1.175)
Sulfur $content^2$	0.215	0.362	0.584**	0.641**
	(0.332)	(0.292)	(0.217)	(0.199)
Ash content			-3.458***	-4.026***
(lb/MMBtu)			(0.703)	(0.656)
Heat content			0.00339**	-0.00175
(BTU/lb)			(0.00115)	(0.000900)
Heat×ash			0.000328***	0.000333***
			(0.0000793)	(0.0000695)
Distance to coal county			0.0893***	-0.00422
(km)			(0.0112)	(0.0167)
$Distance^2$			-0.0000514***	0.00000372
			(0.0000106)	(0.0000120)
Constant	118.9***	153.1***	91.12***	202.8***
	(4.111)	(4.655)	(14.80)	(13.53)
Month FE		Y	Y	Y
Year FE		Y	Υ	Υ
Coal county FE			Υ	Υ
Plant FE				Y
N	145456	145456	144830	144830
Adj. R-sq	0.045	0.442	0.545	0.624

Table 4: Delivered bituminous coal price regressions, in 1995 January US cent/MMBtu, 1995-2004.

Standard errors clustered by utility are in parenthesis. *p < 0.05, **p < 0.01, ***p < 0.001. The delivered coal price is deflated using the Producer Price Index for Crude Energy Materials, retrieved from FRED, Federal Reserve Bank of St. Louis, at https://fred.stlouisfed.org/series/PPICEM on April 26, 2017.

	Units subject to Phase I	Units subject to Phase II
Lagged heat input	0.918***	0.967***
	(0.123)	(0.127)
Constant	66.47 (93.21)	54.45 (144.4)
RMSE N Adj. R-sq	38.37 14 0.735	45.49 14 0.837

Table 5: Heat input regressions for the Southern Company's bituminous coal units, in million MMBtu, 1990-2004.

Robust standard errors are in parenthesis. *p < 0.05, **p < 0.01, ***p < 0.001.

5.2 Heat Input Transitions

The heat input to firm i's bituminous coal units in year t is:

$$H_{i,t} = \gamma_{0,i}^{H} + \gamma_{1,i}^{H} H_{i,t-1} + \epsilon_{i,t}^{H}, \tag{8}$$

where $\epsilon_{i,t}^H \sim \mathcal{N}(0, (\sigma_i^H)^2))$. The parameters of interest are $\gamma_{0,i}^H, \gamma_{1,i}^H$, and σ_i^H . The identification assumptions are:

$$\mathbb{E}(H_{i,t-1}\epsilon_{i,t}^{H}) = 0, \qquad \forall t,$$
$$\mathbb{E}(\epsilon_{i,t}^{H}\epsilon_{i,\tau}^{H}) = 0, \qquad t \neq \tau$$

I estimate the heat input transitions separately for the Phase I units only, and the Phase I and Phase II units taken together. Table 5 reports the estimates of $\gamma_{0,i}^H$, $\gamma_{1,i}^H$, and σ_i^H for the Southern Company as an example.

5.3 Allowance Expenditure Parameter, Internalization Parameters for Allowance and Coal Expenditures

As shown in Section 4, the first-order conditions of the dynamic problem indicate that the marginal cost of lowering the sulfur content equals the marginal cost of acquiring additional current-vintage allowances, as long as the sulfur content remains interior. I use this equality to estimate the allowance expenditure parameters and the internalization parameters for allowance and coal expenditures.

The allowance expenditure is:

$$A(a_{i,t}, a_{i,t}^{nv}; P_t) = a_{i,t}P_t + a_{i,t}^{nv}\delta P_t + \theta_0(a_{i,t}^2 + (a_{i,t}^{nv})^2)$$

where a is the current-vintage net allowance purchase, a^{nv} is the next-vintage net allowance purchase, P is the allowance price, δ is the ratio of the next-vintage allowance price over the current-vintage, and θ_0 is a penalty parameter for large transaction volumes. The price ratio δ is observed in the data. The penalty parameter θ_0 is the allowance expenditure parameter to be estimated.

The internalization functions for the allowance and coal expenditures are:

$$\phi_A(A) = -\theta_A A,$$

$$\phi_C(C) = -\theta_C C,$$

where θ_A and θ_C are the internalization parameters to be estimated. The equality between the marginal costs thus implies:

$$\theta_C(-\gamma_1^{bit} + 2\gamma_2^{bit}x_{i,t}) = \theta_A(P_t + 2\theta_0 a_{i,t}).$$
(9)

Since we can only identify the ratio between the two internalization parameters, I normalize $\theta_C = 1$. Let $\theta_a = \theta_0 \theta_A$. Equation (9) thus becomes:

$$-\gamma_1^{bit} + 2\gamma_2^{bit} x_{i,t} = \theta_A P_t + 2\theta_a a_{i,t}.$$
(10)

As described in Section 2, the allowance trades contain measurement errors. Private communication with the EPA suggests that around 25% of the allowance trades may not have been reported. I use this information to construct a model for the measurement error in the allowance trades. First, I multiply the total allowance trading volume with 25% to approximate the total number of allowance trades that are not reported. Second, I normalize this number by the total allowance allocation, arriving at the per-allocation *expected* measurement error, e = 0.225. The current-vintage measurement error model is:

$$a_{i,t} = a_{i,t}^* + alloc_i \eta_{i,t},\tag{11}$$

where a^* is the true current-vintage net allowance purchase, *alloc* is the firm-specific allowance allocation, and $\eta_{i,t}$ is an i.i.d. normal, per-allocation measurement error with mean zero and standard deviation σ^{η} .¹⁸ I calibrate σ^{η} such that the conditional expectation of positive (and symmetrically, negative) per-allocation measurement error, or $\sigma^{\eta} \frac{\sqrt{2}}{\sqrt{\pi}}$, equals e.

After we replace a with a^* in Equation (10) and substituting in Equation (11), Equation (10) becomes:

$$\gamma_1^{bit} + 2\gamma_2^{bit} x_{i,t} = \theta_A P_t + 2\theta_a a_{i,t} - (2\theta_a alloc_i)\eta_{i,t}.$$

Let $Y_{it} = -\gamma_1^{bit} + 2\gamma_2^{bit} x_{i,t}$. Then:

$$\tilde{Y}_{it} = \theta_A \tilde{P}_t + 2\theta_a \tilde{a}_{i,t} - 2\theta_a \eta_{i,t}, \qquad (12)$$

where $\tilde{Y}, \tilde{P}, \tilde{a}$ are *alloc*_i-normalized versions of Y, P, a. Since $\eta_{i,t}$ is correlated with $\tilde{a}_{i,t}$, which in turn is correlated with \tilde{P}_t , the ordinary least squares estimates of θ_A and θ_a will both be biased.

¹⁸It might appear natural to use a multiplicative measurement error model. For example, $a_{i,t} = a_{i,t}^* \eta_{i,t}$. Many observations of the allowance trades are zero, which, under such a model, would imply either $a_{i,t}^* = 0$ or $\eta_{i,t} = 0$; the former is inconsistent with the possibility that a firm trading a positive number of allowances simply does not report the transaction, and the latter does not pin down a^* .

	OLS	Corrected
Internalization of allowance expenditure (θ_A)	0.636^{***} (0.00986)	$\begin{array}{c} 0.636^{***} \\ (0.00977) \end{array}$
Quadratic allowance cost (θ_a)	$\begin{array}{c} 0.0000562 \\ (0.000645) \end{array}$	$\begin{array}{c} 0.0000572^{*} \\ (0.0000288) \end{array}$
N Adj. R-sq	150 0.968	

Table 6: Estimates of θ_A , the allowance internalization parameter, and θ_a , the quadratic allowance cost parameter.

Standard errors are in parenthesis. p < 0.05, p < 0.01, p < 0.01.

Absent an instrument for a, I correct for the biases in those estimates induced by the measurement error:

$$\begin{aligned} \theta_a^{corrected} &= \theta_a^{OLS} / (1 - \frac{var(\eta)var(\tilde{P})}{var(\tilde{a})var(\tilde{P}) - cov(\tilde{a}, \tilde{P})^2}) \\ \theta_A^{corrected} &= \theta_A^{OLS} + 2\frac{cov(\tilde{a}, \tilde{P})}{var(\tilde{P})} (\theta_a^{OLS} - \theta_a^{corrected}) \end{aligned}$$

Table 6 reports the estimates of θ_a and θ_A and their corrections. To ensure that the sulfur content is interior, as required by the equality between the marginal costs, I restrict the sample to those observations with \tilde{Y} within the 5% and 95% percentiles.

5.4 Inner Loop

The inner loop is a three-dimensional continuous-state mixed-continuous-action dynamic programming problem. The state variables are the allowance stock, the allowance price, and the heat input. The choice variables are annual continuous choices of sulfur content and allowance trades, and the 1999 discrete choice of fuel-switching investment.

Estimation of beliefs held during Phase II requires dynamic programming of the Phase II problem only. Given a parameter vector, for each firm, I use the Endogenous Value Function Iterations method (Bray, 2017a) to rapidly solve for the Phase II value function conditional on the fuel-switching investment chosen in 1999. The solution method exploits the fact that the value function *shape* over the endogenous states in each exogenous state, rather than the value function *levels* in all exogenous and endogenous states, matters for behavior.

Estimation of beliefs held during Phase I requires dynamic programming of both the Phase I (including the investment problem in 1999) and Phase II problems. Given a parameter vector, for each firm, I first use the Relative Value Function Iterations method (Bray, 2017b) to quickly solve for the Phase II value functions given each fuel-switching investment option. The 1999 investment problem requires as inputs the *levels* of the Phase II value functions, which can be backed out from the relative value functions. Indeed, the Relative

Value Function Iterations method focuses on the value function shape over all exogenous and endogenous states, and therefore the resulting relative value function is the full value function shifted by a constant. Having backed out the value functions associated with each investment option, I use backward induction to solve for the value functions specific to each year in Phase I.¹⁹ Appendix D describes further details on the inner loop.

5.5 Outer Loop

The outer loop looks for the parameter values that maximize the simulated likelihood. The likelihood is induced by the i.i.d. distribution of measurement errors in the currentand next-vintage allowance trades and the sulfur content; firms are not required to report transactions of allowances not used for contemporaneous compliance, and the precision of the self-reported sulfur content is not satisfactory.²⁰ The measurement error in the allowance trades induces measurement error in the allowance stock state variable, and the individual likelihood needs to be integrated over possible true allowance stocks. Because the measurement error in the allowance stock affects behavior nonlinearly via the dynamic problem, the likelihood is simulated.

Objective and algorithm. To simulate the likelihood, I first simulate many paths of allowance stock states. The measurement errors in net allowance purchases induce the measurement errors in the allowance stock:

$$\begin{split} W_{i,t+1} &= W_{i,t}^* + alloc_i + a_{i,t} + a_{i,t}^{nv} - deduct_{i,t} \\ &= W_{i,t}^* + alloc_i + a_{i,t}^* + alloc_i \eta_{i,t}^a + (a_{i,t}^{nv})^* + alloc_i \eta_{i,t}^{a^{nv}} - deduct_{i,t} \\ &= (W_{i,t}^* + alloc_i + a_{i,t}^* + (a_{i,t}^{nv})^* - deduct_{i,t}) + alloc_i \eta_{i,t}^a + alloc_i \eta_{i,t}^{a^{nv}} \\ &= W_{i,t+1}^* + alloc_i \eta_{i,t}^a + alloc_i \eta_{i,t}^{a^{nv}}, \end{split}$$

subject to the compliance constraint $W_{i,t}^* + alloc_i + a_{i,t}^* - deduct_{i,t} \ge 0$ and the no-shorting constraint $(a_{i,t}^{nv})^* + alloc_i \eta_{i,t}^{a^{nv}} \ge -alloc_i$. I simulate many paths of $W_i = (W_{i,1995}, W_{i,1996}, \ldots)$ for each firm *i* that is subject to both phases as follows:

- 1. initialize t = 1995, and let $W_{i,1995}^* = W_{i,1995}$;
- 2. draw η^a from its distribution, and compute $a_{i,t}^* = a_{i,t} alloc_i \eta^a$;

3. if
$$W_{i,t}^* + alloc_i + a_{i,t}^* - deduct_{i,t} \ge 0$$
, accept η^a as $\eta_{i,t}^a$; otherwise, go back to Step 2;

- 4. draw $\eta^{a^{nv}}$, and compute $(a_{i,t}^{nv})^* = a_{i,t}^{nv} alloc_i \eta^{a^{nv}}$;
- 5. if $b_{i,t}^* > -alloc_i$, accept $\eta^{a^{nv}}$ as $\eta_{i,t}^{a^{nv}}$; otherwise, go back to Step 4;

¹⁹I implement the inner loop in the AMPL language (Fourer et al., 2003) with the KNITRO (Byrd et al., 2006) solver, on Odyssey, the research computing clusters at the Faculty of Arts and Sciences at Harvard University.

 $^{^{20}}$ For example, the Colbert plant in Alabama reports that the sulfur contents as a percentage of coal weight for its five units in April of 1995 are 0.9, 1, 1, 1, and 2. The measurement error in sulfur content does not affect the coefficient estimates in Equation (12) in the first stage, because this measurement error is for the dependent variable. Yet the standard errors of the coefficients are under-estimated.

- 6. compute $W_{i,t+1}^* = W_{i,t}^* + alloc_i + a_{i,t}^* + b_{i,t}^* deduct_{i,t};$
- 7. let t = t + 1, and go back to Step 2 unless t reaches the last year firm i's behavior is observed.
- 8. repeat Steps 1-7 for N_{sim} times.

Paths of $W_i = (W_{i,2000}, W_{i,2001}, ...)$ for each firm *i* that is subject to only Phase II are simulated similarly. I use $N_{sim} = 100$.

The likelihood of the belief parameters and the capital internalization parameter, denoted by θ , at firm *i*'s observed behavior conditional on the states between t_1 and t_2 is:

$$\begin{split} L_{i}(\theta)[(a_{i,t},a_{i,t}^{nv},x_{it})_{t=t_{1}}^{t_{2}},k_{i}|(W_{i,t},P_{t},H_{i,t})_{t=t_{1}}^{t_{2}}] &= \int_{(W_{i,t})_{t=t_{1}}^{t_{2}}} \Pi_{t=t_{1}}^{t_{2}} \{\phi_{a}[A_{i}(W_{i,t}^{*},P_{t},H_{i,t};\theta) - a_{i,t}] \\ &\times \phi_{a^{nv}}[A_{i}^{nv}(W_{i,t}^{*},P_{t},H_{i,t};\theta) - a_{i,t}^{nv}]\phi_{x}[X_{i}(W_{i,t}^{*},P_{t},H_{i,t};\theta) - x_{i,t}]Pr_{i}^{k}(W_{i,t}^{*},P_{t},H_{i,t};\theta) \} \\ &\times dF[(W_{i,t}^{*})_{t=t_{1}}^{t_{2}}|(a_{i,t},a_{i,t}^{nv},W_{i,t})_{t=t_{1}}^{t_{2}}], \end{split}$$

where $\phi_a(\cdot)$ is the probability density function of measurement errors in the current-vintage net allowance purchase with $\eta_{i,t}^a > \frac{1}{alloc_i}(W_{i,t}^* - deduct_{i,t} + alloc_i + a_{i,t})$ truncated, $\phi_{a^{nv}}(\cdot)$ is the probability density function of measurement errors in the next-vintage net allowance purchase with $\eta_{i,t}^{a^{nv}} > \frac{1}{alloc_i}(alloc_i + a_{i,t}^{nv})$ truncated, $\phi_x(\cdot)$ is the probability density function of measurement errors in the sulfur content. The functions $A_i(\cdot), A_i^{nv}(\cdot), X_i(\cdot)$ are firm-specific policy functions of the current-vintage net allowance purchase, the next-vintage net allowance purchase, and the the sulfur content. The function $Pr_i^{k_i}(\cdot)$ is firm *i*'s state-dependent probability of choosing the observed fuel-switching investment k_i .²¹ Those functions depend on θ in a highly nonlinear way via dynamic programming. The integration is over the possible paths of true allowance states. I compute the integration with simulation.

The log likelihood of θ at all firms' behavior D conditional on states S between t_1 and t_2 is:

$$\log L(\theta)(D_{t_1}^{t_2}|S_{t_1}^{t_2}) = \frac{1}{N(t_2 - t_1)} \sum_{i=1}^N L_i(\theta)[(a_{i,t}, a_{i,t}^{nv}, x_{i,t})_{t=t_1}^{t_2}, k_i|(W_{i,t}, P_t, H_{i,t})_{t=t_1}^{t_2}].$$

To estimate θ , I use the BHHH algorithm (Berndt et al., 1974) with numerical gradients. Grid search informs the choice of the initial value for θ . The standard errors of the estimates take into account the errors that come from the first-stage parameter estimates. The standard errors coming from the second-stage structural estimation are calculated using standard formula with numerical gradients.

Results. Table 7 reports the estimation results from the outer loop, by pooling the observations within the same phase of the Acid Rain Program. As a comparison, I present the estimated parameters of the stochastic process of the allowance price along with the estimated belief parameters. Recall that in the allowance price belief specification in Section 4,

 $^{^{21}}$ For behaviors that are irrelevant to some t (for example, the 1999 fuel-switching investment behavior to 1998), their probability densities are excluded.

		Pha	se I	Pha	se II
		Belief	Process	Belief	Process
Market fundamentals	Constant (b_1)	61.13	-15.39	220.19	250.08
		(22.27)	(12.50)	(19.16)	(16.97)
	Year trend (b_2)	14.979	81.76		
		(6.548)	(9.313)		
Dath dopondopoo	$C_{onstant}(h)$	0.478	0.871	0.425	0 791
i atti dependence	Constant (0_3)	(0.478)	(0.071)	-0.433	-0.721
		(0.107)	(0.0190)	(0.088)	(0.113)
	Year trend (b_{4})	-0.101	-0.463		
		(0.0412)	(0.0518)		
Standard devia	tion, (b_5)	28.66	23.04	29.69	24.31
		(12.89)		(15.69)	
Capital internaliz	vation (θ_K)	0.847			
		(0.314)			

Table 7: Phase-specific belief parameter estimates, allowance price process estimates, and the capital internalization estimate.

The stopping criterion for the estimation algorithm is 10^{-3} . Standard errors that also account for the error introduced by the first-stage parameter estimates are in parenthesis.

the expected allowance price next year is the sum of an intercept, interpreted as the contribution from market fundamentals, and a term dependent on the current allowance price, via a slope coefficient. Furthermore, to capture the time-varying market conditions, both the intercept and the slope have linear time trends. The estimation results suggest that in Phase I, firms believe in a "flatter" allowance price process than it actually is; they under-estimate the role of market fundamentals, compared to the past allowance prices, as a driver of the allowance price.

Figure 8 plots the one-year-ahead allowance price predictions in Phase I using the estimates of the belief and the stochastic process from pooling observations in Phase I. While the latter tracks the observed price trajectory reasonably well, the former predicts allowance prices that are too high in early years of Phase I and too low in later years. Figure 9 plots the estimates from pooling observations in Phase II. Unlike the early years of Phase I, the confidence bands in Phase II now overlap more, and the difference in the mean is much smaller.

Aside from the belief estimates, Table 7 also reports a capital internalization estimate of 0.847. Thus, firms consider the fuel-switching investment cheaper than it really is. This is consistent with the Averch-Johnson effect (Averch and Johnson, 1962), where regulated

Figure 8: One-year-ahead predictions of the allowance price using the estimated belief and the estimated stochastic process, based on the Phase I pooled estimates, in 1995 January dollars, 1995-1999. Shaded areas are 95% confidence intervals.

Figure 9: One-year-ahead predictions of the allowance price using the estimated belief and the estimated stochastic process, based on the Phase II pooled estimates, in 1995 January dollars, 2000-2003. Shaded areas are 95% confidence intervals.

utilities tend to invest more than they should, capitalizing on the guaranteed return from the capital under cost-of-service regulation. The Averch-Johnson effect has also been empirically documented in (Fowlie, 2010) and (Cicala, 2015).

I now explore firm heterogeneity in the beliefs. Figure 10 compares the allowance price predictions by firms with above- and below-median coal capacity, based on the phase-specific pooled estimates; and Figure 11 compares the predictions by firms in states that ultimately deregulated the wholesale electricity market, and in states that remain under cost-of-service regulation. They show that larger firms and firms facing competitive pressure from future deregulation hold beliefs that tend to predict allowance prices better, especially in Phase I. Intuitively, larger firms likely had more resources to devote to studying the allowance market, and firms in states preparing for electricity restructuring either were already market-savvy to begin with, or had started to strengthen their competitiveness.

Figure 10: One-year-ahead predictions of the allowance price according to beliefs by firms with above- and below-median coal capacity, based on the phase-specific pooled estimates, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.

To investigate the evolution of beliefs, Figure 12 plots the adaptive, year-by-year estimates; I estimate the stochastic process of the allowance price by adaptively using data up to the point of prediction, and the beliefs by year by using cross-sectional data from each year separately. The overall pattern is similar to that in the pooled estimates in Figures 8 and 9; the belief is flatter than the price process as firms do not appreciate market trends sufficiently. Furthermore, the annually estimated belief starts wildly apart from the adaptively estimated stochastic process, but the gap appears to get smaller towards the end of Phase I and into Phase II.

This pattern of belief evolution is consistent with the intuition that firms adapt to a new environment as they gain experience. The electric utility industry has traditionally focused

Figure 11: One-year-ahead predictions of the allowance price according to the beliefs by firms in states that ultimately deregulated the electricity market and those that did not, based on the Phase-specific pooled estimates, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.

Figure 12: One-year-ahead predictions of the allowance price, from annual belief estimates and the adaptive estimate of the allowance price process, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.

on satisfying rigid demands from regulators. Engineers have been the major decision makers. As the Acid Rain Program progressed, firms may gradually appreciate the philosophy of market-based environmental regulation, that they are free to choose compliance strategies in their best interest, and that a better grasp of the allowance market can create profits. As advocated in Reinhardt (2000), firms should embrace market-based environmental regulation as an opportunity rather than a constraint. Many electric utilities learned to shift decision making away from engineers and towards those with more experience in market and trading.

6 Implications of Biased Beliefs

In this section, I first examine the effect of biased beliefs on the dynamic payoffs of individual firms. In particular, for each firm, I compare the dynamic payoff it actually obtains in Phase I with that under a counterfactual belief that coincides with the stochastic process of the allowance price. This not only quantifies the importance of beliefs to firms, but also provides a lower bound on the dynamic savings to ratepayers under cost-of-service regulation.

Second, I assess the implications of improving the beliefs of "bias-prone" firms, and reducing the allowance price volatility, for aggregate environmental and economic outcomes. The estimation results in Section 5 suggest that smaller firms and firms with less competitive pressure tend to have more biased beliefs. How would the aggregate SO_2 emissions and coal expenditures change if those firms had the same beliefs as bigger firms and firms with more competitive pressure? This quantifies the effects of policies that enable bias-prone firms to have a better understanding of the allowance market; examples include publishing market information relevant to the allowance market in a timely and transparent fashion, holding workshops to facilitate communications among utilities, brokers, and regulators, and introducing competition to the electricity market. Additionally, I simulate the effects of an allowance price collar (*i.e.* a price floor combined with a price ceiling) that constrains the allowance price between \$50 and \$200. Price floors, ceilings, and collars reduce the range of possible allowance prices. They are common policy tools to reduce allowance price volatility; although not used in the Acid Rain Program, they are present in many cap-andtrade programs. What would be the aggregate environmental and economic implications of a price collar in the Acid Rain Program given the biased beliefs?²²

I conclude this section by comparing a cap-and-trade program and an emission tax in a dynamic context. There has so far been little comparison between these two market-based environmental regulations from a dynamic perspective. Cap-and-trade programs have been said to be about price discovery; in the Acid Rain Program, what does price discovery mean quantitatively? That is, what would be the consequence of committing instead to an emission

²²In those counterfactual simulations, I improve the beliefs held by the bias-prone, but not all, firms in my sample, and use a price collar that is non-binding during Phase I. The purpose is to mitigate the feedback effect of simulated behavior on the allowance price. Indeed, the dynamic model in Section 4 is a single-agent model that treats the allowance price as an exogenous state variable; it models how individual firms respond to allowance prices but not the other way round. Counterfactual change in the behavior of bias-prone firms is unlikely to significantly alter the allowance price trajectory; Table 1 shows that even all firms in my sample merely constitute around half of the allowance trade volume from 1993 to 2003.

tax that had been set initially at the level of the biased price projection? Furthermore, does beliefs help distinguish a cap-and-trade program and an emission tax on efficiency grounds?

6.1 Implications for Firm Payoffs

To quantify the effect of biased beliefs on the dynamic payoffs of individual firms, I first simulate each firm's counterfactual choices of allowance trades and sulfur content iteratively from 1995 to 1998, with belief parameters taking values from the "stochastic process" column of Phase I estimation in Table 7. I then calculate the firm's discounted sum of static payoffs from 1995 and 1998 according to Equation (3). To that discounted sum, I add the discounted expected continuation value from 1999. I obtain this discounted expected continuation value by substituting the actual states of electricity demand and allowance price states, and the counterfactual state of allowance stock, in the year-1999 value function of Equation (6)(before the fuel-switching investment cost shock realizes). Having obtained the counterfactual dynamic payoff of each firm, I repeat the steps above with the actual behavior to obtain the actual dynamic payoff of each firm. The difference, normalized by the number of years in Phase I, is the annual loss in the firm's dynamic payoff due to the biased allowance price belief it held during Phase I. I focus on the beliefs in Phase I, because during that period the estimated belief and the stochastic process are substantially different (Figure 8).

Table 8 reports the cost of biased beliefs to firms that have had compliance obligations under the Acid Rain Program since 1995. The annual forgone dynamic payoffs due to biased beliefs in Phase I range from 0.31 to 20.15 million dollars. To put those numbers in context, Table 8 also lists the total revenue of each utility in 1995, and the forgone dynamic payoff as a percentage of profits using a 10% profit margin. A profit margin of 10% is consistent with selected firms' Form 10-K SEC filings in 1995 (when available). Thus, biased beliefs in the first five years of the Acid Rain Program cause firms to forgo an annual dynamic payoff equivalent to 1.6% to 25.7% of their profits, with an average of around 10%.²³ There is a lot of heterogeneity in the forgone dynamic payoffs because of the different sizes of those utilities.

The forgone dynamic payoff to each firm is the lower bound on the forgone dynamic savings to ratepayers. Indeed, the former is the negative of the discounted sum of *internalized* coal and allowance expenditures, while the latter is the discounted sum of coal and allowance expenditures. Since the internalized portion is less than one, the forgone dynamic payoffs are smaller than the forgone dynamic savings. Policies that improve the belief formation process of individual firms are therefore financially beneficial to the ratepayers. The next subsection examines the effects of such policies on aggregate emissions and production costs.

²³This is not saying that firms forgo an average of 10% profits due to biased beliefs; dynamic payoffs are not profits. The dynamic payoff is the infinitely discounted sum of static payoffs, which drive a firm's behavior. Under cost-of-service regulation, the static payoff to a utility includes more than just the financial profits; for example, it also includes some portion of operating costs to show prudence, as discussed in Section 4.

Utility	Forgone dynamic payoffs (annual, million 1995 \$)	1995 revenue (billion 1995 \$)	Forgone dynamic payoffs as share of profits (%)
American Electric Power	20.15	6.06	3.33
Atlantic City Electric Co	6.77	0.95	7.10
Baltimore Gas & Elec Co	5.63	2.23	2.52
Big Rivers Electric Corp	7.74	0.34	22.77
Cincinnati Gas & Electric Co	7.29	1.39	5.23
Cleveland Electric Illum Co	9.03	1.77	5.10
Dairyland Power Coop	0.31	0.01	25.72
Dayton Power & Light Co	11.65	1.03	11.26
Duquesne Light Co	6.50	1.20	5.42
Holyoke Wtr Pwr Co	1.19	0.06	19.59
Illinois Power Co	15.56	1.37	11.36
Indianapolis Power & Light Co	8.78	0.67	13.03
Kentucky Utilities Co	11.60	0.69	16.89
Metropolitan Edison Co	8.64	0.85	10.10
N Y State Elec & Gas Corp	8.75	1.71	5.12
Ohio Edison Co	10.63	2.18	4.88
Ohio Valley Electric Corp	7.71	0.30	25.73
Pennsylvania Elec Co	13.95	0.98	14.22
Pennsylvania Pwr & Lgt Co	10.38	2.75	3.78
Psi Energy Inc	13.91	1.25	11.14
Public Service Co Of NH	7.02	0.98	7.17
Savannah Electric & Power Co	5.24	0.23	22.83
Southern Company	13.49	8.56	1.58
Southern Indiana G & E Co	5.67	0.28	20.60
Virginia Electric & Power Co	10.21	4.35	2.35

Table 8: Cost of biased beliefs in Phase I to firms that have had compliance obligations under theAcid Rain Program since 1995.

Profits are approximated using a profit margin of 10%, consistent with selected firms' Form 10-K SEC filings in 1995 (when available).

Phase I (1995-1999)	Improve beliefs of		Reduce volatility by		
	Small firms	Firms always regulated	\$50 floor \$200 ceiling	Pre-1995-projected \$200 tax	
Δ Aggregate emissions (thousand tons)	115.86	12609	2099	-9092	
$\%~\Delta$ Aggregate emissions	0.44	47.76	7.95	-34.44	
Δ Aggregate coal cost (million 1995 \$)	-18.53	-411.99	-226.72	905.17	
Δ Per-firm coal cost	-1.68	-29.42	-9.06	36.20	

Table 9: Aggregate effects of improving beliefs and reducing volatility.

6.2 Implications for Aggregate Environmental and Economic Outcomes

Columns 2 and 3 in Table 9 report the changes in the aggregate SO₂ emissions and the aggregate production cost, *i.e.*, coal expenditure, if firms with more biased beliefs shared beliefs with the those with less biased beliefs. Column 2 shows that if smaller firms, or firms with below-median generation capacity, had the same belief as did the bigger firms, aggregate emissions would increase by 0.44% and aggregate production costs decrease by 18.53 million 1995 dollars during Phase I, equivalent to an average saving of 1.68 million dollars per small firm. Column 3 shows that if firms in states that still regulate the wholesale electricity market had the same belief as those in states with more competitive pressure because of pending restructuring, aggregate emissions would increase by 47.76% and aggregate production costs decrease by 412 million dollars, averaging 29.43 million dollars per always-regulated firm. The much larger magnitudes of those changes are due to a few big firms (*e.g.*, the Southern Company) that remain under cost-of-service regulation.

Thus, in Phase I of the Acid Rain Program, as firms improve beliefs, they reduce abatement efforts. Indeed, when firms recognized the time trends in the allowance price process better, they would predict the rises and declines in the allowance price better. Since declines are more dramatic than rises for the first five years, the improvement in the prediction of declines would be more pronounced than that of rises. Now that firms' beliefs would adjust more downwards about declines than upwards about rises, overall the abatement effort would reduce, leading to lower production cost and higher emissions.

Reducing allowance price volatility by price floors and ceilings would change the aggregate emissions and production costs in the same direction. Column 4 in Table 9 shows that an allowance price floor of \$50 combined with a ceiling of \$200 would increase aggregate emissions by 7.95% and reduce the aggregate production costs by 83.97 million dollars, averaging 9.07 million dollars per firm.

6.3 Implications for the Choice between Cap-and-Trade and Tax

An emission tax is an alternative to a cap-and-trade program and another prominent example of market-based environmental regulation. The last column of Table 9 quantifies the price discovery feature of the cap-and-trade program. Under a tax, the biased initial expectation about the marginal abatement cost lingers on. With a cap-and-trade program, however, the market would gradually sort out the truth, as firms start out with biased beliefs but improve those beliefs over time. Indeed, as shown in the last column of Table 9, price discovery in the first five years of the Acid Rain Program avoids an over-abatement of 34.44% aggregate emissions and saves 905 million dollars in aggregate production costs.

Beliefs change the efficiency of cap-and-trade programs relative to that of emission taxes. Efficiency is the benefit minus the cost of emission reduction. Full efficiency requires that each firm reduce their emissions up to the point where its marginal cost of emission reduction equals the marginal benefit it creates. Beliefs create a connection between marginal costs and the marginal benefits under a cap-and-trade program, but are almost irrelevant under an emission tax, which is politically difficult to change dynamically. Thus, if beliefs under a cap-and-trade program align marginal costs with marginal benefits, then cap-and-trade will be more efficient than tax.

To illustrate, Consider a big, rural electric utility such as Southern Company, and a small, urban electric utility such as Atlantic City Electric Company. Suppose that the clean coal price has been declining, so the allowance price tends to go down. One of my empirical findings is that bigger companies have less biased beliefs. Thus, Southern Company, being bigger and having more resources, appreciates that decline in allowance price better, while Atlantic City Electric Company tends to over-predict tomorrow's allowance price. As a result, Atlantic City Electric Company would be more aggressive in reducing its own emissions, leading to a higher marginal cost than the southern company. The higher marginal cost happens to align with its higher marginal benefit, because emission reductions in urban areas are more beneficial than those in rural areas (Muller and Mendelsohn, 2009). In this case, cap-and-trade is more efficient than, as beliefs about the future allowance price align the marginal cost and the marginal benefit of emission reduction.

For example, index Southern Company by 1 and Atlantic City Electric Company by 2. Assume that $MC_1 = 50 + 2r_1$, $MB_1 = 60$, $MC_2 = 20 + 10r_2$, $MB_2 = 140$, where MC denotes marginal cost, MB marginal benefit, and r emission reduction. Table 10 shows in the first row that the net benefit from a tax of 100 is 265. In the immediately following rows, if beliefs align marginal costs with marginal benefits (that is, $MC_1 < MC_2$), the net benefit under a cap-and-trade program will significantly exceed that under a tax under a range of scenarios. Of course, as the remaining rows show, when beliefs misalign the marginal costs and benefits, the net benefit can fall significantly below that under a tax.

Previous literature misses the beliefs as a potential determinant of the relative efficiency of cap-and-trade programs and emission taxes, because of its focus on static models, where firms care about current profits only and beliefs are therefore irrelevant. The literature has proposed that to achieve efficiency, government intervention is necessary, in the form of

Scenario	Net Benefit
$\tan = 100$	265
(MC_1, MC_2) under cap-and-trade	
when marginal cost and marginal benefit align:	
(90,110)	475
(80,120)	625
(70,130)	715
(60,140)	745
(80,110)	600
(80,100)	565
when marginal cost and marginal benefit misalign:	
(110, 90)	-5
(120, 80)	-335
(130,70)	-725
(140,60)	-1175

Table 10: Net benefits from an emission tax and a cap-and-trade program in the SouthernCompany - Atlantic City Electric Company example.

setting up trading ratios. I show that once we move beyond static models and starting using a dynamic model, where beliefs are the key driver of behavior, beliefs are a decentralized channel that affects the efficiency of cap-and-trade relative to tax.

7 Conclusion

This paper studies the evolution of firms' beliefs about future market conditions in a new environment. The empirical context is the U.S. Acid Rain Program, the first cap-and-trade program and a landmark experiment in market-based environmental policy. I use structural dynamic estimation to back out coal-dependent private electric utilities' beliefs about the future SO_2 allowance price from their compliance behavior. The three-dimensional continuous-state mixed-continuous-action dynamic model of allowance trades, coal quality choices, and fuel-switching investment offers the first empirical model of firm behavior in cap-and-trade programs.

I find that firms have biased beliefs about the future allowance price. The beliefs are not consistent with the stochastic process of the allowance price. In particular, firms underestimate the role of market fundamentals as a driver of the allowance price. This cost the firms dynamic payoffs equivalent to an average of about 10% of their profits in the first five years of the program. Under cost-of-service regulation, this is the lower bound on the dynamic savings to ratepayers. Smaller firms and firms facing less competitive pressure have more biased beliefs. Over time, firms' beliefs appear to converge towards the stochastic process. This is consistent with the shift in the management practice in electric utilities during that period; compliance decisions were made less by engineers but more by people experienced in markets and trading.

Beliefs and dynamics add new insights to the debate on the choice between cap-and-trade programs and taxes to regulate pollution. A tax equates marginal abatement costs across firms, while dynamic considerations under a cap-and-trade programs cause firms to equate their marginal abatement costs with their own marginal dynamic value of allowances. As a result, when beliefs about future market conditions align the dynamic value of allowances, and therefore the marginal abatement costs, with the marginal benefit of emissions reduction, cap-and-trade programs are more efficient than taxes. In other words, beliefs are a decentralized channel that can potentially align private interests with public benefits under a cap-and-trade program.

Following the success of the Acid Rain Program, many countries and regions have implemented, or designed for implementation, the cap-and-trade approach to regulating pollution. Those cap-and-trade programs that regulate carbon dioxide emissions, most notably the ongoing European Trading Scheme and the forthcoming cap-and-trade program in China, have huge economic values. Hundreds of billion dollars are at stake, much larger than the annual value of the SO₂ allowance market (which is around one billion). A careful consideration of dynamics and beliefs will be particularly helpful for choosing, designing, and evaluating cap-and-trade programs.

This research points to several future research efforts. First, how do firms' beliefs evolve in those states after they underwent electricity restructuring? Intuitively, competition should improve belief formation. Answering that question requires a competitive model of firm behavior in place of a single-agent model as constructed in this paper.²⁴ Another complication is the potential need to model the dispatch decision between coal and gas, and to model the capital investment to switch to or co-fire with gas, as gas became competitive with coal. Second, does the experience that firms have gained in the Acid Rain Program spill over to later cap-and-trade programs, such as the NO_x Budget Program?²⁵ Third, it would be interesting to collect organizational data to formally investigate the exact forms of shifts in the management practice in electric utilities, and study their impact on belief formation. Fourth, while this paper only estimates the beliefs, it would be useful to investigate what learning rules have led to those beliefs.

 $^{^{24}}$ One of the counterfactual analysis described in the previous section compares the beliefs between utilities when they were still subject to cost-of-service regulation, in states that are always regulated and those that were ultimately deregulated.

²⁵I thank Kenon Smith, who was formerly at EPA, for suggesting this.

References

- Averch, H. and Johnson, L. L. (1962). Behavior of the Firm under Regulatory Constraint. American Economic Review, 52(5):1052.
- Bailey, E. M. (1998). Allowance Trading Activity and State Regulatory Rulings: Evidence from the U.S. Acid Rain Program.
- Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating Dynamic Models of Imperfect Competition. *Econometrica*, 75(5):1331–1370.
- Berndt, E. K., Hall, B. H., Hall, R. E., and Hausman, J. A. (1974). Estimation and Inference in Nonlinear Structural Models. Annals of Economic and Social Measurement, 3(4):653– 665.
- Bray, R. L. (2017a). Markov Decision Processes with Exogenous Variables.
- Bray, R. L. (2017b). Strong Convergence and Dynamic Economic Models.
- Byrd, R. H., Nocedal, J., and Waltz, R. A. (2006). Knitro: An Integrated Package for Nonlinear Optimization.
- Carlson, C., Burtraw, D., Cropper, M., and Palmer, K. (2000). Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade? *Journal of Political Economy*, 108(6):1292–1326.
- Cason, T. N. (1993). Seller Incentive Properties of EPA's Emission Trading Auction. Journal of Environmental Economics and Management, 25(2):177–195.
- Cason, T. N. (1995). An Experimental Investigation of the Seller Incentives in the EPA' s Emission Trading Auction. *American Economics Review*, 85(4):905–922.
- Cason, T. N. and Plott, C. R. (1996). EPA's New Emissions Trading Mechanism: A Laboratory Evaluation. *Journal of Environmental Economics and Management*, 30:133–160.
- Chan, R., Chupp, B. A., Cropper, M., and Muller, N. (2017). The Impact of Trading on the Costs and Benefits of the Acid Rain Program,.
- Cicala, S. (2015). When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation. *American Economic Review*, 105(1):411–444.
- Covert, T. R. (2014). Experiential and Social Learning in Firms : The Case of Hydraulic Fracturing in the Bakken Shale.
- Davis, L. W. (2015). Bonding Requirements for U.S. Natural Gas Producers. Review of Environmental Economics and Policy, 9(1):128–144.
- Doraszelski, U., Lewis, G., and Pakes, A. (2018). Just Starting Out: Learning and Equilibrium in a New Market. *American Economic Review*, 108(3):565–615.

- Ellerman, A. D., Joskow, P. L., Schmalensee, R., Montero, J.-P., and Bailey, E. M. (2000). Markets for Clean Air: The U.S. Acid Rain Program. Cambridge University Press.
- Fieldston Company (1993). Compliance Strategies Review. Technical report, Washington DC.
- Fourer, R., Gay, D. M., and Kernighan, B. W. (2003). A Modeling Language for Mathematical Programming. Thomson Brooks/Cole.
- Fowlie, M. (2010). Emissions Trading, Electricity Restructing, and Investment in Pollution Abatement. *The American Economic Review*, 100(3):837–869.
- Fowlie, M., Reguant, M., and Ryan, S. P. (2016). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy*, 124(1):249–302.
- Goldfarb, A. and Xiao, M. (2011). Who Thinks about the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets. *The American Economic Review*, 101(7):3130–3161.
- Hahn, R. W. and May, C. A. (1994). The Behavior of the Allowance Market: Theory and Evidence.
- Hansen, L. P. (2007). Beliefs, Doubts and Learning : Valuing Macroeconomic Risk. American Economic Review, 97(2):1–30.
- Hortaçsu, A., Puller, S. L., Luco, F., and Zhu, D. (2016). Strategic Ability and Productive Efficiency in Electricity Markets.
- Joskow, P. L. (1972). A Behavioral Theory of Public Utility Regulation. PhD thesis.
- Kolstad, C. and Turnovsky, M. (1998). Cost Functions and Nonlinear Prices: Estimating a Technology with Quality-Differentiated Inputs. *Review of Economics and Statistics*, 80(3):444–453.
- Lile, R. and Burtraw, D. (1998). State-Level Policies and Regulatory Guidance for Compliance in the Early Years of the SO2 Emission Allowance Trading Program. *RFF Discussion Paper*.
- Montero, J.-P. and Ellerman, A. (1998). Explaining Low Sulfur Dioxide Allowance Prices: the Effect of Expectation Errors and Irreversibility. *MIT Center for Energy and Environmental Policy Research Working Paper 98-011.*
- Muller, N. Z. and Mendelsohn, R. (2009). Efficient Pollution Regulation: Getting the Prices Right. American Economic Review, 99(5):1714–1739.
- Reinhardt, F. (2000). Down to Earth : Applying Business Principles to Environmental Managemen. Harvard Business School Press, Boston, MA.
- Rose, K., Burns, R. E., Coggins, J. S., Harunuzzaman, M., and Viezer, T. W. (1992). Public Utility Commission Implementation of the Clean Air Act's Allowance Trading Program.

- Rose, K., Taylor, A. S., and Harunuzzaman, M. (1993). Regulatory Treatment of Electric Utility Clean Air Act Compliance Strateiges, Costs, and Emission Allowances.
- Schmalensee, R. and Stavins, R. N. (2013). The SO2 Allowance Trading System: The Ironic History of a Grand Policy Experiment. *Journal of Economic Perspectives*, 27(1):103–122.
- U.S. Environmental Protection Agency (1994). National Air Pollutant Emission Trends.

Appendix A Data Compilation

A.1 Operations Data

There are three levels of operations data: unit (or boiler), plant, and firm (or operating utility). Units are the regulatory targets in the Acid Rain Program; allowances are allocated to, and compliance defined by, units. Units are identified with string IDs reported by utilities. Plants and utilities are identified with numerical IDs assigned by the Energy Information Administration (EIA).

Firm-plant relationships describe which firm operates which plant in which year. They are obtained from the "plant" file in Form EIA-767, "Annual Steam-Electric Plant Operation and Design Data". This form covers all U.S. plants with a total existing or planned steam-electric unit, with a generator nameplate rating of 10 megawatts or larger, that is fueled by organics, nuclear, and combustible renewables. Plant-unit relationships describe which plant houses which unit in which year. They are obtained from the "boiler" file in Form EIA-767.

Unit-level data. The boiler design, compliance, and emissions data are from the Air Market Programs Data (AMPD) at the Environmental Protection Agency (EPA). The fuel input and data are from the "boiler-fuel" file in Form EIA-767. The generation data are from the "boiler-generator" and "generator" files in Form EIA-767. The scrubber data are from the "boiler-FGD" and "FGD" files in Form EIA-767.

The design data, available in 1990 and then every year since 1995, include the ARP phase designation, primary and secondary fuel types, operating status, commercial operation date, latitude and longitude, and primary and secondary representatives.

The compliance data, available every year since 1995, include the allowance allocation, the allowance holding by the deduction deadline (generally January 30 of the following year), and other deductions.

The emissions data include the operating time, gross load, steam load, heat input, and $SO_2/NO_x/CO_2$ emissions. I use monthly data, which start in 1995 for Phase I units and 1997 for the remaining units, measured at the flue gas outlets using the continuous emission monitoring devices.

The fuel input data, available on a monthly basis, include the heat input, heat content, sulfur content, and ash content of coal used, the heat input of gas used, and the heat input of other fuels (e.g., fuel oil, petroleum coke, biomass, etc.) used. Starting in 2001, coal is further divided into bituminous, sub-bituminous, and lignite coals.

The generation data, available on a monthly basis, include the net generation and the nameplate capacity of generators. The mapping from generators to units is not one-to-one. In order to obtain the amount of generation and capacity that each unit is responsible for, I adopt the following procedure: 1) if multiple generators share only one unit, I assign to that unit the sum of generation or capacity by all those generators; 2) if multiple boilers share only one boiler, I allocate the generation or capacity among units proportionately to each unit's heat input; and 3) if m > 1 generators are associated with n > 1 boilers, I first sum up across m generators and then allocate the sum by each unit's heat input.

The scrubber data, available on an annual basis, include scrubber operation variables and

the design parameters of scrubbers associated with scrubbed units. The operation variables include hours in service, sorbent quantity, energy consumption, and non-energy operating cost. The design parameters include the in-service date, scrubber type, manufacturer, sulfur removal rate, electric power requirement, and non-land nominal installed cost. I multiple units share a scrubber, I allocate the operating cost based on each unit's sulfur input (that is, heat input multiplied by sulfur content), and the installed cost based on the capacity of the generator associated with each unit.

Plant-level data. The plant divestiture data are from Cicala (2015). The fuel shipment data are from Form FERC-423, "Monthly Report of Cost and Quality of Fuels for Electric Plants".

The divestiture data include, for each plant, whether it is divested and when. The fuel shipment data, available on a monthly basis, include fuel type, quantity, quality (sulfur, ash, heat contents), contract type (spot, contract, new contract), contract status (whether expiring in 2 years), and source county of each shipment a plant receives.

Firm-level data. The utility accounting data are from the "TYP1" and "File 1" files in Form EIA-861, "Annual Electric Utility Report". The accounting data, available on an annual basis, include the ownership type (federal, state, municipal, private, co-op, power marketers, and municipal power marketers), net generation, electricity sales to different classes (residential, commercial, industrial, public lighting, wholesale) and the associated revenues.

A.2 Allowance Transfer Data

The allowance transfer data are from EPA AMPD. The data include, for all transfers of allowances, the account numbers and names of the transferor and the transferee and their representatives, date of transfer, number and vintage year of allowances transferred and type of transfer. I obtain from this data the net non-trading transfers and the net trading transfers of each utility in each year. To achieve this goal, I first infer the owning utility of each allowance account, and then identify the non-trading and trading transfers.

Inferring ownership of allowance accounts. There are two types of allowance accounts: unit accounts, and general accounts (in addition to the EPA administrative accounts). Each unit subject to the Acid Rain Program has one and only one unit account. The account number of a unit account contains the plant ID and the unit ID, so that the utility that owns this unit account can be easily identified.

General accounts can be set up by any company and person. The account number of a general account is not informative. Thus, to map general accounts to the owning utilities (if they are owned by utilities at all) effective at the time of the transaction, I rely on the account name and the representative name, supplemented by information on utility name/ownership changes from Form EIA-767 and online searches.

Identifying non-trading allowance activities. Those include allowance allocation, bonus allowances, and other administrative transfers. For reasons discussed below, I take

them as given in utilities' allowance stock paths.

Allowance allocation is identified by the "initial allocation" transfer type. Allowances of vintage years 1995 - 2024 were distributed to all unit accounts on March 23, 1993. Allowances of vintage year 2025 were distributed in 1994. Allowances of vintage years 20 years ahead were distributed in each year starting in 1996. Utilities had perfect foresight over the allowance allocations by the time the program started in 1995, because they had been determined in the Acid Rain Program legislation several years before 1995.

Phase 1 extension bonus allowances are identified by the "phase 1 extension issuance" transfer type. These allowances have vintage years 1995-1999 and were distributed in September 1994 to most Phase I units that install scrubbers to comply with the Acid Rain Program. Utilities had perfect foresight over the bonus allowances by the time the program started in 1995, because they had been determined when utilities began installing scrubbers several years before 1995.

Other administrative transfers include: other bonus allowances (conservation, early reduction, energy biomass, energy geothermal, energy solar, energy wind, small diesel), other deductions (penalty, voluntary, etc.), error correction, state cap related, auction transfers, etc. Those administrative transfers are much smaller than allowance allocations and Phase 1 extension bonus allowances for the utilities in my sample; I take them as given for simplicity.

Identifying trading activities. The remaining allowance transfers are allowance trades. The main challenge is to differentiate between intra- and inter-utility transfers. The intrautility transfers are reallocation among the accounts that a single utility owns, while the inter-utility transfers are the trades of interest for this paper; I use the latter to calculate the net allowance purchase by each utility in each year.

The first issue is to decide whether utilities under a parent company (such as the Southern Company) should be treated as separate decision makers. Some parent holding companies have their own allowance accounts and appear to have centrally managed allowance trading of some of their subsidiary companies. They can be identified by consistent, extremely large volumes of transactions out of a subsidiary company to another, sometimes with a long sequence of vintages. In those cases, I treat the parent company, rather than each of the subsidiary companies, as the decision maker on allowance transactions as well as operations.

The second issue is to infer the start and end dates relevant to the allowance trades to be count towards compliance in each year. Although allowance deduction is based on the emissions incurred in a calendar year, the deduction itself does not occur at the turn of calendar years. Rather, it is aimed that utilities have until January 31 of each year to make sure they have enough allowances of the eligible (prior-year and current-year) vintages in each of the unit accounts to cover that unit's emissions incurred in the previous calendar year. However, both communications with EPA and the data show that the deadline was almost always extended. To infer the effective deadline for each year, I take the last date on which I observe apparently non-outlier private transfer (that is, after which I only observe one or two transactions months ahead) of eligible allowances to the allowance accounts of complying units.

A.3 Other Data

The monthly market price for an SO₂ allowance of the current or earlier vintages is obtained from Denny Ellerman, who collected the data from trade journals and brokerage firms over the years. Three price indices are reported: Cantor Fitzgerald, Emissions Exchange Corporation, and Fieldston. In months when multiple indices are available, they differ very little. I use the Cantor Fitzgerald price index, available starting in August 1994. The monthly market price for an allowance with future vintages, when available, is from the online archive of Cantor Fitzgerald / BGC Environmental Brokerage Services, the biggest allowance broker, available at http://www.bgcebs.com/registered/aphistory.htm.

I use the monthly Producer Price Index by Commodity for Crude Energy Materials, available at https://fred.stlouisfed.org/series/PPICEM, to deflate the fuel cost. I use the monthly Urban Consumer Price Index to deflate the allowance price.

Appendix B Additional Data Patterns

This section presents two more data patterns that inform the structural model, in addition to those already provided in Section 3.

Volumes and vintages of allowance transactions. While utilities can trade allowances of any vintage, almost all transactions during the period of analysis concern the current-year, the prior-year, and the next-year vintages. Figure 13 plots the distribution of volume share of current-, prior-, and next-vintage allowance transactions in any-vintage transactions for utility-years during the compliance periods up to 2003. The utility-year transactions are almost always in current-, prior-, and next-vintage allowances. The model in Section 4 thus focuses on allowance trading in those vintages.

Figure 14 plots the distribution of current- and prior-vintage net allowance purchases for utility-years during the compliance periods up to 2003. Figure 15 plots the counterpart for next-vintage net allowance purchases since one year before the compliance periods up to 2002. The current- and prior-vintage allowance trades are of larger volumes than the next-vintage trades.

Coal and gas prices and the dispatch decision. During the period of analysis, utilities with both coal and gas capacities would typically dispatch coal first. This justifies including those utilities in my sample without modeling their dispatch decisions. Indeed, the natural gas price was much higher than the coal price. Figure 16 plots the distributions of the delivered prices of coal and natural gas reported to Form FERC-423 during the period of analysis. The lower quartile gas price always exceeded the upper quartile coal price, in most years by a lot. The added cost of dispatching coal from its sulfur emissions, given the relatively low allowance price, is unlikely to switch the dispatch order of coal and gas. The boom of shale gas production did not start until around 2007, which is four years after the end of the period of analysis; Figure 17 shows that the massive drop in the natural gas price did not start until 2008.

Figure 13: Distribution of volume share of current-, prior-, and next-vintage allowance transactions in any-vintage transactions during the compliance periods, utility-year, 1995-2003.

Figure 14: Distribution of current- and prior-vintage net allowance purchases during the compliance periods, utility-year, 1995-2003.

Figure 15: Distribution of next-vintage net allowance purchases since one year before the compliance periods, utility-year, 1994-2002.

Figure 16: Delivered price of coal and natural gas, 1995-2003, in current U.S. cents per mmBTU. Calculated from the FERC-423 data.

Figure 17: Natural gas price, 1990-2013, in 2010 U.S. dollars per thousand cubic feet. From Davis (2015).

Appendix C Shape of the Phase II Value Function

This section shows that the Phase II value function $V(W_i, P, H_i)$ is increasing and concave in the allowance stock under mild regularity conditions. Then, the first-order conditions in Section 4 are sufficient for optimality.

The Phase II value function is:

$$V_{i}(W_{i}, P, H_{i}) = \max_{\substack{x_{i} \in X\\a_{i} \ge x_{i}H_{i} - W_{i} - alloc_{i}}} \{\phi_{A}[A(a_{i}; P)] + \phi_{C}[C_{i}(x_{i}; H_{i})] + \beta \int V_{i}(W_{i} + alloc_{i} + a_{i} - x_{i}H_{i}, P', H'_{i}) dF_{P}(P'|P) dF_{H_{i}}(H'_{i}|H_{i})\}$$

,

Letting $W'_i = W_i + a_i + alloc_i - x_iH_i$ be the choice variable in place of a_i , we rewrite the value function as follows:

$$V_{i}(W_{i}, P, H_{i}) = \max_{\substack{x_{i} \in X \\ W_{i}' \geq 0}} \{ \phi_{A}[A(W_{i}' - W_{i} - alloc_{i} + x_{i}H_{i}; P)] + \phi_{C}[C_{i}(x_{i}; H_{i})]$$

+ $\beta \int V_{i}(W_{i}', P', H_{i}') dF_{P}(P'|P) dF_{H_{i}}(H_{i}'|H_{i}) \},$

By the Envelope Theorem, we have:

$$\frac{\partial V_i(W_i, P, H_i)}{\partial W_i} = -\frac{\partial \phi_A(A^*)}{\partial A} \frac{\partial A((W'_i)^* - W_i - alloc_i + x_i^*H_i; P)}{\partial a_i} + \mu^*,$$

where $\mu^* \geq 0$ is the Lagrange multiplier associated with the compliance constraint $W'_i \geq 0$. Since the internalization function, ϕ_A , is decreasing in the allowance expenditure (the larger the magnitude of the allowance expenditure, the more negatively the internalized allowance expenditure enters as a cost in the utility's payoff function), and the allowance expenditure function, A, is increasing in the net allowance purchase, $\frac{\partial V_i(W_i, P, H_i)}{\partial W}$ is positive. Intuitively, more allowances cannot hurt a utility; it can always hold on to the additional allowances it has and replicate the payoff it had received before with fewer allowances.

So far we have shown that the Phase II value function is increasing in the allowance stock. To show concavity, we have:

$$\frac{\partial V_i^2(W_i, P, H_i)}{\partial^2 W_i} = \frac{\partial \phi_A(A^*)}{\partial A} \frac{\partial^2 A((W_i')^* - W_i - alloc_i + x_i^* H_i; P)}{\partial a_i^2}.$$

Since the internalization function, ϕ_A , is decreasing in the allowance expenditure, and the allowance expenditure function is convex in the net allowance purchase (the more net allowance purchase, the higher the marginal cost because of the quadratic allowance cost), $\frac{\partial V_i^2(W_i, P, H_i)}{\partial^2 W_i} \leq 0$, thus concavity of the Phase II value function with respect to the allowance stock.

Appendix D Inner Loop Details

This section describes in detail how the inner loop works, or how I solve for the value function at a given parameter vector. In particular, I explain the application of the Relative and the Endogenous Value Function Iterations methods (Bray, 2017b,a).

Approximating the value function using Chebyshev polynomials. I use Chebyshev polynomials to approximate the value function (Judd, 1998). The alternative method, state discretization, is impractical in this three-dimensional continuous-state problem. Thus:

$$V(W, P, H) \approx \sum_{d_1=0}^{N_d} \sum_{d_2=0}^{N_d} \sum_{d_3=0}^{N_d} coef(d_1, d_2, d_3) T_{d_1}(W) T_{d_2}(P) T_{d_3}(H)$$

where N_d is the degree of Chebyshev polynomials, $coef(d_1, d_2, d_3)$ is the Chebyshev coefficient with degree (d_1, d_2, d_3) , and $T_d(s)$ is the Chebyshev polynomial with degree d at state $s \in [\underline{s}, \overline{s}]$:

$$T_d(s) = \cos(d\cos^{-1}(2\frac{s-\underline{s}}{\overline{s}-\underline{s}}-1)).$$

Computation of $T_d(s)$ uses the recursive formula:

$$\begin{split} T_0(s) &= 1, \\ T_1(s) &= s, \\ T_d(s) &= 2sT_{d-1}(s) - T_{d-2}(s), \ d \geq 2. \end{split}$$

To obtain $coef(\cdot, \cdot, \cdot)$ at a particular iteration, I solve the maximization problem in the Bellman's equations at each approximation node in the state space, and use the optimized value to update the coefficients. To facilitate coefficient updating, I use Chebyshev nodes as approximation nodes; for state variable s, the approximation nodes are $(s_1, s_2, \ldots, s_{N_j})$ such that:

$$s_j = (-\cos(\frac{2j-1}{2N_j}\pi) + 1)(\frac{\bar{s}-\underline{s}}{2}) + \underline{s}, \qquad j = 1, 2, \dots, N_j$$

where N_j is the number of approximation nodes for that state variable. Then, the Chebyshev coefficients are:

$$coef(d_1, d_2, d_3) = \frac{\sum_{j_1=1}^{N_j} \sum_{j_2=1}^{N_j} \sum_{j_3=1}^{N_j} V^*(W_{j_1}, P_{j_2}, H_{j_3}) T_{d_1}(W_{j_1}) T_{d_2}(P_{j_2}) T_{d_3}(H_{j_3})}{\sum_{j_1=1}^{N_j} T_{d_1}(W_{j_1})^2 \sum_{j_2=1}^{N_j} T_{d_2}(P_{j_2})^2 \sum_{j_3=1}^{N_j} T_{d_3}(H_{j_3})^2}, \quad (13)$$

where $V^*(W_{j_1}, P_{j_2}, H_{j_3})$ is the maximized value at the approximation node $(W_{j_1}, P_{j_2}, H_{j_3})$. To keep the number of maximization problems manageable, I use complete polynomials instead of tensor products of polynomials; thus, $coef(d_1, d_2, d_3)$ is updated according to Equation (13) if $d_1 + d_2 + d_3 \leq N_d$, and $coef(d_1, d_2, d_3) = 0$ otherwise. I use $N_d = 3$ and $N_j = 4$.

Each maximization problem involves integrating the current value function iterate over the exogenous state transitions in the allowance price P and the heat input H. I compute the integral using the Gauss-Hermite quadrature:

$$\mathbb{E}[V(W', P', H')|P, H] \approx \sum_{w_2=1}^{N_w} \sum_{w_3=1}^{N_w} \frac{1}{\pi} \omega(w_2) \omega(w_3) \\ \times V(W', \sqrt{2\sigma^P} n(w_2) + \gamma_0^P + \gamma_1^P P, \sqrt{2\sigma^H} n(w_3) + \gamma_0^H + \gamma_1^H H)$$

where N_w is the degree of Gauss-Hermite quadrature, $\omega(w)$ is the the Gauss-Hermite weight at degree w, and n(w) is the the Gauss-Hermite node at degree w. The parameters $(\sigma^P, \gamma_0^P, \gamma_1^P)$ are the standard deviation (b_5) , intercept $(b_1 + b_2(s - 1))$, or $b_1 + b_2 \times 2000$, or b'_1 , and slope of the allowance price transition $(b_3 + b_4(s - 1))$, or $b_3 + b_4 \times 2000$, or b'_3) from Equation (1) or (2). The parameters $(\sigma^H, \gamma_0^H, \gamma_1^H)$ are the standard deviation, intercept, and slope of the heat input transition from Equation (8).

Accelerating dynamic programming using Relative and Endogenous Value Function Iterations. Both Relative and Endogenous Value Function Iterations methods leverage the fact that what matters for behavior is the shape rather than the level of the value function. Thus, we only need to check the iterative difference in the shape, not the level, for convergence. The shape converges at least as fast as does the level, and in many cases much faster. See Bray (2017b,a) for formal results.

The two methods differ in how they define the shape. The Relative Value Function Iterations method uses the shape covering all states. Then, the relative value function is the full value function shifted by a constant. The Exogenous Value Function Iterations method looks at the shape specific to each exogenous state; each shape covers all endogenous states at each exogenous state. To see why the collection of exogenous-state-specific shapes is sufficient for behavior, suppose that at a particular exogenous state, the payoffs at all endogenous states are shifted by the same constant. Since the firm controls the endogenous state but not the exogenous, the relative attractiveness of choices does not change. I implement the Relative Value Function Iterations method by normalizing the value function iterate by the value at the first approximation node, and the Exogenous Value Function Iterations method by normalizing the value function Iterations method by normalizing the value function iterate at each exogenous state by the value at that exogenous state and the first approximation node of the endogenous state.²⁶

Table 11 report the performances of the Endogenous Value Function Iterations, the Relative Value Function Iterations, and the full value function iterations methods for the Phase II problem conditional on the chosen fuel-switching investment at the estimated parameters. The Endogenous Value Function Iterations method takes fewer iterations than does the Relative Value Function Iterations method, which in turn takes many fewer iterations than does the full value function iterations method. The saving in the computing time is substantial: the Exogenous and Relative Value Function Iterations method.

Recovering the Phase II value function from the relative value function for Phase I parameter estimation purposes. As discussed earlier, I solve the Phase II dynamic problem by the Relative, rather than Endogenous, Value Function Iterations method for Phase I parameter estimation purposes. This is because the 1999 investment problem requires the level of value associated with each fuel-switching investment option, which can be backed out from the Relative Value Function Iterations method.

I use generic notations below. Denote the value function by $V^{full}(s)$, which satisfies:

$$V^{full}(s) = \pi(x^*(s), s) + \beta \mathbb{E}[V^{full}(s') | x^*(s), s],$$
(14)

where $\pi(\cdot, \cdot)$ is the static payoff function and $x^*(\cdot)$ is the (optimal) policy function. Let the relative value function at the convergent round, k, be $V^{(k)}(s)$. By the property of the relative value function, $V^{full}(s) = V^{(k)}(s) + L^{(k)}$. Equation (14) now becomes:

$$V^{(k)}(s) + L^{(k)} = \pi(x^*(s), s) + \beta \mathbb{E}[V^{(k)}(s')|x^*(s), s] + \beta L^{(k)}.$$
(15)

By definition of convergence, $V^{(k-1)}$ and $V^{(k)}$ as given by:

$$V^{(k)}(s) = \pi(x^*(s), s) + \beta \mathbb{E}[V^{(k-1)}(s') | x^*(s), s]$$
(16)

have (almost) the same shape. Hence:

$$V^{(k-1)}(s) - V^{(k-1)}(s_0) = V^{(k)}(s) - V^{(k)}(s_0)$$

where s_0 is the state used for normalization. Substituting $V^{(k-1)}(s) = V^{(k)}(s) - V^{(k)}(s_0) + V^{($

²⁶Bray (2017b,a) formulate the Relative and Endogenous Value Function Iterations methods in dynamic models with states that are naturally discrete or easily discretizable. I thank Robert Bray for discussing ways to apply those methods to continuous-state dynamic programming in my context.

Table 11: Performance comparison of the Endogenous Value Function Itera-
tions (Bray, 2017b), the Relative Value Function Iterations (Bray, 2017a), and
the full value function iterations methods, for the Phase II problem conditional
on the chosen fuel-switching investment at the estimated parameters.

Utility	Numbe	Number of iterations		Compu	Computing time (min)	
	Endo.	Rel.	Full	Endo.	Rel.	Full
Carolina Power & Light Co	7	47	980	3.2	24	410
Detroit Edison Co	7	36	977	3.2	19	407
Duke Energy Corp	8	85	989	3.7	41	414
South Carolina Electric&Gas Co	8	32	918	3.7	16	372
Kentucky Utilities Co	8	50	925	3.8	25	389
Psi Energy Inc	8	132	962	3.7	63	405
Virginia Electric & Power Co	8	23	958	3.9	12	415
Southern Company	12	30	925	6.5	16	390
American Electric Power	10	55	>1000	5.1	27	>404
Dairyland Power Coop	8	62	755	4.0	30	304
Dayton Power & Light Co	8	18	935	3.9	9	366
Atlantic City Elec Co	8	17	754	4.0	9	303
Cincinnati Gas & Electric Co	7	40	927	3.4	20	360
Indianapolis Power & Light Co	8	77	758	3.9	36	304
Public Service Co Of Nh	8	25	752	3.9	13	302
Savannah Electric & Power Co	8	40	752	3.9	20	300
Southern Indiana G & E Co	9	18	763	4.4	9	305
Central Hudson Gas & Elec Corp	8	18	754	3.9	9	301
Central Illinois Light Co	8	24	750	3.9	12	300
Central Operating Co	8	41	754	3.9	19	301
Empire District Electric Co	8	19	754	3.9	10	277
Interstate Power Co	8	40	754	3.9	20	274
Madison Gas & Electric Co	8	80	754	3.9	39	303
Minnesota Power Inc	8	74	755	3.9	36	302
Northern Indiana Pub Serv Co	8	29	755	3.9	15	303
Northern States Power Co	8	56	759	3.9	26	297
Rochester Gas & Elec Corp	8	39	751	3.9	19	303
Southern California Edison Co	7	18	908	3.3	9	384
St Joseph Lgt & Pwr Co	8	76	754	3.9	35	300
Tampa Electric Co	8	20	760	3.9	10	303
Holyoke Wtr Pwr Co	8	17	751	3.9	8	298
Ohio Valley Electric Corp	8	18	761	3.9	9	302

Each firm's problem is computed by 1 of 64 cores on a 256GB RAM machine on the Odyssey research computing clusters at the Faculty of Arts and Sciences at Harvard University. The number of approximation nodes is 4, the degree of the Chebyshev polynomials is 3, and the stopping criterion is 10^{-6} .

 $V^{(k-1)}(s_0)$ in Equation (16), and comparing with Equation (15), we have:

$$L^{(k)} = \frac{\beta}{1-\beta} (V^{(k)}(s_0) - V^{(k-1)}(s_0)).$$

Thus, the iterative value difference in the normalizing state upon convergence scaled by $\frac{\beta}{1-\beta}$ yields the level difference between the relative value function and the full value function.