

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

Cuicui Chen*

January 15, 2019

Abstract

I study firms' belief formation in a new market by examining private electric utilities' beliefs about the future sulfur dioxide allowance price under the Acid Rain Program, the first large-scale cap-and-trade program. I estimate their beliefs from 1995 to 2003 using a firm-level dynamic model of allowance trades, coal quality, and pollution reduction investment. I find that firms initially underestimate the role of market fundamentals as a price driver, but over time their beliefs converge toward the stochastic process of allowance prices. Beliefs have important implications for implementing cap-and-trade programs under utility regulation, and change their efficiency relative to taxes.

*Harvard University. Email: cuicui_chen@hks.harvard.edu. This is the first chapter in my doctoral dissertation. 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.

Beliefs matter. Investors choose portfolios under beliefs about future prices (Keynes, 1936), manufacturers install technologies under beliefs about future competitiveness (Dixit and Pindyck, 1994), and retailers manage inventories under beliefs about future demand (Bajari et al., 2018). Forming beliefs is no easy matter, and likely more challenging when firms find themselves in a new market. This paper studies firms' beliefs about the future price of a new commodity created by government regulation, and draws implications for firm value, consumer expenditure, and environmental policy.

My setting is the U.S. Acid Rain Program (1995-present), the first large-scale cap-and-trade program and a grand experiment in market-based environmental policy (Stavins, 1998). The program created an unprecedented market for sulfur dioxide allowances, each permitting the allowance holder to emit one ton of sulfur dioxide in production. Every year, electric utilities receive allowances from the government for free. They are required to hold enough allowances to cover their sulfur dioxide emissions, using whatever compliance methods they choose: burning clean fuel, installing pollution-reduction technologies, or buying allowances. Because allowances can be saved for the future, firms' beliefs about the future allowance price are the key to their compliance behavior.

I use dynamic structural estimation to back out the beliefs held by coal-dependent private electric utilities about the future allowance price that rationalize their compliance behavior during the first decade of the Acid Rain Program. The intuition is that the higher the firms expect tomorrow's allowance price to be, the more aggressively they stock up on allowances today, by reducing the sulfur content of coal, making more emission-reduction investment, and buying more allowances. Having estimated those beliefs, I compare them with full-information beliefs (Rust, 1994) and adaptive beliefs (Sargent, 1993). Full-information beliefs coincide with the stochastic process of the allowance price estimated with a full series of data (as in Collard-Wexler (2013)), and adaptive beliefs with data only up to the point of firms' decision making in an iterative fashion (as in Jeon (2018)).

I have three findings. First, firms' beliefs deviate from both the full-information and the adaptive benchmarks; firms underestimate the role of market fundamentals as a driver of the allowance price, by putting too little weight on the conditions of those markets that drive the allowance price (such as electricity and coal markets) and too much on the allowance market itself (historical allowance prices).¹ Hence, the usual practice of structural industrial organization that endows firms with full-information beliefs may not apply to new markets.

¹Some may wonder why the allowance price did not turn out to be consistent with the expectations of those coal-dependent private electric utilities according to the rational expectations hypothesis. Note, however, that those firms are only a subset of the participants in the allowance market; other important participants include coal-independent utilities, public utilities, non-utility energy firms, and brokerage firms.

Given that the data relevant for price prediction were timely and widely circulated,² this belief bias implies that firms may not be using their information effectively. Second, bigger firms, and firms facing more competitive pressure, exhibit less biased beliefs. This is in line with the findings from Tanaka et al. (2018) that larger and more cyclically sensitive firms make more accurate forecasts of future market conditions (GDP). Third, over time, those beliefs converge toward the adaptive beliefs. This is consistent with the shift in the management practice in electric utilities that compliance decisions are made less by engineers but more by people experienced in markets and trading (Reinhardt, 2000).³

Beliefs have important implications for the well-being of firms and consumers. Biased beliefs in the first five years of the Acid Rain Program cost firms an average expected discounted sum of payoffs equivalent to around 10% of their profits. Under cost-of-service regulation of electric utilities, those forgone payoffs to utilities are the lower bound on the forgone savings to ratepayers. Therefore, policies that improve the belief formation process would be financially beneficial to ratepayers. Such policies include holding workshops and conferences to facilitate communication among market participants and the regulator, introducing more competition, and using price ceilings and floors to constrain the volatility of the allowance price. Those policies are highly relevant today as cap-and-trade programs, the predominant climate policy, are being adopted by many regions and countries (such as China and EU) as their first market-based environmental policy and usually start with the electric generating sector, the primary source of anthropogenic carbon dioxide emissions. More broadly, as government regulation routinely creates large markets for commodities that have never been traded before (*e.g.*, medallions for taxicabs, Renewable Identification Number credits for renewable fuels, and Corporate Average Fuel Economy credits), it is important to address the implications of biased beliefs for firms and consumers, especially in the early years of those markets.

Another policy implication is that those beliefs change the relative efficiency of cap-and-trade programs and taxes. The former is more efficient when beliefs align the marginal cost of emission reduction with the marginal benefit of doing so.⁴ For example, given my

²Unlike the early environmental regulation, information transparency has been emphasized from the inception of the Acid Rain Program (Napolitano et al., 2007); allowance trading data were made available every month (before the Internet made them instantly available) by the government and allowance price data by brokerage firms. Furthermore, the electric utility industry has easy access to electricity and fuel market data, and is itself a close-knit community with frequent communication.

³Some have suggested that public utilities commissions, overseeing the utilities under cost-of-service regulation, may also have adapted in how they regulate utilities in the presence the Acid Rain Program (Bailey, 1998). Most of that adaptation should have already occurred before 1995, the beginning of my period of analysis, by which time the utilities would have submitted compliance plans to public utilities commissions for review, and many researchers debated (*e.g.*, Rose et al. (1992, 1993)), and regulators implemented, necessary changes to cost-of-service regulation (summarized by Lile and Burtraw (1998)).

⁴Full efficiency would require each firm reduce its emissions up to the point where its marginal cost of

empirical finding that larger electric utilities tend to forecast better by not over-predicting allowance prices as much as do smaller utilities, the former would be less aggressive in reducing emissions. The resulting lower marginal cost aligns with their lower marginal benefit of emission reduction, as larger utilities tend to locate in less populated areas. This alignment pushes cap-and-trade to be more efficient than tax. This finding that beliefs act as a decentralized efficiency-changing channel adds to the debate on the relative merits of cap-and-trade and tax, which has so far used only static models and proposed government intervention to change efficiency (Muller and Mendelsohn, 2009).⁵

This paper makes three technical contributions. First, it constructs the first fully dynamic empirical model of firm behavior in cap-and-trade programs. Dynamics is indispensable because all existing cap-and-trade programs allow firms to save allowances for the future, creating inter-temporal incentives.⁶ In my model, 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 constraint that it hold enough allowances to cover the sulfur dioxide emissions. The allowance price beliefs are the perceived law of motion for the allowance price state variable. The model can be adapted to modeling firm behavior in other commodity markets created by government regulation, such as markets for medallions and renewable credits. 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 computational burden of estimating and simulating my three-dimensional continuous-state dynamic model. Those techniques will be useful to any structural estimation and counterfactual simulation that involve solving a dynamic problem. Third, this paper and Toyama (2018) are the first to match extensive allowance trading data to electric operations data for academic research.⁷ This data exercise will enable studies of firms' trading behavior in any

emission reduction equals the marginal benefit it creates.

⁵There is a strand of literature inspired by the seminal paper of Weitzman (1974) on the efficiency comparison between cap-and-trade and tax under uncertainty. In that literature, the government faces uncertainty about firms' marginal cost of emission reduction, but firms know their own marginal costs perfectly. In this paper, however, firms themselves face uncertainty about the future allowance price, and thus their marginal costs.

⁶Fowle et al. (2016) develop a fully dynamic model of the cement industry under various environmental policies, including a cap-and-trade program, but they focus on entry and exit decisions and assume away allowance trading and allowance price uncertainty. Toyama (2018) develops a finite-horizon model of utility behavior in the Acid Rain Program. Other empirical models of firm behavior in cap-and-trade programs are static, including Fowle (2010), Cicala (2015), and Chan et al. (2017).

⁷The administrator of the Acid Rain Program, the 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.

cap-and-trade program.

My paper contributes to the burgeoning literature on firm behavior in new markets. 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. Hortacısu and Puller (2008) and Hortacısu et al. (2017) study firms' bidding behavior in the newly 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. This literature has thus focused on static decisions in competitive environments. This paper focuses on dynamic decisions in a single-agent environment.

The literature on belief formation (Coibion and Gorodnichenko, 2015; Gillen et al., 2017; Coibion et al., 2017; Tanaka et al., 2018), mostly relying on surveys of forecasts, has consistently documented biased beliefs in firms. For example, Coibion et al. (2017) find that managers in New Zealand systematically over-predict inflation, and Gillen et al. (2017) find that Intel tends to overstate expected sales. A common concern with the use of stated beliefs is whether the survey respondents report those beliefs truthfully. Tanaka et al. (2018), for example, conduct multiple checks to argue for truthful reporting. My "revealed preference" approach of recovering beliefs from behavior obviates this concern, as does Greenwood and Hanson (2015).

Several papers have studied various forms and sources of inefficiency in the Acid Rain Program and other cap-and-trade programs. Carlson et al. (2000) document 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. Fowlie (2010) (in the context of the NO_x Budget Program, a later cap-and-trade program) and Cicala (2015) point to electric utilities' distorted incentive created by cost-of-service regulation, and Toyama (2018) finds substantial transaction cost in the allowance market. I document another source of inefficiency, namely biased beliefs, after controlling for the distorted incentive under cost-of-service regulation and the transaction cost. This paper also complements the finding in Montero and Ellerman (1998) that the early projections of allowance prices published by the trade association of the electricity sector contain expectation errors that over-predict the low-sulfur coal price and therefore overstate the future allowance price.

The next section describes the institutional background, the data, and suggestive evidence for biased beliefs. That motivates a structural model that features flexible beliefs about the future allowance price in Section 2. Section 3 estimates the model using a nested-fixed-point algorithm with a simulated maximum likelihood estimator. Section 4 investigates the implications of biased beliefs for firms, consumers, and environmental policy. Section 5 concludes.

1 Background and Data

1.1 Background on the Acid Rain Program

Legislated in the 1990 Clean Air Act Amendments and administered by the Environmental Protection Agency (EPA), the Acid Rain Program was designed to cut acid rain by reducing sulfur dioxide emissions from electric generating plants to about half their 1980 level. Phase I (1995-1999) covers 263 largest, dirtiest units, almost all coal-fired. Phase II (2000-present) covers all fossil-fuel-based units exceeding 25MW generating capacity. Around 180 units that were originally covered by Phase II only voluntarily opted into Phase I.⁸ 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, which have been covered since 2000.

Each unit receives a pre-determined number of allowances every year in its compliance phase.⁹ To comply, each unit needs to hold enough allowances by the end of each compliance year to cover its emissions. The EPA deducts the number of allowances equal to the emissions every year. Note that a utility (firm) typically operates several electricity-generating plants, each of which typically has several units. Firms can trade and bank allowances.

Almost all sulfur dioxide emissions from electric utilities come from units that burn bituminous coal. Utilities can reduce sulfur dioxide emissions by using lower-sulfur bituminous coal, switching to the ultra-low-sulfur, sub-bituminous coal, or retrofitting the unit(s) with a flue-gas desulfurization equipment, or a “scrubber”.¹⁰

⁸Montero (1999) studies the voluntary participation in detail.

⁹The 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 sulfur dioxide per MMBtu of heat input for Phase I and around 1.2 lb/MMBtu for Phase II.

¹⁰Reduced utilization is rarely used as a means of compliance. Indeed, utilities are required to meet the electricity demand under cost-of-service regulation, and moving production from Phase I units to Phase II units while in Phase I was heavily penalized by the Acid Rain Program because of emissions leakage. Toyama (2018) find that the economic significance of the reduction in the capacity factor is very limited, and that emission reduction is achieved mainly through reduction in emission rates, by using lower-sulfur bituminous coal and switching to sub-bituminous coal.

For the first emission-reduction option, note that bituminous coal has a continuous and wide range of sulfur contents, from 1.5 to 8 pounds sulfur dioxide per MMBtu of energy input. Utilities burning higher-sulfur bituminous coal can thus purchase lower-sulfur bituminous coal from the spot market, or request contract coal suppliers to reduce the sulfur content in future shipments (by mining at a different depth of coal seam, or washing the mined coal differently in coal preparation plants). Little capital modification is needed for bituminous-coal boilers to burn lower-sulfur bituminous coal. Therefore, I model the sulfur content of bituminous coal as a continuous choice, and determine the cost of bituminous coal from an estimated coal price function with sulfur content as an attribute.

The second emission-reduction option is switching to sub-bituminous coal, with very low sulfur content ranging from 0.5 to 1.5 lb/MMBtu. Utilities pay an average of \$45 (in 1998 dollar) per kW installed capacity for the necessary capital adjustment.¹¹ Because the installed capacity of units is naturally discrete, I model fuel-switching as a discrete choice problem. Furthermore, since utilities typically plan abatement investments such that they would take effect by the start of a compliance phase,¹² and given that my period of analysis starts after Phase I begins,¹³ I model the discrete fuel-switching investment as one-time choice in a year prior to Phase II.

The third option, installing a scrubber, is the most expensive investment and requires three years to install. It works best with high-sulfur coal. The scrubbing decision is irrelevant to my model for two reasons. First, my period of analysis starts after the inception of Phase I, three years after when utilities made the decision to install scrubbers for Phase I. Second, scrubbing was irrelevant for firms with only Phase II units, because by definition, they are much cleaner than Phase I units and would not find scrubbing worthwhile. Indeed, all scrubbers built for compliance started operating from late 1994 to late 1995, just in time for Phase I.

¹¹Sub-bituminous coal is predominantly mined in the Powder River Basin in Wyoming. The price at the mine mouth 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. Montero and Ellerman (1998) find that the early projections of allowance prices published by the trade association of the electricity sector did not adequately appreciate this price decline.

¹²There are two reasons for this. 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 any capital 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, virtually all fuel-switching investment for compliance purposes took effect around 1995 and around 2000.

¹³Thus, I take as given the pre-1995 scrubbing and fuel-switching investments. This is because they reflected firms' beliefs about the allowance price when the allowance price observations were insufficient to generate a meaningful comparison benchmark.

Apart from reducing emissions, utilities can buy allowances to comply. Allowance trades are mostly handled by brokerage firms, without a central exchange. Allowances have vintages, and only those allowances of the current and the prior vintages can be used for current-year compliance.¹⁴¹⁵

I focus on private electric utilities in this paper. They were responsible for the vast majority of electricity generation and sulfur dioxide emissions in Phase I, and the ones that still exist after the electricity sector restructuring that occurred around 2000 (the start of Phase II) are disproportionately coal-fired. Private electric utilities are subject to cost-of-service regulation by public utilities commissions. Utilities must meet the electricity demand at a pre-determined electricity rate, which aims to cover prudently-incurred operating costs and a fair rate of return on capital. Unanticipated costs are passed on to ratepayers upon the commission's approval (Joskow, 1972).

While modeling private electric utilities simplifies the analysis because they operate in a single-agent decision-making environment due to cost-of-service regulation, it introduces distorted incentives. Lile and Burtraw (1998) describe in detail the regulatory treatment of the costs and revenues incurred while firms comply with the Acid Rain Program. Generally, both the prudently-incurred expenditure on lower-sulfur coal and allowances and the revenue from selling allowances are passed on to ratepayers. Approved capital investment cost is added to the rate base. My model takes into account the distortion of incentives by those cost-recovery rules.

1.2 Data Sources and Summary Statistics

I compile data from 1995 to 2003 on electricity production and environmental compliance at the unit, plant, and utility levels from publicly available sources; post-2004 allowance prices varied wildly due to unanticipated regulatory uncertainty, which is beyond the scope of this paper.¹⁶ 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), also based on EIA data. The Acid Rain Program compliance data are from EPA's Air Markets Program Data

¹⁴For example, if a unit emitted 4000 tons of sulfur dioxide in 1997, the unit account needs to hold at least 4000 allowances of vintage 1997 or prior.

¹⁵The EPA also holds a small allowance auction every March to auction off 2.8% of the annually allocated allowances and returns revenue to compliance units. Because of the small volume, I do not model allowance auctions.

¹⁶The 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.

(AMPD) system. Part of AMPD data, the allowance transfers, contain measurement errors, because firms are required to report to the AMPD only the allowances they intend to use for current-year compliance. The allowance transfers data may thus 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. My estimation addresses those measurement errors. Lastly, 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, the biggest allowance broker. Appendix A further describes data collection and processing.

My sample includes 42 coal-dependent private electric utilities that operate un-scrubbed bituminous-coal units subject to at least one phase of the Acid Rain Program.¹⁷ A utility is coal-dependent if its sulfur dioxide emissions from coal account for more than 90% of its total sulfur dioxide emissions; this avoids the need to also model firms' non-coal choices. The restriction to utilities that operate un-scrubbed bituminous-coal units ensures a non-trivial compliance problem. The sample accounts for around 53% of the allowance trading volume during my period of analysis.

Table 1 reports firms' compliance requirement and pollution reduction investment behavior.¹⁸ Of those that need to comply 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 that need to comply only after 2000, none chose scrubbing, and five chose switching. Those that neither scrub or switch fuel reduced the sulfur content of bituminous coal or bought allowances to comply. Table 2 reports summary statistics, including only the un-scrubbed bituminous-coal units subject to the Acid Rain Program.

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

¹⁷As 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 sulfur dioxide 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, 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 sulfur dioxide 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.

¹⁸In each row in Panel A, the number of private utilities is larger in 1995 than 2000 because of the electricity restructuring that happened around 2000.

Table 1: Compliance requirement and pollution reduction investment.

Panel A: Compliance requirement (Phase designation of subsidiary coal units)	Number of firms	
	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	16
Total	42	32

Panel B: Pollution reduction investment	Number of firms	
	complying since 1995	complying since 2000
Scrubbing	13	0
Fuel switching	10	5
None	3	11
Total	26	16

A firm is a coal-dependent private electric utility that operates at least one un-scrubbed bituminous-coal unit subject to at least one phase of the Acid Rain Program.

Table 2: Summary statistics.

Variable	Mean	Median	Std.dev.	Min	Max
Firms that comply since 1995, 1995-1999					
Allowance stock	90408	34944	189755	-96704	1336528
Heat input (trillion Btu)	101.3	48.5	150.8	5.1	784.0
Sulfur dioxide emissions (ton)	128380	57780	164746	7144	721013
Allowance allocation	157099	90749	214698	8223	998840
Sulfur content (ton/trillion Btu)	1339.9	1322.3	474.0	467.1	2653.6
Net allowance purchase	12503	1633	63361	-303027	343634
Firms that comply since 1995, 2000-2003					
Allowance stock	235411	106567	363746	2887	1699997
Heat input (trillion Btu)	240.4	66.8	363.4	9.9	1298.2
Sulfur dioxide emissions (ton)	187813	97652	260308	4790	950592
Allowance allocation	100220	56655	138320	5611	617632
Sulfur content (ton/trillion Btu)	990.1	957.6	372.5	378.4	1988.9
Net allowance purchase	23854	0	67004	-92877	277180
Firms that comply since 2000, 2000-2003					
Allowance stock	74570	39064	117031	-6603	578980
Heat input (trillion Btu)	102.4	50.1	130.0	1.2	412.1
Sulfur dioxide emissions (ton)	67891	32383	84565	377	277493
Allowance allocation	60650	23078	85053	838	472237
Sulfur content (ton/trillion Btu)	770.1	666.1	420.1	307.4	1749.5
Net allowance purchase	1764	0	30035	-124224	106001

Only non-scrubbed bituminous-coal units subject to the Acid Rain Program are included.




Figure 1: Left: Observed and predicted one-year-ahead Phase I allowance prices based on 1993-1999 data. Right: Observed and predicted one-year-ahead Phase II allowance prices based on 1993-2002 data.

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) in Table 3 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:

$$P_t = a + cP_{t-1} + \epsilon_t, \quad t \leq 2000$$

for Column (1) and

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

for Column (2). Adding pre-2000 trends substantially improves the fit. Column (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. Figure 1 plots 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 figure shows that trends appear important for Phase I allowance price prediction, but not for Phase II. This informs the specification of beliefs in the dynamic model in Section 2. Appendix B provides additional data patterns that inform the model in Section 2.

Table 3: One-year-ahead current- and prior-vintage allowance price.

	(1)	(2)	(3)	(4)
	93-99	93-99	93-02	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)	48.85 (11.46)
N	78	78	114	114
Adj. R-sq	0.029	0.557	0.404	0.329

Robust standard errors are in parentheses. All prices are in 1995 January dollar.

1.3 Suggestive Evidence for Biased Beliefs

I offer three pieces of suggestive evidence for biased beliefs held by firms about the future allowance price.

Biases in early price projections. Figure 2 shows that before the start of Phase I in 1995, all major predictions of the Phase I (1995-1999) average allowance price, summarized by Hahn and May (1994), exceeded the actual prices. Most of those predictions exceeded the actual prices by a large margin. The expected 1995, 1997, and 1999 allowance prices surveyed by Fieldston Company (1993) were no better. To be clear, some of the discrepancies necessarily come from the difference between ex-ante expectations and ex-post realizations. Nevertheless, the fact that all those early price expectations consistently and substantially exceeded the actual prices implies that more forces may be at play. My private communication with utility executives, regulators, and allowance brokers reveals the utilities' initial difficulty in appreciating the nascent 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." My analysis looks at their beliefs about the future allowance price after Phase I began, as market information became more available and utilities gained experience with the allowance market.

Biases in future-vintage prices. Firms trade and use current- and prior-vintage allowances to cover the current-year emissions, but they can also trade allowances of future vintages. If firms had reasonable beliefs about the allowance market evolution, the future-vintage allowance prices would on average be comparable to the discounted actual price that realize in that vintage year. Figure 3 uses partial data on the next-vintage allowance price to show that such comparability is likely absent; the next-vintage price from the previous year differs consistently and substantially from the actual price in the current year.

Seemingly irrational compliance behavior. Figure 4 plots the distributions of sulfur contents and allowance trades against the actual allowance price in Phase I. The allowance price dipped from 1995 to 1996 and then started a steady increase. As shown in the left panel of Figure 1, the dip and the increase are easily predictable. If firms indeed expected this dip and increase in 1996, they would buy more allowances than they did in 1995. By doing so, firms would benefit both statically (because the allowance price was lower) and dynamically (because the allowance price was going up). The left panel in Figure 4 shows that the industry-level median net allowance purchase decreased in 1996 from 1995. Of the 26 firms that comply since 1995, 15 reduced their net allowance purchase in 1996.

Furthermore, as the allowance price was predictably going up in 1997, firms should reduce the sulfur content in 1997 relative to 1996. By doing so, they would benefit both statically (as the allowance price was higher) and dynamically (as the price was still going up). The




Figure 2: Phase I average allowance price forecasts by various institutions (Hahn and May, 1994), 1995, 1997, and 1999 allowance price forecasts by survey respondents (Fieldston Company, 1993), and the actual allowance price from the beginning of the allowance market.




Figure 3: Monthly average bids for a current- and prior-vintage allowance and, when available, for a next-vintage allowance for immediate settlement in the previous year.

right panel in Figure 4 shows, on the contrary, that the median sulfur content peaked in 1997. Of the 26 firms complying since 1995, 12 increased their sulfur content in 1996.




Figure 4: Left: Distribution of net purchase of current- and prior-vintage allowances in Phase I. Right: Distribution of sulfur content in Phase I.

Table 4 confirms this puzzle using regressions. I regress sulfur content and net allowance purchase on the current allowance price and the predicted one-year-ahead allowance prices according to full-information beliefs in columns (1) and (3), and adaptive beliefs in columns (2) and (4), after controlling for firm fixed effects. The idea is that fixing today’s allowance price, the higher the expected price tomorrow, the lower the sulfur content and the more the net allowance purchase. Table 4 shows that neither response is statistically significant and has the right sign. My structural analysis to follow will explicitly compare the estimated beliefs and benchmark beliefs while accounting for alternative sources of friction: the allowance transaction cost and the distortion of incentives under cost-of-service regulation.

2 Model

This section presents a discrete-time infinite-horizon single-agent dynamic model of a coal-dependent private electric utility subject to the Acid Rain Program. Each year, the firm pays for coal to produce electricity that meets the demand, as required by cost-of-service regulation. The coal price is a function of sulfur content. The firm chooses the sulfur content, which, together with the heat input (the amount of energy in coal), determines the sulfur dioxide emissions. The firm also trades allowances, which incurs expenditure or revenue depending on whether the net purchase is positive. The firm inherits allowance stock from the previous year, and receives allowances for the current year; it needs to make sure that the sum of those and the net allowance purchase covers its sulfur dioxide emissions. It

Table 4: Dependence of sulfur content and net allowance purchase on current allowance prices and predicted one-year-ahead allowance prices.

	Sulfur content (ton/trillion Btu)		Net allowance purchase	
	(1)	(2)	(3)	(4)
Current allowance price	-0.774 (0.437)	-0.768 (0.447)	-211.98 (149.59)	-225.23 (149.74)
Predicted one-year-ahead allowance price	0.419 (0.626)	0.296 (0.478)	-138.85 (176.27)	-73.32 (121.80)
Constant	2626.0 (76.7)	2641.2 (63.8)	43842.5 (21570.5)	37021.9 (17727.9)
Firm FE	Y	Y	Y	Y
N	120	120	102	102
Adj. R-sq	0.912	0.912	0.077	0.074

Robust standard errors are in parentheses. All prices are in 1995 January dollar. Columns (1) and (3) use predicted one-year-ahead allowance price from full-information beliefs, and columns (2) and (4) from adaptive beliefs.

can also invest in 1999 to switch some capacity to sub-bituminous coal.

Formally, the decision variables are the sulfur content of coal, the net allowance purchase, and the fuel-switching investment. The state variables are the allowance stock, the heat input required to meet the electricity demand, and the allowance price. The allowance stock is controlled by the firm through sulfur dioxide emissions and the net allowance purchase, and evolves deterministically. The heat input is exogenous because consumers face a fixed electricity price under cost-of-service regulation and thus have exogenous electricity demand. I adopt the standard assumption that the firm’s beliefs about the future heat input coincide with its stochastic process. The allowance price is exogenous to the firm.¹⁹ Unlike the beliefs about the future heat input, those about the future allowance price do not have to coincide with its stochastic process.

Below I first describe a flexible specification of the beliefs about the future allowance price. I then state the firm’s problem in Phase II (2000-present). A key observation is that the firm would optimally choose the sulfur content and the net allowance purchase such that their “marginal costs” equalize, where the marginal costs incorporate the allowance transaction cost and the distortion of flow-cost incentives under cost-of-service regulation. This equation between those marginal costs allows me to identify the transaction cost parameter and the flow-cost distortion parameter, while the magnitude of those marginal costs reveals their

¹⁹In existing research and qualitative accounts of the Acid Rain Program, firms have not been found to exercise market power in the allowance market.

beliefs about the future allowance price and helps identify the belief parameters. Finally, I state the firm's problem in Phase I (1995-1999) and takes the Phase II value functions as its terminal values. During Phase I, in addition to choosing sulfur content and net allowance purchase, the firm makes a fuel-switching investment decision in 1999. This helps identify the capital-cost distortion parameter.

2.1 Beliefs about the Future Allowance Price

I assume that the firm believes in year $t \in \{1995, 1996, \dots, 2000\}$ that the allowance price in year $s > t$ evolves according to the following process:

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

where ϵ_s is an i.i.d. normal error with mean zero and standard deviation b_5 . Thus, the firm forecasts future allowance prices based on historical allowance prices via the slope ($b_1 + b_2 \times \min\{s - 1, 2000\}$), and the conditions of related markets (*e.g.*, coal and electricity markets) via the intercept ($b_3 + b_4 \times \min\{s - 1, 2000\}$). I allow the belief to be non-stationary via the time trends in the slope, b_2 , and in the intercept, b_4 . This captures the notion that Phase I is the first five years of the novel allowance market, when the price formation process may still be evolving. The time trends vanish in 2000, both to capture the perception that allowance price formation will settle down by Phase II and to facilitate the dynamic programming of the infinite-horizon Phase II problem, which requires stationary state transitions.

The parameters of interest are $(b_1, b_2, b_3, b_4, b_5)$, which capture the firm's beliefs about the future allowance price while the firm is in Phase I. They will come from structural dynamic estimation in Section 3, rather than from the estimation of the stochastic process in Section 1. To make the estimated beliefs comparable to the full-information beliefs, I will use all data from 1995 and 1999 to recover pooled estimates of the belief parameters. To compare to the adaptive beliefs, I will use each year of the data from 1995 to 1999 to recover year-specific estimates of the beliefs parameters.

For the beliefs about the future allowance price while the firm is in Phase II, I assume that in year $t \in \{2000, \dots, 2003\}$, the firm believes that the allowance price in year $s > t$ follows following process:

$$P_s = b'_1 + b'_3 P_{s-1} + \epsilon_s \quad (2)$$

where ϵ_s is an i.i.d. normal error with mean zero and standard deviation b'_5 . This is a stationary process. 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.

2.2 The Infinite-Horizon Phase II Problem

In each year $t \in \{2000, 2001, \dots, \infty\}$, 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 :

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

where the internalization functions translate the monetary value of those expenditures into what the firm deems payoff. They flexibly capture the distortionary effect of cost-of-service regulation. As a benchmark, a firm that is not subject to cost-of-service regulation would have $\phi_A(A) = -A$ and $\phi_C(C) = -C$, thereby fully internalizing the allowance and coal expenditures. Under cost-of-service regulation, however, allowance and coal expenditures are reimbursed, so full internalization is not likely. Neither is zero internalization, because reimbursement requires regulatory approval.²⁰

For instance, if the firm pays A for 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, and therefore $\phi_A = -0.2A$. 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 goodwill with the regulator. The firm receives $0.9C - C = -0.1C$ as the static payoff from burning coal, and therefore $\phi_C = -0.1C$.

The ability to save allowances for the future 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 has the number of allowances equal to the sulfur dioxide emissions $x_{it}H_{it}$ deducted from its allowance account. Thus, the allowance stock evolves as follows:

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

²⁰I have omitted the electricity revenue from the firm’s payoff. This is because the revenue is mostly out of the firm’s control under cost-of-service regulation, as both the electricity demand and the electricity rate are mostly exogenous.

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

$$V_i(W_i, P, H_i) = \max_{\substack{x_i \in X \\ a_i \geq x_i H_i - W_i - alloc_i}} \{ \phi_A[A(a_i; P)] + \phi_C[C_i(x_i; H_i)] \\ + \beta \int V_i(W_i + alloc_i + a_i - x_i H_i, P', H'_i) dF_P(P'|P) dF_{H_i}(H'_i|H_i) \},$$

where F_P is the allowance price belief of Equation (2), and F_{H_i} is the heat input process. 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’’, $a_i \geq x_i H_i - W_i - alloc_i$, requires that enough allowances be held to cover the sulfur dioxide emissions each year; after rearranging it becomes $W'_i \equiv W_i + alloc_i + a_i - x_i H_i \geq 0$.

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

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

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

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

where μ^* is the Lagrange multiplier for the compliance constraint.

The first-order conditions explain optimal behavior intuitively.²¹ 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 equals the marginal cost of acquiring additional allowances. This is regardless of whether the compliance constraint binds. Indeed, sulfur content reductions and allowances are perfect substitutes both statically (in terms of current production) and dynamically (in terms of their contribution to the future allowance stock). In Section 3, I use this equality to estimate the transaction cost and internalization parameters with-

²¹Appendix F 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.

out doing dynamic programming. Appendix C explains how I extend the model to include future-vintage allowance trading.

2.3 The Finite-Horizon Phase I Problem

The Phase I problem involves a fuel-switching investment problem in 1999. I model this investment as a discrete choice, as fuel-switching typically involves all bituminous units within a plant and plants are discrete.²² 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 realize. Third, the firm chooses which plant(s) to switch fuel. The investment materializes in 2000. To account for the distortion by cost-of-service regulation, the capital cost K is replaced by $\phi_K(K)$, where ϕ_K is the capital internalization function.

Having computed the 1999 value functions, the firm performs backward induction to solve each of the 1995-1998 problems. Appendix C lays out all the Bellman's Equations in Phase I.

2.4 Identification

We need to estimate the parameters in coal expenditure, allowance transaction cost, heat input process, coal and allowance internalization, capital internalization, and beliefs.

The first four parameters 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 transaction cost parameter and the coal and allowance internalization parameters are identified by the equality between the marginal cost of lower sulfur content and the marginal cost of allowances from Equations (4) and (5). The transaction cost parameter concerns how the

²²Since 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.

marginal cost rises with the transaction volume.²³ The heat input transition parameters are identified under standard identification assumptions of first-order autoregressive models.

Identifying the last two parameters requires solving the dynamic problem. Exogenous variations in the allowance price and the heat inputs lead to temporal and cross-sectional variations in the expected marginal value of allowances, which determines behavior. The last two parameters are separately identifiable because they have different implications for behavior. Higher capital internalization leads to less switching, more allowance purchase, and lower sulfur content, while a higher allowance price forecast leads to more switching, more allowance purchase, and lower sulfur content.²⁴

3 Estimation

My estimation strategy proceeds in two steps. First, I estimate the parameters that do not require dynamic programming: the coal expenditure parameters for each firm, estimated using fixed-effect regressions; the heat input transition parameters for each firm, estimated using time-series regressions; the transaction cost parameter and the internalization parameters for allowance and coal expenditures, estimated using an ordinary least-squares regression with a 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.²⁵ 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.

²³The 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.

²⁴I have adopted the standard assumption that firms are risk neutral. If we relaxed this assumption, then beliefs and risk attitude may not be separately identifiable.

²⁵I do not use the estimation approach (Bajari et al., 2007), because Phase I has a finite horizon and is thus non-stationary, and my data do not permit reliable estimation of year-specific policy functions.

3.1 Coal Price and Expenditure

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

$$p_{jkt}^{bit} = \gamma_0^{bit} + \gamma_1^{bit} x_{jkt} + \gamma_2^{bit} x_{jkt}^2 + \gamma_Z^{bit} Z_{jkt} + dummies + \epsilon_{jkt}^{bit} \quad (6)$$

where x is the sulfur content and Z is a flexible control variable vector, including the heat content, the ash content, the distance between the plant and the source county, and their interactions. The quadratic specification in 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 5 reports the estimation results. I use the estimates from Column (4).

The bituminous 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 have omitted from the coal expenditure the non-sulfur components of the bituminous coal price function (6); they do not affect the firm's choices in Phase I. Note that the coal expenditure in Phase I refers to the bituminous coal expenditure only, because I have defined a firm in my sample as a collection of its un-scrubbed and bituminous coal units (as of Phase I) as opposed to all units.

In the Phase II problem, firm i that switches capacity k_j from bituminous to sub-bituminous coal incurs both bituminous and sub-bituminous coal expenditures. The bituminous coal expenditure in Phase II 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 new sub-bituminous units, where $\bar{\alpha}_i$ is the heat input rate of those units. This expenditure now includes the non-sulfur components in the bituminous coal price, summarized by $\bar{\gamma}_i^{bit}$. This is because those components now 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 sub-bituminous coal expenditure in Phase II is calculated as follows. 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 the sub-bituminous coal price, p_{jkt}^{sub} , over the sub-bituminous coal shipments received by plants operated by firm

Table 5: Delivered bituminous coal price.

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

Standard errors clustered by utility are in parentheses. Coal price is in 1995 January US cent/MMBtu, 1995-2004.

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 below. First, I conduct a sub-bituminous coal price regression in the style of Equation (6), 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 is the industry-average sulfur content used by sub-bituminous coal units, and the other predictors are the average values of the sub-bituminous coal from the coal county yielding the lowest predicted coal price for firm i .

3.2 Heat Input Transitions

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

$$H_{i,t} = \gamma_{0,i}^H + \gamma_{1,i}^H H_{i,t-1} + \epsilon_{i,t}^H, \quad (7)$$

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:

$$\begin{aligned} \mathbb{E}(H_{i,t-1} \epsilon_{i,t}^H) &= 0, & \forall t, \\ \mathbb{E}(\epsilon_{i,t}^H \epsilon_{i,\tau}^H) &= 0, & t \neq \tau. \end{aligned}$$

I estimate the heat input transitions separately for the Phase I units only, and the Phase I and Phase II units taken together. Appendix D reports the estimates for Southern Company as an example.

3.3 Transaction Cost and Internalization of Flow Expenditures

I leverage the first-order conditions in Section 2 to estimate the transaction cost parameter 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 (observed) ratio of the next-vintage allowance price over the current-vintage, and θ_0 is the transaction cost parameter to be estimated.

The internalization functions for the allowance and coal expenditures are:

$$\begin{aligned} \phi_A(A) &= -\theta_A A, \\ \phi_C(C) &= -\theta_C C, \end{aligned}$$

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

$$\theta_C(-\gamma_1^{bit} + 2\gamma_2^{bit}x_{i,t}) = \theta_A(P_t + 2\theta_0a_{i,t}). \quad (8)$$

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 (8) now becomes:

$$-\gamma_1^{bit} + 2\gamma_2^{bit}x_{i,t} = \theta_AP_t + 2\theta_aa_{i,t}. \quad (9)$$

As described in Section 1, the allowance trades contain measurement errors. Appendix E explains how I correct for those measurement errors in estimation. Table 6 reports the results. 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.

Table 6: The allowance internalization parameter θ_A and the transaction cost parameter θ_a .

	OLS	Corrected
Allowance internalization parameter (θ_A)	0.636 (0.00986)	0.636 (0.00977)
Transaction cost parameter (θ_a)	0.0000562 (0.000645)	0.0000572 (0.0000288)
N	150	
Adj. R-sq	0.968	

Standard errors are in parentheses.

3.4 Beliefs and Internalization of Capital Expenditure

I estimate the belief parameters and the capital internalization parameter using a nested-fixed-point algorithm with a simulated maximum likelihood estimator. The inner loop of this algorithm 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. The outer loop of this algorithm looks for parameter values that maximize the simulated likelihood. The likelihood is induced by the i.i.d. distribution of measurement errors in the current- and

next-vintage allowance trades and the sulfur content.²⁶ 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. Appendix G describes the estimation procedure in detail, where I explain the use of the Relative and the Endogenous Value Function Iterations methods (Bray, 2017b,a) to vastly reduce the computational burden.

Table 7 reports the estimation results from pooling the observations within the same phase. I also present the estimated parameters of the stochastic process of the allowance price from Section 1, or the full-information beliefs. Recall that in the allowance price belief specification in Section 2, 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 underestimate the role of market fundamentals as an allowance price driver.

Figure 5 plots the one-year-ahead predictions of the allowance price under the pooled beliefs and the full-information beliefs using Phase I data (left) and Phase II data (right). 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. In Phase II, however, the confidence bands overlap more, and the difference in the means is much smaller.

Aside from the belief estimates, Table 7 also reports a capital internalization estimate of 0.847. This means that electric utilities deem fuel-switching investment cheaper than it really is. This is consistent with the Averch-Johnson effect (Averch and Johnson, 1962), where regulated 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 (Fowle, 2010) and (Cicala, 2015).

I now explore firm heterogeneity in the beliefs. Figure 6 compares the allowance price predictions by firms with above- and below-median coal capacity (left), and by firms in states that ultimately deregulated the wholesale electricity market and in states that remain

²⁶As noted in Section 1, 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. 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 (9) in the first stage, because this measurement error is in the dependent variable. Yet the standard errors of the coefficients are underestimated.

Table 7: Pooled beliefs, full-information beliefs, and capital internalization.

		<u>Phase I</u>		<u>Phase II</u>	
		Pooled	Full-information	Pooled	Full-information
Market fundamentals	Constant (b_1)	61.13 (22.27)	-15.39 (12.50)	220.19 (19.16)	250.08 (16.97)
	Year trend (b_2)	14.979 (6.548)	81.76 (9.313)		
Path dependence	Constant (b_3)	0.478 (0.107)	0.871 (0.0796)	-0.435 (0.088)	-0.721 (0.113)
	Year trend (b_4)	-0.101 (0.0412)	-0.463 (0.0518)		
	Standard deviation, (b_5)	28.66 (12.89)	23.04	29.69 (15.69)	24.31
	Capital internalization (θ_K)	0.847 (0.314)			

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 parentheses. Allowance prices are in 1995 January dollar.




Figure 5: One-year-ahead predictions of the allowance price under the pooled beliefs and the full-information beliefs, using Phase I data (left) and Phase II data (right), with 95% confidence intervals.




Figure 6: One-year-ahead predictions of the allowance price according to beliefs by firms with above- versus below-median capacity (left), and in states that ultimately deregulated the wholesale electricity market versus in states that remain under cost-of-service regulation (right), with 95% confidence intervals.

under cost-of-service regulation (right), based on the pooled estimates. It shows 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.

To investigate the evolution of beliefs, Figure 7 plots the year-specific estimates along with the adaptive beliefs. I estimate the former year by year using cross-sectional data from each year, and have estimated the latter by adaptively using data up to the point of decision making in Section 1. The overall pattern is similar to that in the pooled estimates in Figure 5; the effective beliefs are flatter than the adaptive beliefs as firms do not appreciate market trends sufficiently. Furthermore, the year-specific beliefs start wildly apart from the adaptive




Figure 7: One-year-ahead predictions of the allowance price under the year-specific beliefs and the adaptive beliefs with 95% confidence intervals.

beliefs, 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 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 bring in 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.

4 Implications of Biased Beliefs

In this section, I first examine the impact of biased beliefs on the dynamic payoffs of firms. 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 consumers under cost-of-service regulation. Second, I compare the efficiency of cap-and-trade program and an emission tax in a dynamic context, where beliefs about the future allowance price are the key to firms' compliance behavior. Appendix H presents additional policy experiments assessing the implications of biased beliefs for aggregate environmental and economic outcomes.

4.1 Firm Payoffs and Consumer Expenditure

To quantify the effect of biased beliefs on the dynamic payoffs of firms, I first simulate each firm's counterfactual choices of allowance trades and sulfur content iteratively from 1995 to 1998 under full-information beliefs. I then calculate the firm's discounted sum of static payoffs from 1995 to 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 1999 value function. Having obtained the counterfactual dynamic payoff of each firm, I repeat the steps above with the actual behavior to obtain the actual dynamic payoffs. The difference, normalized by the number of years in Phase I, is the average annual loss in the firm's dynamic payoff due to the biased allowance price belief it held during Phase I.

Table 8 reports the cost of biased beliefs to firms that have complied 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 (the profit margin is consistent with selected firms' Form 10-K SEC filings in 1995). Therefore, 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%.²⁷ There is a lot of heterogeneity in the forgone dynamic payoffs because of the utility sizes.

The forgone dynamic payoff to each firm is the lower bound on the forgone dynamic savings to ratepayers. This is because 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, where internalization is less than one. Policies that improve the belief formation process of individual firms are therefore financially beneficial to the

²⁷This 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 2.

Table 8: Cost of biased beliefs in Phase I to firms that have complied since 1995.

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

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

ratepayers.

4.2 Relative Efficiency of Cap-and-Trade and Tax

Beliefs change the efficiency of cap-and-trade relative to tax. Efficiency is the benefit minus the cost of emission reduction. Full efficiency would require that each firm reduce their emissions until its marginal cost of emission reduction equals the marginal benefit. Beliefs create a link between marginal costs and the marginal benefits under a cap-and-trade program, while they are mostly irrelevant under tax; frequent tax rate adjustment is politically unappealing. Therefore, when beliefs under a cap-and-trade program align marginal costs with marginal benefits, cap-and-trade is 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 tax, as beliefs about the future allowance price align the marginal cost and the marginal benefit of emission reduction.

To put this illustration in numbers, 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 9 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.

Table 9: Net benefits from an emission tax and a cap-and-trade program in the Southern Company - Atlantic City Electric Company example.

Scenario	Net Benefit
tax = 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

Previous literature misses beliefs as a potential determinant of the relative efficiency of cap-and-trade programs and emission taxes. This is because of its focus on static models, where firms care about current profits only and beliefs are therefore irrelevant. Muller and Mendelsohn (2009) propose that to improve the efficiency of cap-and-trade programs, government intervention is necessary, in the form of setting up trading ratios. I show that once we move beyond static models and use 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.

5 Conclusion

This paper studies firms' belief formation in a new market. The context is the Acid Rain Program, the first cap-and-trade program and a landmark experiment in market-based environmental policy. I structurally estimate coal-dependent private electric utilities' beliefs about the future sulfur dioxide allowance price. My dynamic model of allowance trades, coal quality, and fuel-switching investment is the first fully dynamic empirical model of firm behavior in cap-and-trade programs.

I find that firms have biased beliefs about the future allowance price. They 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 of the allowance price.

Beliefs and dynamics add new insights to the efficiency comparison between cap-and-trade programs and taxes. A tax simply equates marginal abatement costs across all firms, while dynamic considerations under a cap-and-trade program cause firms to equate their marginal abatement costs with firm-specific marginal dynamic value of allowances. Therefore, when beliefs about future market conditions align the latter with the marginal benefit of emissions reduction, so do the marginal abatement costs, making cap-and-trade programs more efficient than taxes. Therefore, beliefs are a decentralized channel that can align private interests with public benefits under a cap-and-trade program.

Following the broad success of the Acid Rain Program, many countries have adopted, or plan to adopt, the cap-and-trade approach to regulating pollution. Those cap-and-trade programs that regulate carbon dioxide emissions, most notably the ongoing EU Emissions 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 sulfur dioxide allowance market around one billion. A careful consideration of dynamics and beliefs will be helpful in choosing, designing, and evaluating cap-and-trade programs.

I conclude with four open questions. First, how do firms' beliefs evolve in the post-deregulation era in states that have deregulated the electricity sector? Competition should improve belief formation. Answering that question requires a competitive model of firm behavior in place of a single-agent model as in this paper.²⁸ 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?²⁹ If it does, then early cap-and-trade programs have learning value that a proper benefit-cost evaluation should include. Third, it would be interesting to formally identify the sources of belief biases and improvement; detailed organizational data could be collected to pin down the mechanisms through which shifts in the management practice in electric utilities impact belief formation. Last but not least, while this paper only estimates the beliefs as they are, it would be intriguing to investigate what learning rules may have led to those beliefs in the first place.

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.
- Bajari, P., Chernozhukov, V., Hortaçsu, A., and Suzuki, J. (2018). The Impact of Big Data on Firm Performance - An Empirical Investigation.
- 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.

²⁸It also potentially requires modeling the dispatch decision between coal and gas, and the investment decision to switch to or co-fire with gas, as gas became competitive with coal several years after deregulation.

²⁹I thank Kenon Smith, a former EPA employee, for suggesting this.

- 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.
- 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.
- Coibion, O. and Gorodnichenko, Y. (2015). Information Rigidity and the Expectations Formation Process : A Simple Framework and New Facts. *American Economic Review*, 105(8):2644–2678.
- Coibion, O., Gorodnichenko, Y., and Kumar, S. (2017). How Do Firms Form Their Expectations? New Survey Evidence.
- Collard-Wexler, A. (2013). Demand Fluctuations in the Ready-Mix Concrete Industry. *Econometrica*, 81(3):1003–1037.
- 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.
- Dixit, A. K. and Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press, Princeton, N.J.
- 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.
- Fowle, M. (2010). Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement. *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.
- Gillen, B. J., Plott, C. R., and Shum, M. (2017). A pari-mutuel-like mechanism for information aggregation: A field test inside intel. *Journal of Political Economy*, 125(4):1075–1099.
- Goldfarb, A. and Xiao, M. (2011). Who Thinks about the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets. *American Economic Review*, 101(7):3130–3161.
- Greenwood, R. and Hanson, S. G. (2015). Waves in Ship Prices and Investment. *Quarterly Journal of Economics*, pages 55–109.
- Hahn, R. W. and May, C. A. (1994). The Behavior of the Allowance Market: Theory and Evidence.
- Hortaçsu, A. and Puller, S. L. (2008). Understanding strategic bidding in multi-unit auctions : a case study of the Texas electricity spot market. *The RAND Journal of Economics*, 39(1):85–114.
- Hortaçsu, A., Puller, S. L., Luco, F., and Zhu, D. (2017). Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets.
- Jeon, J. (2018). Learning and Investment under Demand Uncertainty in Container Shipping.
- Joskow, P. L. (1972). *A Behavioral Theory of Public Utility Regulation*. PhD thesis.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. Macmillan, London.
- 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. (1999). Voluntary Compliance with Market-Based Environmental Policy : Evidence from the U . S . Acid Rain Program. *Journal of Political Economy*, 107(5):998–1033.
- 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.
- Napolitano, S., Schreifels, J., Stevens, G., Witt, M., LaCount, M., Forte, R., and Smith, K. (2007). The U.S. Acid Rain Program: Key Insights from the Design, Operation, and Assessment of a Cap-and-Trade Program. *Electricity Journal*, 20(7):47–58.

- 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.
- Rust, J. (1994). Structural Estimation of Markov Decision Processes. In Engle, R. and McFadden, D., editors, *Handbook of Econometrics*, volume IV, pages 3139–3139.
- Sargent, T. J. (1993). *Bounded Rationality in Macroeconomics.* Clarendon Press, Oxford, UK.
- 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.
- Stavins, R. N. (1998). What can we learn from the Grand Policy Experiment? Lessons from SO₂ Allowance Trading. *Journal of Economic Perspectives*, 12(3):69–88.
- Tanaka, M., Bloom, N., David, J. M., and Koga, M. (2018). Firm Performance and Macro Forecast Accuracy. 18(9):1–38.
- Toyama, Y. (2018). Incentives to Invest, Storable Permits, and Transaction Costs in Market-Based Environmental Regulation. pages 1–54.
- Weitzman, M. L. (1974). Price vs. quantities. *Review of Economic Studies*, 41(4):477–491.

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 SO₂/NO_x/CO₂ 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. If 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 intra-utility 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 1.

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 8 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 2 thus focuses on allowance trading in those vintages.




Figure 8: 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 9 plots the distribution of current- and prior-vintage net allowance purchases for utility-years during the compliance periods up to 2003. Figure 10 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 11 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




Figure 9: Distribution of current- and prior-vintage net allowance purchases during the compliance periods, utility-year, 1995-2003.




Figure 10: Distribution of next-vintage net allowance purchases since one year before the compliance periods, utility-year, 1994-2002.




Figure 11: Delivered price of coal and natural gas.

[Source] FERC Form 423.

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 12 shows that the massive drop in the natural gas price did not start until 2008.

Appendix C Technical Details on the Model

C.1 Allowing for Future-Vintage Allowance Trading

The model in the main text 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:

$$\begin{aligned}
 V_i(W_i, P, H_i) = & \max_{\substack{x_i \in X \\ a_i \geq x_i H_i - W_i - alloc_i \\ a_i^{nv} \geq -alloc_i}} \{ \phi_A[A(a_i, a_i^{nv}; P)] + \phi_C[C_i(x_i; H_i)] \\
 & + \beta \int V_i(W_i + alloc_i + a_i + a_i^{nv} - x_i H_i, P', H_i') dF_P(P'|P) dF_{H_i}(H_i'|H_i) \}.
 \end{aligned}$$




Figure 12: Natural gas price, 1990-2013, in 2010 U.S. dollars per thousand cubic feet. [Source] Davis (2015).

The additional constraint $a_i^{nv} \geq -alloc_i$ 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.

C.2 The Phase I Bellman Equations

Let firm i 's investment choice set be J_i , in which each choice j is characterized by the capacity to be switched, k_j .

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:

$$\begin{aligned}
V_i^j(W_i, P, H_i) = & \max_{\substack{x_i \in X \\ a_i \geq x_i \max\{H_i - \bar{\alpha}_i k_j, 0\} + \bar{x}_i \min\{\bar{\alpha}_i k_j, H_i\} - W_i - alloc_i \\ a_i^{nv} \geq -alloc_i}} \{ \phi_A[A(a_i, a_i^{nv}; P)] + \phi_C[C_i(x_i; H_i)] \\
& + \beta \int V_i(W_i + a_i + a_i^{nv} + alloc_i - x_i \max\{H_i - \bar{\alpha}_i k_j, 0\} - \bar{x}_i \min\{\bar{\alpha}_i k_j, H_i\}, P', H_i') \\
& \times dF_P(P'|P) F_{H_i}((H_i)'|H_i) \},
\end{aligned}$$

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:

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

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.

Therefore, the Bellman's Equation for Firm i 's 1999 problem is:

$$\begin{aligned}
V_i^{1999}(W_{i,1999}, P_{1999}, H_{i,1999}) = & \max_{\substack{x_i \in X \\ a_i \geq x_i H_{i,1999} - W_{i,1999} - alloc_{i,1999} \\ a_i^{nv} \geq -alloc_{i,2000}}} \{ \phi_A[A(a_i, a_i^{nv}; P_{1999})] \\
& + \phi_C[C_i(x_i; H_{i,1999})] + \log\left[\sum_{j=1}^{|J_i|} \exp(-\phi_K(k_j c^k) \right. \\
& \left. + \beta V_i^j(W_{i,1999} + alloc_{i,1999} + a_i + a_i^{nv} + x_i H_{i,1999}, P_{2000}, \tilde{H}_{i,2000}) \right. \\
& \left. \times dF_P(P_{2000}|P_{1999}) dF_{\tilde{H}_i}(\tilde{H}_{i,2000}|g_i(H_{i,1999}))\right] \},
\end{aligned} \tag{10}$$

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

Table 10: Heat input of Southern Company’s bituminous coal units.

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

Robust standard errors are in parentheses. Heat input is in million MMBtu, 1990-2004.

to Phase I units to that to Phase II ones.³⁰

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

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

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*.

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.

Appendix D Heat Input Transitions: Example Estimates

Table 10 reports the heat input transition estimates, $\gamma_{0,i}^H$, $\gamma_{1,i}^H$, and σ_i^H , for the Southern Company as an example.

³⁰The use of g_i avoids tracking the heat input to Phase II units as an additional state variable.

Appendix E Measurement Error Correction

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:

$$a_{i,t} = a_{i,t}^* + alloc_i \eta_{i,t}, \quad (11)$$

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 .³¹ 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 (9) and substituting in Equation (11), Equation (9) 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:

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

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.

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

$$\theta_a^{corrected} = \theta_a^{OLS} / \left(1 - \frac{var(\eta)var(\tilde{P})}{var(\tilde{a})var(\tilde{P}) - cov(\tilde{a}, \tilde{P})^2}\right)$$

$$\theta_A^{corrected} = \theta_A^{OLS} + 2 \frac{cov(\tilde{a}, \tilde{P})}{var(\tilde{P})} (\theta_a^{OLS} - \theta_a^{corrected})$$

³¹It 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^* .

Appendix F 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 2 are sufficient for optimality.

The Phase II value function is:

$$V_i(W_i, P, H_i) = \max_{\substack{x_i \in X \\ a_i \geq x_i H_i - W_i - alloc_i}} \{ \phi_A[A(a_i; P)] + \phi_C[C_i(x_i; H_i)] \\ + \beta \int V_i(W_i + alloc_i + a_i - x_i H_i, P', H'_i) dF_P(P'|P) dF_{H_i}(H'_i|H_i) \},$$

Letting $W'_i = W_i + a_i + alloc_i - x_i H_i$ 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 G Details on the Estimation of Belief and Capital Internalization Parameters

G.1 Inner Loop

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.³²

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:

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

where N_d is the degree of Chebyshev polynomials, $\text{coef}(d_1, d_2, d_3)$ 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 $s \in$

³²I 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.

$[\underline{s}, \bar{s}]$:

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

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

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

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 $(s_1, s_2, \dots, s_{N_j})$ such that:

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

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:

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

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 (7).

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 iterate at each exogenous state by the value at that exogenous state and the first approximation node of the endogenous state.³³

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 methods take an average of 1.2% and 6.2% as long as the full value function iteration 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

³³Bray (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 Iterations (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	Number of iterations			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} .

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:

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

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:

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

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

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

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

G.2 Outer Loop

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 al-

lowance stock:

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

subject to the compliance constraint $W_{i,t}^* + alloc_i + a_{i,t}^* - deduct_{i,t} \geq 0$ and the no-shorting constraint $(a_{i,t}^{nv})^* + alloc_i \eta_{i,t}^{a^{nv}} \geq -alloc_i$. I simulate many paths of $W_i = (W_{i,1995}, W_{i,1996}, \dots)$ 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 $W_{i,t}^* + alloc_i + a_{i,t}^* - deduct_{i,t} \geq 0$, 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;
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 $W_i = (W_{i,2000}, W_{i,2001}, \dots)$ 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:

$$\begin{aligned}
L_i(\theta) &[(a_{i,t}, a_{i,t}^{nv}, x_{it})_{t=t_1}^{t_2}, k_i | (W_{i,t}, P_t, H_{i,t})_{t=t_1}^{t_2}] = \int_{(W_{i,t})_{t=t_1}^{t_2}} \Pi_{t=t_1}^{t_2} \{ \phi_a [A_i(W_{i,t}^*, P_t, H_{i,t}; \theta) - a_{i,t}] \\
&\quad \times \phi_{a^{nv}} [A_i^{nv}(W_{i,t}^*, P_t, H_{i,t}; \theta) - a_{i,t}^{nv}] \phi_x [X_i(W_{i,t}^*, P_t, H_{i,t}; \theta) - x_{i,t}] Pr_i^k(W_{i,t}^*, P_t, H_{i,t}; \theta) \} \\
&\quad \times dF[(W_{i,t}^*)_{t=t_1}^{t_2} | (a_{i,t}, a_{i,t}^{nv}, W_{i,t})_{t=t_1}^{t_2}],
\end{aligned}$$

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 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 .³⁴ 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.

Appendix H Implications of Biased Beliefs for Aggregate Environmental and Economic Outcomes

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 3 suggest that smaller firms and firms with less competitive pressure tend to have more biased beliefs. How would the aggregate sulfur dioxide 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

³⁴For behaviors that are irrelevant to some t (for example, the 1999 fuel-switching investment behavior to 1998), their probability densities are excluded.

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-and-trade programs. What would be the aggregate environmental and economic implications of a price collar in the Acid Rain Program given the biased beliefs?³⁵

Columns 2 and 3 in Table 12 report the changes in the aggregate sulfur dioxide 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 12 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.

³⁵In 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 2 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; even all firms in my sample merely constitute around half of the allowance trade volume from 1993 to 2003.

Table 12: Aggregate effects of improving beliefs and reducing volatility.

Phase I (1995-1999)	<u>Improve beliefs of</u>		<u>Reduce volatility by</u>	
	Small firms	Firms always regulated	\$50 floor \$200 ceiling	Pre-1995-projected \$200 tax
Δ Aggregate emissions (thousand tons)	115.86	12609	2099	-9092
% Δ 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