



Essays on Environmental Economics and Industrial Organization

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:40050030>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Essays on Environmental Economics and Industrial Organization

A dissertation presented

by

Cuicui Chen

to

The Committee on Higher Degrees in Public Policy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Public Policy

Harvard University

Cambridge, Massachusetts

April 2018

© 2018 Cuicui Chen

All rights reserved.

Dissertation Advisors:
Professor Ariel Pakes
Professor Joe Aldy

Author:
Cuicui Chen

Essays on Environmental Economics and Industrial Organization

Abstract

This dissertation comprises three essays on environmental economics and industrial organization. A common theme for the essays is that they draw motivation from the real world and the theory literature to study individual behavior in environment and energy problems.

The first essay explores firm behavior in a new environment by examining private electric utilities' beliefs about sulfur dioxide allowance prices following the implementation of the U.S. Acid Rain Program in 1995. It questions the standard assumption in the empirical industrial organization literature that firms have reasonable expectations about future market conditions. This is partially motivated by strong qualitative evidence I have gathered from interviewing regulators, brokers, and electric utility executives that utilities were struggling to make sense of the allowance market, especially in the early years of this first large-scale cap-and-trade program.

In the second essay, co-authored with Richard Zeckhauser, we emphasize that in reality, nations are in substantially asymmetric situations. We then leverage this under-appreciated fact to construct a mechanism that leads to voluntary and Pareto efficient contributions to global public goods from individual nations.

In the third essay, I introduce repeated games with imperfect public monitoring from the theory literature to empirical industrial organization and explore partial identification of model primitives. This allows us to test firm conduct and detect collusion in contexts where firms do not observe each other's behavior perfectly, such as the operations of the Organization of the Petroleum Exporting Countries.

Contents

Abstract	iii
Acknowledgments	xi
Introduction	1
1 Slow Focus: Belief Evolution in the U.S. Acid Rain Program	3
1.1 Introduction	3
1.2 Institutional Background	10
1.3 Data	14
1.4 Model	22
1.4.1 Allowance Price Belief	22
1.4.2 The Phase II Problem	23
1.4.3 Fuel-Switching Investment	27
1.4.4 Finite-Horizon Phase I	29
1.4.5 Identification	30
1.5 Estimation	31
1.5.1 Coal Expenditure Parameters	32
1.5.2 Heat Input Transitions	34
1.5.3 Allowance Expenditure Parameter, Internalization Parameters for Al- lowance and Coal Expenditures	35
1.5.4 Inner Loop	38
1.5.5 Outer Loop	38
1.6 Implications of Biased Beliefs	46
1.6.1 Implications for Firm Payoffs	47
1.6.2 Implications for Aggregate Environmental and Economic Outcomes	50
1.6.3 Implications for the Choice between Cap-and-Trade and Tax	51
1.7 Conclusion	53
2 Collective Action in an Asymmetric World	56
2.1 Introduction	56
2.2 Climate Change Mitigation - A Traditional Alliance Problem	59

2.3	Cheap Riding in the Paris Agreement	61
2.4	Cheap Riding and Prior Proposals	64
2.5	Cheap-Riding Efficient Equilibrium (CREE)	66
2.5.1	Model Setup	66
2.5.2	Nash Equilibrium	67
2.5.3	Lindahl Equilibrium	69
2.5.4	CREE	71
2.5.5	Summary of the Numerical Example	76
2.6	Conclusion	78
3	Partial Identification in Repeated Games with Imperfect Public Monitoring	80
3.1	Introduction	80
3.2	Model	84
3.2.1	Set-up	84
3.2.2	Idea of Identification	85
3.3	Examples	87
3.3.1	Example 1: The Folk Theorem Obtains	87
3.3.2	Example 2: The Folk Theorem Does Not Obtain	91
3.3.3	Discussion	96
3.3.4	The Identified Set	97
3.4	Simulation	101
3.4.1	Unknown Signal Distribution Only	102
3.4.2	Mis-measured Private Histories Only	104
3.4.3	Both Unobserved Signal Distribution and Mis-measured Private Histories	106
3.5	Future Work	108
	References	109
	Appendix A Appendix to Chapter I	116
A.1	Data Compilation	116
A.1.1	Operations Data	116
A.1.2	Allowance Transfer Data	118
A.1.3	Other Data	121
A.2	Additional Data Patterns	121
A.3	Shape of the Phase II Value Function	125
A.4	Inner Loop Details	126

Appendix B Appendix to Chapter 2	131
B.1 Technical Details of Section 2.3	131
B.2 Supply-Demand Arrangement	133

List of Tables

1.1 Sample construction.	17
1.2 Sample composition and abatement investment.	17
1.3 Predicting one-year-ahead current- and prior-vintage allowance prices, in 1995 January dollar.	19
1.4 Delivered bituminous coal price regressions, in 1995 January US cent/MMBtu, 1995-2004.	33
1.5 Heat input regressions for the Southern Company's bituminous coal units, in million MMBtu, 1990-2004.	35
1.6 Estimates of θ_A , the allowance internalization parameter, and θ_a , the quadratic allowance cost parameter.	37
1.7 Phase-specific belief parameter estimates, allowance price process estimates, and the capital internalization estimate.	42
1.8 Cost of biased beliefs in Phase I to firms that have had compliance obligations under the Acid Rain Program since 1995.	49
1.9 Aggregate effects of improving beliefs and reducing volatility.	50
1.10 Net benefits from an emission tax and a cap-and-trade program in the Southern Company - Atlantic City Electric Company example.	53
2.1 Pearson rank correlation test results.	63
2.2 Kendall partial rank correlation test results.	64
2.3 Allocations at various solutions in our example.	76
3.1 The stage game where firms produce substitute goods.	88
3.2 The stage game where firms produce complement goods.	92
3.3 Identified set of c at different discount factors, when only action profiles $(0, 0)$ and $(1, 1)$ are observed, with the known signal distribution and no measurement error.	101
3.4 Identification result for $c = 0.238$ at various discount factors and number of signal distributions drawn.	105
3.5 Identification result for $c = 0.20$ at various discount factors.	108

A.1 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.	129

List of Figures

1.1 Predicted 1995-1999 average allowance price by various institutions (Hahn and May, 1994), predicted 1995, 1997, and 1999 allowance prices by survey respondents (Fieldston Company, 1993), and the actual allowance price from the beginning of the allowance market to 2004.	6
1.2 Plants with Phase I units (left) and those with only Phase II units (right) as of 2000. Data source: EPA's Air Market Program Data.	10
1.3 Monthly market price of current- and prior-vintage allowances, in current US dollars, 1993-2004. Data are provided by Denny Ellerman.	14
1.4 Observed one-year-ahead pre-2000 allowance price, and the predictions with <i>v.s.</i> without trends, based on 1993-1999 data, in 1995 January dollar.	20
1.5 Observed one-year-ahead post-2000 allowance price, and the predictions with trends in Phase I but with <i>v.s.</i> without trends in Phase II, based on 1993-2002 data, in 1995 January dollar.	20
1.6 Distribution of sulfur content in non-scrubbed bituminous coal units during compliance, 1995-1999.	21
1.7 Distribution of net purchase of current-vintage allowances by utilities during compliance, 1995-1999.	21
1.8 One-year-ahead predictions of the allowance price using the estimated belief and the estimated stochastic process, based on the Phase I pooled estimates, in 1995 January dollars, 1995-1999. Shaded areas are 95% confidence intervals.	43
1.9 One-year-ahead predictions of the allowance price using the estimated belief and the estimated stochastic process, based on the Phase II pooled estimates, in 1995 January dollars, 2000-2003. Shaded areas are 95% confidence intervals.	43
1.10 One-year-ahead predictions of the allowance price according to beliefs by firms with above- and below-median coal capacity, based on the phase-specific pooled estimates, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.	44

1.11 One-year-ahead predictions of the allowance price according to the beliefs by firms in states that ultimately deregulated the electricity market and those that did not, based on the Phase-specific pooled estimates, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.	45
1.12 One-year-ahead predictions of the allowance price, from annual belief estimates and the adaptive estimate of the allowance price process, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.	46
2.1 Allocations at various solutions in our example.	77
3.1 The feasible and individually rational payoff set, V^* , where firms produce substitute goods.	89
3.2 The feasible and individually rational payoff set, V^* , when firms produce complement goods.	93
3.3 The limit perfect public equilibrium payoff set, $Q(c)$, when firms produce complement goods and have relatively low promotion cost, $c < 0.6$. The graph is generated using $c = 0.55$	95
3.4 The limit perfect public equilibrium payoff set, $Q(c)$, when firms produce complement goods and have relatively low promotion cost, $c \geq 0.6$. The graph is generated using $c = 0.7$	96
A.1 Distribution of volume share of current-, prior-, and next-vintage allowance transactions in any-vintage transactions during the compliance periods, utility-year, 1995-2003.	122
A.2 Distribution of current- and prior-vintage net allowance purchases during the compliance periods, utility-year, 1995-2003.	123
A.3 Distribution of next-vintage net allowance purchases since one year before the compliance periods, utility-year, 1994-2002.	123
A.4 Delivered price of coal and natural gas, 1995-2003, in current U.S. cents per mmBTU. Calculated from the FERC-423 data.	124
A.5 Natural gas price, 1990-2013, in 2010 U.S. dollars per thousand cubic feet. From Davis (2015).	124

Acknowledgments

I would like to thank my advisors Ariel Pakes, Elie Tamer, Joe Aldy, Richard Zeckhauser and Rob Stavins for extensive guidance, helpful comments, and continual support in writing this dissertation. This dissertation would not have been possible without their expertise and generosity.

This research benefitted from discussions with Scott Barrett, Rob Bray, Wolfgang Buchholz, Keith Dougherty, Myrto Kalouptsi, Chris Knittel, Robin Lee, Forest Reinhardt, Jim Roumasset, and participants at the Harvard Environmental Economics Lunch, Seminar in Environmental Economics and Policy, Industrial Organization Lunch, and Industrial Organization Workshop. For Chapter 1, special thanks go to Dallas Burtraw, Denny Ellerman, Kenon Smith, and the industry interviewees for providing institutional details, and to Paul Edmon, Mike Ethier, and the AMPL company for providing computational support. Comments from two anonymous referees greatly improved Chapter 2. I gratefully acknowledge the financial support from the Environment and Natural Resources Program at Harvard Kennedy School and the Harvard Institute for Quantitative Social Science.

Finally, I want to thank my husband, daughter, parents, and the rest of my family for their unconditional support and love. Most of them may not know what exactly I did for the six years of my life; all they wanted to know is that I was happy throughout.

To Ruize and Iris

Introduction

This dissertation comprises three essays on environmental economics and industrial organization. A common theme for the essays is that they draw motivation from the real world and the theory literature to study individual behavior in environment and energy problems.

The first essay explores firm behavior in a new environment by examining private electric utilities' beliefs about sulfur dioxide allowance prices following the implementation of the U.S. Acid Rain Program in 1995. It questions the standard assumption in the empirical industrial organization literature that firms have reasonable expectations about future market conditions. This is partially motivated by strong qualitative evidence I have gathered from interviewing regulators, brokers, and electric utility executives that utilities were struggling to make sense of the allowance market, especially in the early years of this first large-scale cap-and-trade program. I estimate firms' beliefs about the allowance price from 1995 to 2003 using a firm-level dynamic model of allowance trades, coal quality, and emission reduction investment. I find that firms initially underestimate the role of market fundamentals as a driver of allowance prices, but over time their beliefs appear to converge toward the stochastic process of allowance prices. Such beliefs in the first five years of the program cost firms around 10% of their profits. Beliefs also change the relative efficiency of cap-and-trade programs and emission taxes.

In the second essay, co-authored with Richard Zeckhauser, we leverage the asymmetric incentives of individual nations in an international public good provision game to achieve outcomes that are both Pareto efficient and individually rational. A central authority possessing tax and expenditure responsibilities can readily provide an efficient level of a public good.

Providing international public goods, absent a central authority, is challenging due to the substantial asymmetries among nations; small-interest nations have strong incentives to ride cheaply. Our empirical results reveal cheap riding intentions in providing for climate change mitigation, a critical international public good, using individual nations' Intended Nationally Determined Contributions voluntarily pledged for the Paris Climate Change Conference. We find that larger nations made much larger pledges in proportion to both their Gross National Incomes and their historical emissions. To achieve Pareto optimality despite disparate cheap-riding incentives, we propose the Cheap-Riding Efficient equilibrium. That solution takes the Nash equilibrium as a base point, and then applies the principles of either the Nash Bargaining solution or the Lindahl equilibrium to proceed to the Pareto frontier.

Game theoretic models used for structural estimation in industrial organization assume that firms observe each other's behavior perfectly, and yet important problems such as collusion feature imperfect monitoring. For example, a model where firms only observe a noisy signal of all firms' production, in the form of the market price, aptly captures the collusion cases such as the operations of the Organization of the Petroleum Exporting Countries. In the third essay, I introduce repeated games with imperfect public monitoring to structural industrial organization and explore partial identification of model primitives. I use incentive compatibility constraints, in which continuation payoffs come from the limit equilibrium payoff set, to deliver bounds on the structural parameters. By simulating a repeated duopoly competition game, I illustrate the identification procedure, and explore the implications of unknown public signal distributions and mis-measured firms' actions.

Chapter 1

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

1.1 Introduction

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

George Bernard Shaw

Everybody's Political What's What, 1945

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

Peter Passell for *The New York Times*

Paying to pollute: a free market solution that's yet to be tested, 1996

New economic environments abound. Technology creates new markets, regulation defines new boundaries, and industry dynamics shape new competitive landscapes. As firms start out

in an environment they lack experience in, they may not form reasonable expectations about future market conditions. While adaptability has been much advocated for in the business literature, how firms actually behave in a new environment is little understood in economics.

This paper investigates the beliefs of private electric utilities about the sulfur dioxide (SO₂) allowance price under the U.S. Acid Rain Program from 1995 to 2003. As the first large-scale cap-and-trade program, the Acid Rain Program thrust the electric utility industry into an unprecedented market-based environmental regulation. In planning to comply with the Acid Rain Program, utilities need to form beliefs about the future SO₂ allowance price. Given the novelty of this environment, do utilities hold beliefs that are at least as good as what they could have inferred from available data? Do beliefs improve over time as utilities gain experience in the allowance market? What do the beliefs imply for firm profits and social welfare? Answers to those questions are important for choosing, designing and evaluating cap-and-trade programs, now the predominant policy for addressing climate change, as well as various forms of other quantity-based instruments¹. The lessons are particularly relevant for countries, such as China, that expose traditional industries to market-based environmental regulation for the first time.

I use dynamic structural estimation to back out those beliefs about the future allowance price from firm-level allowance trades, coal quality choices, and emission reduction investment. The intuition is simple: the higher a utility expects tomorrow's allowance price to be, the more aggressively it saves allowances, by buying more (or selling fewer) allowances, reducing the sulfur content of coal, and making more emission reduction investment. In estimating the beliefs, I make three contributions. First, I build the first empirical dynamic model of firm behavior in cap-and-trade programs. Dynamics is necessary for understanding firm behavior in cap-and-trade programs, because all compliance decisions have dynamic implications via the saving of allowances.² Second, I provide the first empirical application of two novel dynamic

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

²Fowle et al. (2016) develops a dynamic model of the cement industry under various environmental policies,

programming acceleration methods, the Relative Value Function Iterations method (Bray, 2017b) and the Endogenous Value Function Iterations method (Bray, 2017a). They vastly reduce the computation burden. Third, this is the first time that extensive SO₂ allowance trading data have been matched to electric operations data for academic research³

Existing evidence suggests that the initial beliefs about the future allowance price were quite unreasonable. Figure 1.1 shows that by 1995, all major predictions of the 1995-1999 average allowance price (summarized by Hahn and May (1994)) exceeded the actual prices, most of which by a large margin. The expected 1995, 1997, and 1999 allowance prices surveyed by Fieldston Company (1993) were no better. It is not convincing that those disparities, being consistent and vast, are merely due to the difference between expectation and realization. Neither do they arise simply from regulatory capture; the rules had been set years ago. A plausible culprit, however, is that utilities may not have made good use of available information; Montero and Ellerman (1998) attribute 60% of the over-prediction to the failure to incorporate the decline in the low-sulfur coal price that had been going on since 1980. The inefficient use of information is likely because of utilities' inexperience with the unprecedented allowance market. My private communication with utility executives, regulators, and allowance brokers confirms the utilities' initial difficulty with the allowance market; utilities were wondering about "how is there going to be a market," "what are we going to do," and many small firms appeared "not up for the idea of market and trading."

To investigate whether the beliefs have improved since 1995 as more information became available,⁴ I first estimate the beliefs from a dynamic model of cap-and-trade compliance. A coal-dependent private electric utility chooses the net purchase of allowances, the sulfur content of coal, and capital investment to allow fuel-switching, based on electricity demand, allowance

including a cap-and-trade program; they focus on entry and exit decisions and assume away allowance trading and allowance price uncertainty. Other empirical models of firm behavior in cap-and-trade programs are static, including Fowle (2010), Chan et al. (2017), and Cicala (2015).

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

⁴That said, when the initial beliefs in Figure 1.1 were formed, there had been almost two years of allowance price observations by 1995, because the allowance market had started operating in 1993.




Figure 1.1: Predicted 1995-1999 average allowance price by various institutions (Hahn and May, 1994), predicted 1995, 1997, and 1999 allowance prices by survey respondents (Fieldston Company, 1993), and the actual allowance price from the beginning of the allowance market to 2004.

price, and allowance stock, subject to the requirement that it have enough allowances to cover the SO_2 emissions. The model takes into account the allowance transaction cost, the distortion of incentives by cost-of-service regulation, the vintage structure of allowances, and the two-phase compliance structure of the Acid Rain Program. I estimate this three-dimensional continuous-state dynamic model, where the allowance price beliefs are the perceived law of motion for the allowance price state variable, using a nested-fixed-point-style algorithm with a maximum simulated likelihood estimator. The allowance price belief parameters are identified from inter- and intra-firm variations in compliance behavior induced by variations in the expected marginal dynamic value of allowances, which in turn are induced by exogenous variations in the allowance price and electricity demand.

Having estimated the beliefs, I then compare them with two alternative beliefs widely used in structural work in industrial organization: full-information rational expectations, and adaptive-learning beliefs. The former comparison asks whether utilities know the “overall” stochastic process of allowance prices as does an econometrician who estimates the process

using a full series of data. Some argue that full-information rational expectations require too much statistical knowledge from economic agents, and propose to put economic agents and econometricians on comparable footing (Hansen, 2007). Thus, the latter comparison is based on the “adaptive” stochastic process of allowance prices, which is estimated iteratively, as time rolls forward, using those data available only up to each rolling year.⁵

I find that firms underestimate the role of market fundamentals as a driver of the allowance price; they pay too much attention to what is happening inside the allowance market (the historical allowance prices) and too little to the conditions in related markets (such as electricity and coal markets) that drive the allowance price. Consequently, they predict a flatter allowance price trajectory than full-information rational expectations; the latter predicts the rises and falls in the actual allowance price trajectory much better.⁶ Smaller firms, and firms facing less competitive pressure, exhibit flatter beliefs. Over the years, those beliefs appear to converge toward the adaptive stochastic process of the allowance price. This trend is consistent with the shift in the management practice in electric utilities during that period,⁷ where compliance decisions are made less by engineers but more by people with experience in markets and trading; utilities seem to have gradually grasped the philosophy of market-based environmental regulation that compliance is about profit maximization, not cost minimization (Reinhardt, 2000).⁸

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

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

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

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

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

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

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

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

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

The rest of this paper proceeds as follows. Section 1.2 describes the institutional details of the Acid Rain Program, compliance options, the allowance price trajectory, and cost-of-service regulation under which private electric utilities operate. Section 1.3 explains data sources, sample construction, and suggestive evidence of biased beliefs about the future allowance price. Section 1.4 presents the model, followed by the estimation strategy and results in

Section [1.5](#). Section [1.6](#) presents counterfactual experiment results on the consequences of biased beliefs, and compares cap-and-trade programs with emissions taxes from a dynamic perspective. Section [2.6](#) concludes.

1.2 Institutional Background

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

Phase I (1995-1999) of the Acid Rain Program covers 263 largest, dirtiest units, almost all coal-fired, and Phase II (2000-) covers all fossil-fuel-based units exceeding 25MW generating capacity. Around 180 units from Phase II voluntarily opted into Phase I. I do not differentiate between those voluntary units and the original 263 units, and call them “Phase I units”. I use “Phase II units” to refer to the rest of the units, those that have compliance requirement only since 2000; note, however, that Phase I units have compliance obligations in Phase II as well. Figures [1.2](#) shows the geographic locations of plants with Phase I units and those with only Phase II units.

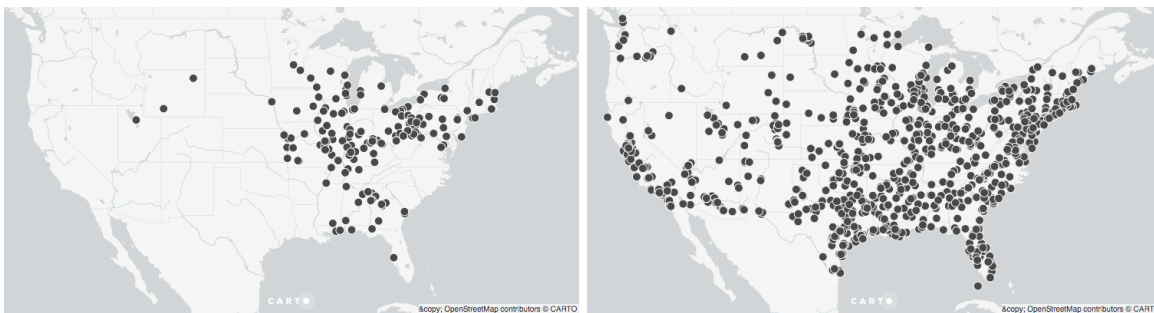


Figure 1.2: *Plants with Phase I units (left) and those with only Phase II units (right) as of 2000. Data source: EPA’s Air Market Program Data.*

Each SO_2 allowance represents one ton of SO_2 . The number of allowances allocated to

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

Compliance options. Almost all SO₂ emissions from electric utilities come from units that burn bituminous coal. Utilities can reduce SO₂ emissions by using lower-sulfur bituminous coal (the intensive margin), switching to the ultra-low-sulfur, sub-bituminous coal (the extensive margin), or retrofitting the unit(s) with a flue-gas desulfurization equipment, or a “scrubber”.

The sulfur content of bituminous coal ranges from 1.5 to 5 lb/MMBtu. It is mined widely around the Illinois Basin and the Appalachia areas, and is considered high-quality coal because of its high heat content and low moisture content. Bituminous coal has the most developed spot market, and many firms choose lower-sulfur, Central Appalachian bituminous coal to reduce SO₂ emissions (Ellerman et al., 2000). Little capital investment is necessary.

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

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

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

The flexibility of the cap-and-trade program allows firms to buy allowances without cutting back emissions. A firm with excess allowances can sell them for revenue. The EPA does not impose any restriction on who can trade, when, and how. Allowances are labeled with vintages, and in a given compliance year, only the allowances of the current and the prior vintages can be used for compliance (that is, the EPA deducts only the current- and the prior-vintage allowances). This precludes borrowing from the future. For example, if a unit emitted 4000

tons of SO₂ in 1997, the unit account needs to hold at least 4000 allowances of vintage 1997 or prior by the emission deduction deadline. The EPA also holds a small allowance auction in every March to auction off 2.8% of the annually allocated allowances and returns revenue to compliance units.

The allowance price. Trade journals and brokerage firms publish monthly allowance prices. Figure 1.3 shows the market price of current- and prior-vintage allowances reported by three major sources from 1993-2005. The allowance price began at almost \$200 before Phase I started, but it quickly plunged to as low as \$60 during the first year of the program. It then oscillated for the next decade. I focus on the allowance prices before the end of 2003; the post-2004 prices varied wildly due to unanticipated regulatory uncertainty, which is beyond the scope of this paper.¹⁰

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

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

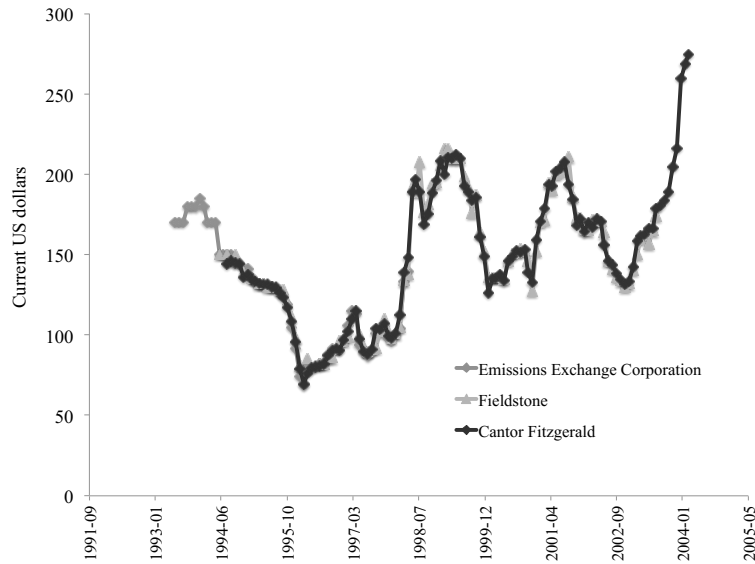


Figure 1.3: Monthly market price of current- and prior-vintage allowances, in current US dollars, 1993-2004. Data are provided by Denny Ellerman.

The Acid Rain Program creates new costs and revenues for private utilities. [Lile and Burtraw \(1998\)](#) describe in detail the regulatory treatment of those costs and revenues by each state's public utilities commission. Approved scrubbing or fuel-switching investment would be added to the rate base, earning profits for the utility at the pre-specified rate of return. Prudently-incurred expenditure on lower-sulfur coal, of acquiring allowances, and the revenue from selling allowances would be passed on to ratepayers.

1.3 Data

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

The main sources of electricity production data are multiple surveys administered by the

Energy Information Administration (EIA) and the Federal Energy Regulatory Commission (FERC). Data on plant divestiture are from [Cicala \(2015\)](#), which are also based on EIA data. The Acid Rain Program compliance data are from the Air Markets Program Data (AMPD) system at the Environmental Protection Agency. Monthly current- and prior-vintage allowance prices are from Denny Ellerman, who had collected it from trade journals and brokerage firms. Monthly future-vintage prices are from the online archive of Cantor Fitzgerald / BGC Environmental Brokerage Services, the biggest allowance broker. Appendix [A.1](#) describes data collection and processing details.

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

Sample construction. The period of analysis is from 1995 to 2003. I focus on private electric utilities that operate un-scrubbed bituminous-coal units subject to at least one phase of the Acid Rain Program and rely little on other SO₂-emitting fuels. The sample includes 42 firms, accounting for over half of the allowance trading volumes during the period of analysis.

Table [1.1](#) summarizes the sample construction process. First, from all electric utilities, I select those with coal units subject to at least one phase of the Acid Rain Program. This is because otherwise the utilities would have trivial compliance burden; coal is responsible for almost all SO₂ emissions from electric utilities^{[11](#)} Those included utilities, however, can have gas units in addition to coal units. Indeed, gas operations interfered little with coal operations during the period of analysis; gas emits zero SO₂, and as will be described later in this section, gas prices were too high to affect the electricity generation from coal.

Second, from all electric utilities with coal units subject to at least one phase of the Acid Rain Program, I select those that are privately-owned, which are under cost-of-service

¹¹In 1985, about 70% of the U.S. SO₂ emissions came from electric utilities, of which 96% were emitted by coal-fired generation units ([U.S. Environmental Protection Agency, 1994](#)).

regulation. Private electric utilities were responsible for the vast majority of electricity generation prior to the electricity restructuring. Studying the beliefs after the electricity restructuring is for future research, which requires a competitive, instead of single-agent, model.

Third, I exclude those utilities with non-coal-non-gas units, such as oil and petroleum coke, that constitute more than 10% of the utility capacity and SO₂ emissions. This is because their compliance problem would then involve choices of non-coal-non-gas fuels, complicating the strategy space. Finally, I exclude those utilities that do not operate un-scrubbed bituminous-coal units. This is because utilities with all coal units scrubbed or burning sub-bituminous coal have trivial compliance burden.¹²

Table 1.2 reports the sample composition and the emission abatement investment. In the “composition of coal units” panel, the number of private utilities with a given type of unit composition is larger in 1995 than 2000, because of the electricity restructuring that happened between 1998 and 2002. Of those with compliance obligations since 1995, 13 chose to scrub at least one unit, and 11 chose to switch at least one unit to sub-bituminous coal. Of those with compliance obligations since 2000, none chose scrubbing, and five chose switching. Those that neither scrub or switch fuel reduced the sulfur content of bituminous coal or bought allowances to comply.

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

Table 1.1: *Sample construction.*

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

Table 1.2: *Sample composition and abatement investment.*

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

Evolution of allowance prices. Table [1.3](#) reports various estimates of the stochastic process of the current- and prior-vintage allowance price. Those estimates would typically serve as inputs to the estimation of a dynamic model. The point of this paper, however, is to compare those estimates with the effective beliefs that firms hold. In other words, I do not take the former as inputs to dynamic estimation; instead I back out the latter from data.

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

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

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.

Figures [1.4](#) and [1.5](#) plot the observed one-year-ahead allowance price against the predictions from Columns (1) and (2), and from Columns (3) and (4). Consistent with Table [1.3](#), the figures show that trends appear important for Phase I allowance price prediction, but not for Phase II. This informs the specification of belief in the dynamic model. Appendix [A.2](#) provides additional data patterns that inform the model in Section [1.4](#).

Suggestive evidence of biased beliefs. Figures [1.6](#) and [1.7](#) plot the distributions of sulfur contents and allowance trades against the allowance price trajectory in Phase I. The allowance price started a sharp increase in 1996. As shown in Figure [1.4](#), this increase is predictable based on available data. If firms had reasonable expectations of this increase, they would cut back on SO₂ emissions and buy allowances. By doing so, they would benefit both statically, because the allowance price was low, and dynamically, because the allowance price was going up. Figures [1.6](#) and [1.7](#) show otherwise.

Table 1.3: *Predicting one-year-ahead current- and prior-vintage allowance prices, in 1995 January dollar.*

	(1) 93-99	(2) 93-99	(3) 93-02	(4) 93-02
Current price	0.179* (0.0869)	0.871*** (0.0796)	0.569*** (0.0635)	0.490*** (0.0690)
Current price \times phase I year		-0.463*** (0.0518)	-0.258*** (0.0231)	
Phase I year		81.16*** (9.313)	43.21*** (3.760)	
Current price \times year				-0.204*** (0.0210)
Year				33.53*** (3.408)
Constant	101.6*** (14.01)	-15.39 (12.50)	34.03** (10.48)	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 parenthesis. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The allowance price is deflated using the Urban Consumer Price Index.




Figure 1.4: Observed one-year-ahead pre-2000 allowance price, and the predictions with v.s. without trends, based on 1993-1999 data, in 1995 January dollar.




Figure 1.5: Observed one-year-ahead post-2000 allowance price, and the predictions with trends in Phase I but with v.s. without trends in Phase II, based on 1993-2002 data, in 1995 January dollar.




Figure 1.6: *Distribution of sulfur content in non-scrubbed bituminous coal units during compliance, 1995-1999.*

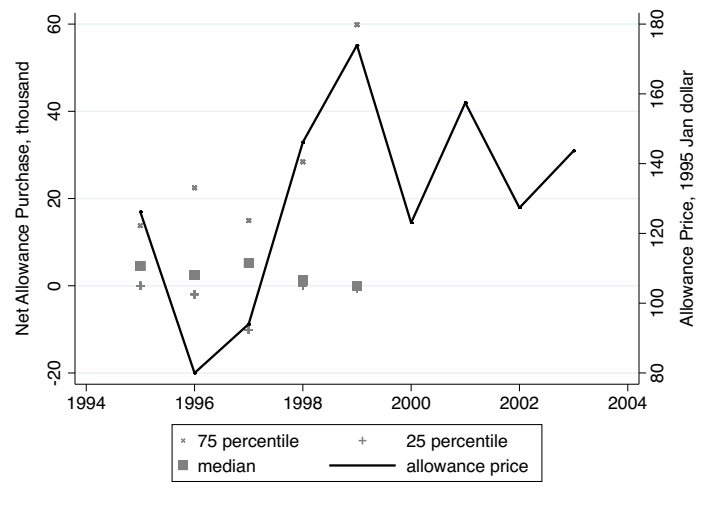


Figure 1.7: *Distribution of net purchase of current-vintage allowances by utilities during compliance, 1995-1999.*

1.4 Model

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

1.4.1 Allowance Price Belief

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

$$P_s = b_1 + b_2 \times \min\{s - 1, 2000\} + (b_3 + b_4 \times \min\{s - 1, 2000\})P_{s-1} + \epsilon_s, \quad (1.1)$$

where ϵ_s is an i.i.d. normal error with mean zero and standard deviation b_5 . Again, the values for the parameters $(b_1, b_2, b_3, b_4, b_5)$, which capture the firm's belief about the allowance price, will come from structural dynamic estimation in Section [1.5](#), rather than from the estimation of the stochastic process in Section [1.3](#).

Thus, the firm forecasts future allowance prices based on historical allowance prices via the slope coefficients, and the conditions of related markets (*e.g.*, coal and electricity markets) via the intercepts. Importantly, the firm believes that the allowance price process is not stationary

until Phase II; Equation (1.1) has time trends in both the slope and the intercept only if the forecast year $s \leq 2000$. I impose stationarity on the Phase II forecasts to capture the perception that the allowance price formation process will settle down by Phase II. This also facilitates the dynamic programming of the infinite-horizon Phase II problem, as it requires stationary state transitions.

The parameters of interest are $(b_1, b_2, b_3, b_4, b_5)$. In Section 1.5, I will present both the pooled estimates, using all data between 1995 and 1999, and the year-specific estimates, using the data in each year separately. The former corresponds to rational expectations and the latter adaptive-learning beliefs. For the latter, I adopt the standard assumption for the adaptive estimation that the firm is myopic about the possible future belief changes; it does not internalize those possible changes when it makes its decisions today.

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

$$P_s = b'_1 + b'_3 P_{s-1} + \epsilon_s \quad (1.2)$$

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

1.4.2 The Phase II Problem

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

$$\pi_i(a_{it}, x_{it}; P_t, H_{it}) = \phi_A[A(a_{it}; P_t)] + \phi_C[C_i(x_{it}; H_{it})], \quad (1.3)$$

where the internalization functions translate the monetary value of those expenditures into what the firm deems payoff. As an example, a firm that is not subject to cost-of-service regulation and produces output D would have $\phi_A(A) = -A$ and $\phi_C(C) = -C + P(D)D$, where P is the inverse demand function. That firm would internalize 100% of the allowance and coal expenditures is because the revenue is independent of them. However, the firms in this paper are subject to cost-of-service regulation, and their revenue depends on their cost. As described in Section [1.2](#), the public utilities commission passes the costs that those firms prudently incurs to ratepayers. As a result, those firms may not internalize every dollar spent on allowances and coal.

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

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

allowance stock is deterministic and endogenous:

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

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

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

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

where F_P is the allowance price belief of Equation (1.2), and F_{H_i} is the transition of heat input.

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

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

$$\frac{\partial \phi_C(C^*)}{\partial C} \frac{\partial C_i(x_i^*; H_i)}{\partial x_i} / H_i - \beta E \frac{\partial V((W'_i)^*, P', H'_i)}{\partial W'_i} - \mu^* = 0, \quad (1.4)$$

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

$$\frac{\partial \phi_A(A^*)}{\partial A} \frac{\partial A(a_i^*; P)}{\partial a_i} + \beta E \frac{\partial V((W'_i)^*, P', H'_i)}{\partial W'_i} + \mu^* = 0, \quad (1.5)$$

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

The first-order conditions explain optimal behavior intuitively.¹³ Equation (1.4) 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

¹³ Appendix A.3 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.

lower sulfur content both avoids drawing from the allowance stock and relaxes the compliance constraint. Equation (1.5) 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 (1.4) and (1.5) together imply that, at the optimum, the marginal cost of lowering the sulfur content in coal equals the marginal cost of acquiring additional allowances. This is regardless of whether the compliance constraint binds or not. Indeed, sulfur content reductions and allowances are perfect substitutes. In Section 1.5, I use this equality to estimate the payoff parameters without solving the dynamic programming problem.

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

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

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

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

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 [A.2](#) 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.

1.4.3 Fuel-Switching Investment

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

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

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

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

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:

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

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.

1.4.4 Finite-Horizon Phase I

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

$$\begin{aligned}
V_i^{1999}(W_{i,1999}, P_{1999}, H_{i,1999}) = & \max_{\substack{x_i \in X \\ a_i \geq x_i H_{i,1999} - W_{i,1999} - alloc_{i,1999} \\ a_i^{nv} \geq -alloc_{i,2000}}} \{ \phi_A[A(a_i, a_i^{nv}; P_{1999})] \\
& + \phi_C[C_i(x_i; H_{i,1999})] + \log[\sum_{j=1}^{|J_i|} \exp(-\phi_K(k_j c^k)] \\
& + \beta V_i^j(W_{i,1999} + alloc_{i,1999} + a_i + a_i^{nv} + x_i H_{i,1999}, P_{2000}, \tilde{H}_{i,2000}) \\
& \times dF_P(P_{2000}|P_{1999}) dF_{\tilde{H}_i}(\tilde{H}_{i,2000}|g_i(H_{i,1999})) \}, \tag{1.6}
\end{aligned}$$

where F_P is the Phase I allowance price transition in Equation (1.1) with $s = 2000$, and $F_{\tilde{H}_i}$ is the transition of the heat input to Phase II units; the function g_i converts the heat input to Phase I units to that to Phase II ones.¹⁵

The Bellman's Equation for the year $t \in \{1995, \dots, 1998\}$ 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 (1.1) with $s = t + 1$, and F_{H_i} is the transition of heat input to Phase I units.

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

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

1.4.5 Identification

The parameters to be estimated are:

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

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

The allowance expenditure parameter and the internalization parameters in ϕ_A, ϕ_C are identified based on the equality between the marginal cost of lower sulfur content and the marginal cost of allowances from Equations (1.4) and (1.5). The allowance expenditure parameter concerns how the marginal cost rises with the transaction volume.¹⁶ As will be discussed in Section 1.5, this equality is subject to the measurement error in the reported allowance trades. I correct for the measurement error to obtain unbiased estimates of those parameters.

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

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

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

1.5 Estimation

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

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

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

simulation numerically integrates out the measurement errors in the allowance stock state variable induced by the measurement errors in the net allowance purchase.

1.5.1 Coal Expenditure Parameters

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

$$p_{jkt}^{bit} = \gamma_0^{bit} + \gamma_1^{bit} x_{jkt} + \gamma_2^{bit} x_{jkt}^2 + \gamma_Z^{bit} Z_{jkt} + dummies + \epsilon_{jkt}^{bit} \quad (1.7)$$

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

The parameters of interest are γ_1^{bit} and γ_2^{bit} . They measure how the delivered coal price changes with the sulfur content of coal. [Table 1.4](#) reports the estimates of γ_1^{bit} and γ_2^{bit} at various specifications. I use the estimates from Column (4) for the rest of the analysis.

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

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

where x is the sulfur content, and H is the heat input. I omit from the coal expenditure the non-sulfur components of the coal price function [\(1.7\)](#); they do not affect the firm's choices in Phase I.

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

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

where $\bar{\alpha}_i k_j$ is the heat input to the sub-bituminous units switched to, in which $\bar{\alpha}_i$ is the heat input rate of those units. Unlike in the Phase I problem, the coal expenditure in the Phase II problem includes the non-sulfur components in the bituminous coal price, summarized by $\bar{\gamma}_i^{bit}$. The non-sulfur price components affect the value of switching capacity k_j , which in turn

Table 1.4: *Delivered bituminous coal price regressions, in 1995 January US cent/MMBtu, 1995-2004.*

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

affects the terminal value functions for Phase I. I construct the non-sulfur bituminous coal price $\bar{\gamma}_i^{bit}$ by taking the heat-input-weighted average of $(p_{jkt}^{bit} - \gamma_1^{bit}x_{jkt} - \gamma_2^{bit}x_{jkt}^2)$ over the bituminous coal shipments received by plants operated by firm i during Phase II.

The coal expenditure in the Phase II problem also includes sub-bituminous coal expenditure. Let $\bar{\gamma}_i^{sub}$ denote the sub-bituminous coal price. For firm i that receives sub-bituminous coal shipments in Phase II, I construct $\bar{\gamma}_i^{sub}$ by taking the heat-input-weighted average of p_{jkt}^{sub} over the sub-bituminous coal shipments received by plants operated by firm i during Phase II. For firm i that does not receive sub-bituminous coal shipments in Phase II, I construct the counterfactual $\bar{\gamma}_i^{sub}$ as follows. First, I conduct a sub-bituminous coal price regression in the style of Equation (1.7), using data from plants that receive sub-bituminous coal shipments. Second, I construct the predictors for plants that do not receive such shipments. The sulfur content would be the industry-average sulfur contents used by sub-bituminous coal units; the other predictors would be the average values of the sub-bituminous coal produced by the coal county that yields the lowest predicted coal price for firm i .

1.5.2 Heat Input Transitions

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

$$H_{i,t} = \gamma_{0,i}^H + \gamma_{1,i}^H H_{i,t-1} + \epsilon_{i,t}^H, \quad (1.8)$$

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

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

I estimate the heat input transitions separately for the Phase I units only, and the Phase I and Phase II units taken together. Table 1.5 reports the estimates of $\gamma_{0,i}^H$, $\gamma_{1,i}^H$, and σ_i^H for the Southern Company as an example.

Table 1.5: *Heat input regressions for the Southern Company's bituminous coal units, in million MMBtu, 1990-2004.*

	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 parenthesis. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

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

As shown in Section [1.4](#), the first-order conditions of the dynamic problem indicate that the marginal cost of lowering the sulfur content equals the marginal cost of acquiring additional current-vintage allowances, as long as the sulfur content remains interior. I use this equality to estimate the allowance expenditure parameters and the internalization parameters for allowance and coal expenditures.

The allowance expenditure is:

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

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

The internalization functions for the allowance and coal expenditures are:

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

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

$$\theta_C(-\gamma_1^{bit} + 2\gamma_2^{bit} x_{i,t}) = \theta_A(P_t + 2\theta_0 a_{i,t}). \quad (1.9)$$

Since we can only identify the ratio between the two internalization parameters, I normalize $\theta_C = 1$. Let $\theta_a = \theta_0 \theta_A$. Equation (1.9) thus becomes:

$$-\gamma_1^{bit} + 2\gamma_2^{bit} x_{i,t} = \theta_A P_t + 2\theta_a a_{i,t}. \quad (1.10)$$

As described in Section 1.2, the allowance trades contain measurement errors. Private communication with the EPA suggests that around 25% of the allowance trades may not have been reported. I use this information to construct a model for the measurement error in the allowance trades. First, I multiply the total allowance trading volume with 25% to approximate the total number of allowance trades that are not reported. Second, I normalize this number by the total allowance allocation, arriving at the per-allocation *expected* measurement error, $e = 0.225$. The current-vintage measurement error model is:

$$a_{i,t} = a_{i,t}^* + alloc_i \eta_{i,t}, \quad (1.11)$$

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

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

Table 1.6: Estimates of θ_A , the allowance internalization parameter, and θ_a , the quadratic allowance cost parameter.

	OLS	Corrected
Internalization of allowance expenditure (θ_A)	0.636*** (0.00986)	0.636*** (0.00977)
Quadratic allowance cost (θ_a)	0.0000562 (0.000645)	0.0000572* (0.0000288)
N	150	
Adj. R-sq	0.968	

Standard errors are in parenthesis. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

After we replace a with a^* in Equation (1.10) and substituting in Equation (1.11), Equation (1.10) becomes:

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

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

$$\tilde{Y}_{it} = \theta_A \tilde{P}_t + 2\theta_a \tilde{a}_{i,t} - 2\theta_a \eta_{i,t}, \quad (1.12)$$

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:

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

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

Table 1.6 reports the estimates of θ_a and θ_A and their corrections. To ensure that the sulfur content is interior, as required by the equality between the marginal costs, I restrict the sample to those observations with \tilde{Y} within the 5% and 95% percentiles.

1.5.4 Inner Loop

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

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

Estimation of beliefs held during Phase I requires dynamic programming of both the Phase I (including the investment problem in 1999) and Phase II problems. Given a parameter vector, for each firm, I first use the Relative Value Function Iterations method (Bray, 2017b) to quickly solve for the Phase II value functions given each fuel-switching investment option. The 1999 investment problem requires as inputs the *levels* of the Phase II value functions, which can be backed out from the relative value functions. Indeed, the Relative Value Function Iterations method focuses on the value function shape over all exogenous and endogenous states, and therefore the resulting relative value function is the full value function shifted by a constant. Having backed out the value functions associated with each investment option, I use backward induction to solve for the value functions specific to each year in Phase I.¹⁹ Appendix A.4 describes further details on the inner loop.

1.5.5 Outer Loop

The outer loop looks for the parameter values that maximize the simulated likelihood. The likelihood is induced by the i.i.d. distribution of measurement errors in the current- and next-

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

vintage allowance trades and the sulfur content; firms are not required to report transactions of allowances not used for contemporaneous compliance, and the precision of the self-reported sulfur content is not satisfactory.²⁰ The measurement error in the allowance trades induces measurement error in the allowance stock state variable, and the individual likelihood needs to be integrated over possible true allowance stocks. Because the measurement error in the allowance stock affects behavior nonlinearly via the dynamic problem, the likelihood is simulated.

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

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

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 $W_{i,t+1}^* = W_{i,t}^* + alloc_i + a_{i,t}^* + b_{i,t}^* - deduct_{i,t}$;

²⁰For 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 (1.12) in the first stage, because this measurement error is for the dependent variable. Yet the standard errors of the coefficients are under-estimated.

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:

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

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 ²¹. Those functions depend on θ in a highly nonlinear way via dynamic programming. The integration is over the possible paths of true allowance states. I compute the integration with simulation.

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

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

To estimate θ , I use the BHHH algorithm (Berndt et al., 1974) with numerical gradients.

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

Grid search informs the choice of the initial value for θ . The standard errors of the estimates take into account the errors that come from the first-stage parameter estimates. The standard errors coming from the second-stage structural estimation are calculated using standard formula with numerical gradients.

Results. Table [1.7](#) reports the estimation results from the outer loop, by pooling the observations within the same phase of the Acid Rain Program. As a comparison, I present the estimated parameters of the stochastic process of the allowance price along with the estimated belief parameters. Recall that in the allowance price belief specification in Section [1.4](#), the expected allowance price next year is the sum of an intercept, interpreted as the contribution from market fundamentals, and a term dependent on the current allowance price, via a slope coefficient. Furthermore, to capture the time-varying market conditions, both the intercept and the slope have linear time trends. The estimation results suggest that in Phase I, firms believe in a “flatter” allowance price process than it actually is; they under-estimate the role of market fundamentals, compared to the past allowance prices, as a driver of the allowance price.

Figure [1.8](#) plots the one-year-ahead allowance price predictions in Phase I using the estimates of the belief and the stochastic process from pooling observations in Phase I. While the latter tracks the observed price trajectory reasonably well, the former predicts allowance prices that are too high in early years of Phase I and too low in later years. Figure [1.9](#) plots the estimates from pooling observations in Phase II. Unlike the early years of Phase I, the confidence bands in Phase II now overlap more, and the difference in the mean is much smaller.

Aside from the belief estimates, Table [1.7](#) also reports a capital internalization estimate of 0.847. Thus, firms consider the fuel-switching investment cheaper than it really is. This is consistent with the Averch-Johnson effect ([Averch and Johnson, 1962](#)), where regulated 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](#)).

Table 1.7: *Phase-specific belief parameter estimates, allowance price process estimates, and the capital internalization estimate.*

		<u>Phase I</u>		<u>Phase II</u>	
		Belief	Process	Belief	Process
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 parenthesis.




Figure 1.8: One-year-ahead predictions of the allowance price using the estimated belief and the estimated stochastic process, based on the Phase I pooled estimates, in 1995 January dollars, 1995-1999. Shaded areas are 95% confidence intervals.

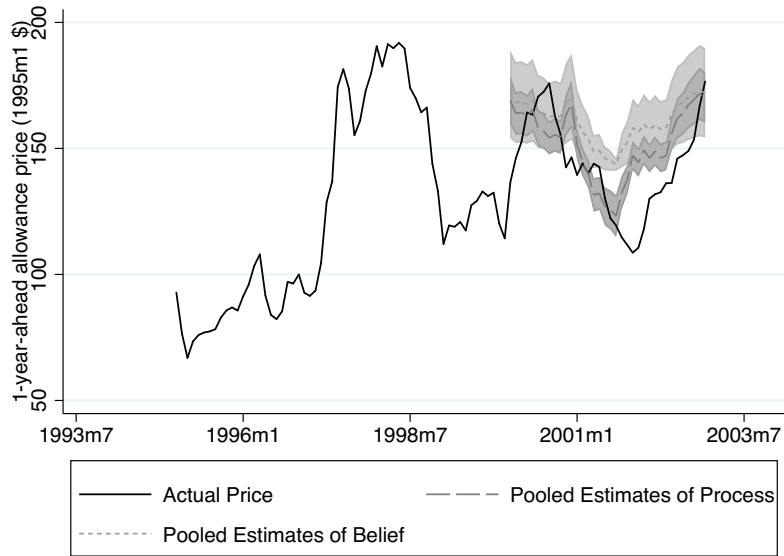


Figure 1.9: One-year-ahead predictions of the allowance price using the estimated belief and the estimated stochastic process, based on the Phase II pooled estimates, in 1995 January dollars, 2000-2003. Shaded areas are 95% confidence intervals.

I now explore firm heterogeneity in the beliefs. Figure [1.10](#) compares the allowance price predictions by firms with above- and below-median coal capacity, based on the phase-specific pooled estimates; and Figure [1.11](#) compares the predictions by firms in states that ultimately deregulated the wholesale electricity market, and in states that remain under cost-of-service regulation. They show that larger firms and firms facing competitive pressure from future deregulation hold beliefs that tend to predict allowance prices better, especially in Phase I. Intuitively, larger firms likely had more resources to devote to studying the allowance market, and firms in states preparing for electricity restructuring either were already market-savvy to begin with, or had started to strengthen their competitiveness.




Figure 1.10: *One-year-ahead predictions of the allowance price according to beliefs by firms with above- and below-median coal capacity, based on the phase-specific pooled estimates, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.*

To investigate the evolution of beliefs, Figure [1.12](#) plots the adaptive, year-by-year estimates; I estimate the stochastic process of the allowance price by adaptively using data up to the point of prediction, and the beliefs by year by using cross-sectional data from each year separately. The overall pattern is similar to that in the pooled estimates in Figures [1.8](#) and [1.9](#); the belief is flatter than the price process as firms do not appreciate market trends sufficiently.




Figure 1.11: *One-year-ahead predictions of the allowance price according to the beliefs by firms in states that ultimately deregulated the electricity market and those that did not, based on the Phase-specific pooled estimates, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.*

Furthermore, the annually estimated belief starts wildly apart from the adaptively estimated stochastic process, but the gap appears to get smaller towards the end of Phase I and into Phase II.

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




Figure 1.12: *One-year-ahead predictions of the allowance price, from annual belief estimates and the adaptive estimate of the allowance price process, in 1995 January dollars, 1995-2003. Shaded areas are 95% confidence intervals.*

1.6 Implications of Biased Beliefs

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

Second, I assess the implications of improving the beliefs of “bias-prone” firms, and reducing the allowance price volatility, for aggregate environmental and economic outcomes. The estimation results in Section [1.5](#) suggest that smaller firms and firms with less competitive pressure tend to have more biased beliefs. How would the aggregate SO_2 emissions and coal expenditures change if those firms had the same beliefs as bigger firms and firms with more competitive pressure? This quantifies the effects of policies that enable bias-prone firms to have a better understanding of the allowance market; examples include publishing market information relevant to the allowance market in a timely and transparent fashion,

holding workshops to facilitate communications among utilities, brokers, and regulators, and introducing competition to the electricity market. Additionally, I simulate the effects of an allowance price collar (*i.e.* a price floor combined with a price ceiling) that constrains the allowance price between \$50 and \$200. Price floors, ceilings, and collars reduce the range of possible allowance prices. They are common policy tools to reduce allowance price volatility; although not used in the Acid Rain Program, they are present in many cap-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?²²

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

1.6.1 Implications for Firm Payoffs

To quantify the effect of biased beliefs on the dynamic payoffs of individual firms, I first simulate each firm's counterfactual choices of allowance trades and sulfur content iteratively from 1995 to 1998, with belief parameters taking values from the "stochastic process" column of Phase I estimation in Table 1.7. I then calculate the firm's discounted sum of static payoffs from 1995 and 1998 according to Equation (1.3). 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

²²In those counterfactual simulations, I improve the beliefs held by the bias-prone, but not all, firms in my sample, and use a price collar that is non-binding during Phase I. The purpose is to mitigate the feedback effect of simulated behavior on the allowance price. Indeed, the dynamic model in Section 1.4 is a single-agent model that treats the allowance price as an exogenous state variable; it models how individual firms respond to allowance prices but not the other way round. Counterfactual change in the behavior of bias-prone firms is unlikely to significantly alter the allowance price trajectory; Table 1.1 shows that even all firms in my sample merely constitute around half of the allowance trade volume from 1993 to 2003.

counterfactual state of allowance stock, in the year-1999 value function of Equation (1.6) (before the fuel-switching investment cost shock realizes). Having obtained the counterfactual dynamic payoff of each firm, I repeat the steps above with the actual behavior to obtain the actual dynamic payoff of each firm. The difference, normalized by the number of years in Phase I, is the annual loss in the firm's dynamic payoff due to the biased allowance price belief it held during Phase I. I focus on the beliefs in Phase I, because during that period the estimated belief and the stochastic process are substantially different (Figure 1.8).

Table 1.8 reports the cost of biased beliefs to firms that have had compliance obligations under the Acid Rain Program since 1995. The annual forgone dynamic payoffs due to biased beliefs in Phase I range from 0.31 to 20.15 million dollars. To put those numbers in context, Table 1.8 also lists the total revenue of each utility in 1995, and the forgone dynamic payoff as a percentage of profits using a 10% profit margin. A profit margin of 10% is consistent with selected firms' Form 10-K SEC filings in 1995 (when available). Thus, biased beliefs in the first five years of the Acid Rain Program cause firms to forgo an annual dynamic payoff equivalent to 1.6% to 25.7% of their profits, with an average of around 10%.²³ There is a lot of heterogeneity in the forgone dynamic payoffs because of the different sizes of those utilities.

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

²³This is not saying that firms forgo an average of 10% profits due to biased beliefs; dynamic payoffs are not profits. The dynamic payoff is the infinitely discounted sum of static payoffs, which drive a firm's behavior. Under cost-of-service regulation, the static payoff to a utility includes more than just the financial profits; for example, it also includes some portion of operating costs to show prudence, as discussed in Section 1.4

Table 1.8: *Cost of biased beliefs in Phase I to firms that have had compliance obligations under the Acid Rain Program 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).

Table 1.9: *Aggregate effects of improving beliefs and reducing volatility.*

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

1.6.2 Implications for Aggregate Environmental and Economic Outcomes

Columns 2 and 3 in Table [1.9](#) report the changes in the aggregate SO₂ emissions and the aggregate production cost, *i.e.*, coal expenditure, if firms with more biased beliefs shared beliefs with the those with less biased beliefs. Column 2 shows that if smaller firms, or firms with below-median generation capacity, had the same belief as did the bigger firms, aggregate emissions would increase by 0.44% and aggregate production costs decrease by 18.53 million 1995 dollars during Phase I, equivalent to an average saving of 1.68 million dollars per small firm. Column 3 shows that if firms in states that still regulate the wholesale electricity market had the same belief as those in states with more competitive pressure because of pending restructuring, aggregate emissions would increase by 47.76% and aggregate production costs decrease by 412 million dollars, averaging 29.43 million dollars per always-regulated firm. The much larger magnitudes of those changes are due to a few big firms (*e.g.*, the Southern Company) that remain under cost-of-service regulation.

Thus, in Phase I of the Acid Rain Program, as firms improve beliefs, they reduce abatement efforts. Indeed, when firms recognized the time trends in the allowance price process better, they would predict the rises and declines in the allowance price better. Since declines are more

dramatic than rises for the first five years, the improvement in the prediction of declines would be more pronounced than that of rises. Now that firms' beliefs would adjust more downwards about declines than upwards about rises, overall the abatement effort would reduce, leading to lower production cost and higher emissions.

Reducing allowance price volatility by price floors and ceilings would change the aggregate emissions and production costs in the same direction. Column 4 in Table [1.9](#) shows that an allowance price floor of \$50 combined with a ceiling of \$200 would increase aggregate emissions by 7.95% and reduce the aggregate production costs by 83.97 million dollars, averaging 9.07 million dollars per firm.

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

An emission tax is an alternative to a cap-and-trade program and another prominent example of market-based environmental regulation. The last column of Table [1.9](#) quantifies the price discovery feature of the cap-and-trade program. Under a tax, the biased initial expectation about the marginal abatement cost lingers on. With a cap-and-trade program, however, the market would gradually sort out the truth, as firms start out with biased beliefs but improve those beliefs over time. Indeed, as shown in the last column of Table [1.9](#), price discovery in the first five years of the Acid Rain Program avoids an over-abatement of 34.44% aggregate emissions and saves 905 million dollars in aggregate production costs.

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

To illustrate, Consider a big, rural electric utility such as Southern Company, and a small,

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

For example, index Southern Company by 1 and Atlantic City Electric Company by 2. Assume that $MC_1 = 50 + 2r_1$, $MB_1 = 60$, $MC_2 = 20 + 10r_2$, $MB_2 = 140$, where MC denotes marginal cost, MB marginal benefit, and r emission reduction. Table 1.10 shows in the first row that the net benefit from a tax of 100 is 265. In the immediately following rows, if beliefs align marginal costs with marginal benefits (that is, $MC_1 < MC_2$), the net benefit under a cap-and-trade program will significantly exceed that under a tax under a range of scenarios. Of course, as the remaining rows show, when beliefs misalign the marginal costs and benefits, the net benefit can fall significantly below that under a tax.

Previous literature misses the beliefs as a potential determinant of the relative efficiency of cap-and-trade programs and emission taxes, because of its focus on static models, where firms care about current profits only and beliefs are therefore irrelevant. The literature has proposed that to achieve efficiency, government intervention is necessary, in the form of setting up trading ratios. I show that once we move beyond static models and starting using a dynamic model, where beliefs are the key driver of behavior, beliefs are a decentralized channel that affects the efficiency of cap-and-trade relative to tax.

Table 1.10: *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

1.7 Conclusion

This paper studies the evolution of firms' beliefs about future market conditions in a new environment. The empirical context is the U.S. Acid Rain Program, the first cap-and-trade program and a landmark experiment in market-based environmental policy. I use structural dynamic estimation to back out coal-dependent private electric utilities' beliefs about the future SO₂ allowance price from their compliance behavior. The three-dimensional continuous-state mixed-continuous-action dynamic model of allowance trades, coal quality choices, and fuel-switching investment offers the first empirical model of firm behavior in cap-and-trade programs.

I find that firms have biased beliefs about the future allowance price. The beliefs are not consistent with the stochastic process of the allowance price. In particular, firms underestimate the role of market fundamentals as a driver of the allowance price. This cost the firms dynamic payoffs equivalent to an average of about 10% of their profits in the first five years of the program. Under cost-of-service regulation, this is the lower bound on the dynamic savings to ratepayers. Smaller firms and firms facing less competitive pressure have more biased beliefs.

Over time, firms' beliefs appear to converge towards the stochastic process. This is consistent with the shift in the management practice in electric utilities during that period; compliance decisions were made less by engineers but more by people experienced in markets and trading.

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

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

This research points to several future research efforts. First, how do firms' beliefs evolve in those states after they underwent electricity restructuring? Intuitively, competition should improve belief formation. Answering that question requires a competitive model of firm behavior in place of a single-agent model as constructed in this paper²⁴. Another complication is the potential need to model the dispatch decision between coal and gas, and to model the capital investment to switch to or co-fire with gas, as gas became competitive with coal.

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

Second, does the experience that firms have gained in the Acid Rain Program spill over to later cap-and-trade programs, such as the NO_x Budget Program?²⁵ Third, it would be interesting to collect organizational data to formally investigate the exact forms of shifts in the management practice in electric utilities, and study their impact on belief formation. Fourth, while this paper only estimates the beliefs, it would be useful to investigate what learning rules have led to those beliefs.

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

Chapter 2

Collective Action in an Asymmetric World¹

2.1 Introduction

Prior to the 21st Conference of the Parties (COP21) to the United Nations Framework Convention on Climate Change in Paris in December 2015, the European Union and almost all individual nations had submitted their Intended Nationally Determined Contributions (INDCs). Each INDC lays out the climate action the submitter intends to take under the new international climate agreement, the Paris Agreement. An important component of the INDCs are the nations' intended reductions in greenhouse gases (GHGs). Despite enthusiasm for the Paris Agreement, including from economists involved, those pledged reductions (assuming optimistically that they are met) are unlikely to come close to controlling GHGs to the overall level that the scientific community generally agrees is needed to prevent the climate change problem from becoming significantly worse over time.²

¹Co-authored with Richard Zeckhauser.

²Indeed, Hohne et al. (2015) finds that the unconditional reduction pledges in INDCs would lead to a median warming of around 2.7 degrees Celsius by 2100 with the full range of 2.2 to 3.4 degrees Celsius. European Commission (2017) also states that "[the INDCs] are not yet enough to keep global warming below 2 degrees Celsius".

Traditional and widely proposed mechanisms for pursuing efficient global public good provision in international agreements will fail to control GHGs. Solutions that impose a common requirement, such as a uniform global carbon tax or a uniform GHG reduction obligation, will not get widespread agreement among asymmetrically-situated nations. The alternative approach, which is to get nations together to agree on differentiated burden-sharing, may secure agreement, as we have seen for example with the Kyoto Protocol and the Paris Agreement. However, there remains too strong an element of voluntary decisions in such agreements, as no overarching government unit is available to enforce actions in providing for the global public good. Behaviors of individual nations in this situation will then be best characterized by a Nash equilibrium, which, as is well known in the literature, leads to under-provision of the global public good.

The first goal of this paper is get to the crux of the under-provision problem by looking more closely than does the literature at the behavior of asymmetrically-situated nations in global public good provision. Doing so allows us to gain a deeper understanding of the problem and to better assess the prior solution proposals. We formalize the notion in [Olson \(1965\)](#) of “exploitation of the great by the small”. We empirically test this notion in the context of global climate change mitigation using the INDCs submitted prior to the Paris conference. As hypothesized, “small-interest” nations made disproportionately smaller pledges than did “large-interest” nations. The measure of “interest” turns out to be predominantly based on the Gross National Incomes (GNIs) and past emissions, and neither vulnerability to climate change nor per capita income seem to be relevant. To our best knowledge, this is the first time that the important role of asymmetric incentive of small-interest and large-interest nations has been recognized in the context of global climate agreements.

The paper then proposes a new solution, the Cheap-Riding Efficient equilibrium (CREE).³ It fully recognizes such asymmetries. It defines the relative contributions of players of differing

³While the term “free riding” is typically used to describe such situations, in many cases we see less extreme behavior, which we term “cheap riding”. If the potential contributor gets a substantial portion of the benefits from a public good, or if s/he enjoys separate benefits from the action of contributing that are quite apart from the public good, then s/he will likely contribute a positive amount, although that amount would still be below what efficiency would require.

size in a manner that both caters to the strong incentive of small-interest players to ride cheaply, yet still achieves Pareto optimality. In this equilibrium, we establish the Nash equilibrium as the starting point. From there, we consider two Pareto-improving paths to the Pareto frontier. One adapts the Nash Bargaining solution; the other relies on the principle of the Lindahl equilibrium. In our illustrative numerical example, the two outcomes are remarkably similar. We recognize that such mechanisms are not currently in place, and that most international efforts at providing global public goods lead to total outputs that are well below what would be optimal. In other words, there are alternative agreements, CREE included, that require greater contributions from all that would bring about a substantial Pareto improvement.

Numerous papers have examined the under-provision problem of public goods with no central authority, including the effect of group size (for example, [Isaac and Walker \(1988\)](#)), the implications of individual heterogeneity (for example, [Boadway and Hayashi \(1999\)](#), [Callander and Harstad \(2015\)](#)), the role of uncertainty (see [Kolstad \(2007\)](#) and the literature reviewed therein), and the validity of the Nash assumption (for example, [Cornes and Sandler \(1984\)](#), [Sugden \(1985\)](#)).

A number of solutions have been proposed in the literature. [Arce M. and Sandler \(2001\)](#) propose setting up a super-national organization that sends signals to nations in order to induce correlated equilibria that are Pareto superior to Nash equilibria. [Gerber and Wichardt \(2009\)](#) study a two-stage mechanism where players commit to the public good by paying a deposit prior to the contribution stage. They show that properly designed deposits support prior commitment and full ex post contributions as a sub-game perfect Nash equilibrium. [Barrett \(1994\)](#) represents public good provision in a repeated prisoners' dilemma game and shows that cooperation can be both individually and collectively rational. Similarly, [Heitzig et al. \(2011\)](#) cast the public good provision game in a repeated game setting and argue that dynamic concerns can enforce efficiency. [Abul Naga and Jones \(2012\)](#) discuss the role of other-regarding preferences, such as altruism, in bringing about efficient provision. One strand of literature studies matching schemes (for example, [Barrett \(1990\)](#), [Falkinger et al. \(1996\)](#), [Boadway et al. \(2011\)](#), [Buchholz et al. \(2011, 2014\)](#)), in which players decide on the

unconditional and conditional (matching) contributions to the public good. This process can lead to interior matching equilibria at which all agents make strictly positive unconditional contributions. However, none of these solutions pays sufficient attention to the asymmetry inherent in players' situations.

We proceed as follows. Section [2.2](#) connects the climate change mitigation problem, or provision of a global public good in general, to the traditional Alliance problem. Section [2.3](#) presents empirical evidence demonstrating cheap riding in the INDCs that nations submitted for the Paris Agreement. Section [2.4](#) discusses weaknesses of prior proposals for solving the under-provision problem; they do not adequately recognize the disparate incentives of players to ride cheaply. We propose our solution, the CREE, in Section [2.5](#). Section [2.6](#) concludes.

2.2 Climate Change Mitigation - A Traditional Alliance Problem

The provision of global public goods is challenging for two reasons. First, forces that at times motivate contributions by individuals - such as warm glow, prestige or self interest - will rarely be sufficient to motivate nearly sufficient contributions by nations to a public good that entails substantial expenditures, as does GHG reduction. Second, and more importantly, the potential providers are very differently situated. Some are rich, some poor; some are large, some small, and some bear much greater responsibility than others for past emissions. At any level of the public good, some will secure much greater benefits from its provision - both overall and at the margin - than will others. Thus, China with a population of 1.3 billion, the world's second largest economy, a significant pollution problem, and the intention and ability to lead in the production of green energy technologies, will benefit greatly from the effective control of GHGs. Landlocked Laos, with fewer than 7 million citizens, would benefit little.

The theory of alliances ([Olson and Zeckhauser, 1966](#)) was developed to address just such a situation. Its principal lesson is that larger nations, as measured by national incomes, will contribute disproportionately more to the alliance good (e.g., the defense budget of NATO)

than smaller nations. This lesson has been generalized to any global public good: nations with larger interests in a global public good will contribute disproportionately more to its provision. Here, we use “interest” to refer to the marginal benefit that the nation would get from the global public good if it contributed more to it by itself. Because this marginal benefit generically varies with the allocation of contributions, we need to specify the latter to complete the definition of “interest”. A useful specification is the proportional contribution plan, with the proportion based on some salient measures of the contributing nations. Thus, a nation has a larger interest than another if it would be willing to contribute more than its proportional share justified by the salient measure.

In combating common security threats, such as in the context of NATO, the salient measure is the national incomes, that are at stake if the threat eventuates. Thus, a nation that has twice as great a national income as another would contribute more than twice as much at a self-interested, Nash equilibrium. That nations with greater national incomes tend to contribute disproportionately more is also observed in the combat against ISIS, in which the U.S. provides vastly disproportionately; Pentagon officials complained that some of the 64 partner nations and regional groups, mostly the Arab allies which though threatened more are smaller nations, are not doing nearly enough (Hennigan and Bennett, 2016). This pattern also emerges in non-defense contexts, for example, where Saudi Arabia cut its oil production far more below its optimum than did other OPEC nations when OPEC was still hanging together to cut production. Indeed, Griffin and Xiong (1997) find that small producers (such as Gabon, Qatar, Algeria, Libya, Indonesia and Nigeria) were subsidized at the expense of large producers (especially Saudi Arabia) in the OPEC arrangement, a phenomenon they call the “small producer bias”.

In the context of climate change mitigation, the salient measure can be much more than just national incomes. It may include other measures of tangible benefits from climate change mitigation, such as green orientation, per capita income, and the degree to which a nation would be affected - whether due to damages or costs of avoiding those damages - by climate change. Greater interest could also arise from intangible benefits. Thus, a nation that has

emitted a lot in the past might commit to greater reductions due to a feeling of responsibility, a concern for reputation, or an orientation toward fairness. In the next section we will empirically investigate which factors prove important.

Given the under-provision of a global public good due to the strong cheap-riding incentive that small-interest nations have in a Nash equilibrium, the question arises as to why the nations do not get together and bargain their way to an efficient equilibrium (thus bypassing the Nash equilibrium). Such an approach might have potential if the members were all similarly situated, and the number of members were not too large. In a negotiation, a natural focal point in the sense of Schelling (1960) would be that each member contributes the same; none could expect to pay less than the others. In such a symmetric situation, with only a few players, they could merely identify and agree to the optimal per-capita contribution of individuals, or per-nation contribution in the international context. Positing that contributions could be monitored, efficiency would be achieved.

In the real-life situation, however, matters are far from symmetric. Even if the individuals across the nations were identical, for example in income and preferences, in a negotiation, the small-interest nations could expect the large-interest nations to contribute more. They would argue, correctly, that given proportional contributions large-interest nations benefit much more at the margin. However, determining the appropriate ratio of contributions via negotiation would present a challenge. Large- and small-interest nations would respectively advance arguments as to why the ratio should be smaller or larger. As a result, agreement is unlikely. With each nation following its own principles, under provision is to be expected.

2.3 Cheap Riding in the Paris Agreement

In this section, we empirically document the existence of cheap riding in the INDCs submitted for the Paris Conference. First, we convert the reduction goals in INDCs to absolute amounts of carbon emissions reduction, or the contributions, by estimating each nation's business-as-usual (BAU) emissions in the target year and comparing them to the target emissions. Second, we examine the rank correlations between the individual nations' contributions, properly

normalized, and their national income, climate vulnerability, pollution level, historical carbon emissions, and measures of their environmental concerns. We describe the technical details in Appendix [B.1](#).

Table [2.1](#) shows the Pearson rank correlation test results. The rank correlation coefficient between the *reduction per dollar of GNI* and GNI is positive and highly significant. This means that nations with larger GNI's make disproportionately greater pledges of reductions relative to their GNI's. That GNI appears to be an important determinant of the cheap-riding incentive is consistent with the findings in the NATO defense context ([Olson and Zeckhauser, 1966](#)) and the OPEC production-reduction context ([Griffin and Xiong, 1997](#)). Per capita GNI produces a positive rank order correlation as well, in accord with the observation that environment is a superior good ([Kahn and Matsusaka, 1997](#)).

From Table [2.1](#), some observers might be surprised that the vulnerability measures do not produce big contributions. Indeed, they point in the other direction. Some might argue that while GNP seems an obvious measure of size that matters in the defense context, it should matter much less in the context of climate change mitigation. The Maldives, a minuscule nation relative to the United States is often cited in this context. Its very existence would be at risk given sea level rise. Thus, the argument goes, it should contribute a lot in proportion to its GNP to the mitigation cause. We believe, however, that there are clear parallels between the defense and the climate-change-mitigation contexts. The national income at stake is still a more important measure of size than vulnerability for either problem. If the counterargument is based on vulnerability (as with the Maldives argument) in the climate-change-mitigation context, then in the 1960s NATO context Germany and France should have contributed much more relative to the United States. After all, they were much more vulnerable to aggression by the Soviet Union. The large nation in NATO contributed significantly more disproportionately despite having an ocean of protection.⁴ If an alternative counterargument is based on per capita income, then Luxembourg, having a very high per capita income, contributed vastly

⁴That said, vulnerable nations might have incentive to appear to contribute a lot so as to serve as role models. They might say: "Look, if we do not contribute much, others will ask if we are really crying wolf, and do not feel threatened, or do not care."

less than the United States. Moreover, with NATO in the 1960s, there was an established organization that had some ability to get small-interest cheap riders to contribute more. Maybe Germany, France, and Luxembourg would have contributed relatively even less had this overarching organization not existed. To return to climate change, one phrasing of the Maldives' implicit argument might be: "We are a tiny blameless nation. Nothing we could do could alter our dire fate. The blameworthy nations of the world have a responsibility to take vigorous action to save us."

Nations with higher cumulative 1970-2012 carbon emissions⁵ also tend to pledge disproportionately more. This suggests that nations might have incorporated historical responsibility in their preferences when they formulate emission reduction pledges. However, since historical emissions and GNI are highly correlated, a controlled analysis is needed to single out the marginal effect of each.

Table 2.1: *Pearson rank correlation test results.*

Correlation with per-unit reduction	Coefficient	p-value	# of obs
GNI	0.3433	0.000	106
Per capita GNI	0.2547	0.008	106
% vulnerable rural population	-0.1854	0.038	126
% vulnerable urban population	-0.0641	0.478	125
% population exposed to disaster	-0.0563	0.527	129
Annual temperature	-0.1143	0.189	134
Historical carbon emissions 1970-2012	0.7311	0.000	128

We now look more closely at the roles of GNI and historical emissions, as well as per capita GNI, in determining the cheap-riding incentive in a controlled analysis. We conduct the Kendall partial rank correlation test for each one with the properly normalized contribution while controlling for the other two. The analysis also controls for the vulnerability, pollution, and environmental concern measures.

Table 2.2 shows that both GNI, as a measure of tangible benefit, and historical emissions, as a measure of intangible benefit, survive as robust determinants of the cheap-riding incentive

⁵We thank a referee for suggesting the role of historical emissions in driving the pledges.

after adding the controls. Per capita GNI has no relationship. Thus, for two hypothetical nations with the same historical emissions (and per capita GNI, climate vulnerability, pollution level, and environmental concerns), nations with larger GNIs would pledge disproportionately more and thus suffer cheap riding from nations with small GNIs. Similarly, for two hypothetical nations with the same GNI (and other controls), nations that have cumulatively emitted more would pledge disproportionately more.⁶

Table 2.2: *Kendall partial rank correlation test results.*

Partial correlation with per-unit reduction	GNI	GNI	Per capita GNI	Per capita GNI	Historical emission	Historical emission
Coefficient	0.1084	0.1703	0.0943	0.0219	0.6168	0.5920
p-value	0.045	0.020	0.197	0.857	0.000	0.003
Environmental concern	N	Y	N	Y	N	Y
Other controls	Y	Y	Y	Y	Y	Y
# of obs	59	29	59	29	59	29

Notes. “Other controls” for a particular outcome variable include the other two outcome variables, as well as vulnerability and pollution measures.

2.4 Cheap Riding and Prior Proposals

We evaluate two important prior proposals to try to achieve efficient levels of global climate change control. Despite the fact that cheap-riding incentives loom large in this context, as we have shown in Section 2.3, neither proposal recognizes this problem.

First, carbon taxes, along with tradable permits, are the economist’s preferred regulatory tool for environmental externalities. To achieve efficiency, a necessary condition is that the tax or the permit price be the same for all nations. But small-interest nations are getting much less benefit relative to what they are spending on the margin. Thus, a uniform tax or permit price ignores the cheap-riding incentive; small-interest nations are likely to find it individually irrational to participate in the uniform-price regime. Transferring vast resources

⁶Historical emissions are fairly strongly related to GNI. Normalizing both quantities so that the median country is at 100, eight countries have normalized emissions more than 2.5 times GNI. They are the Russian Federation and three former Soviet countries, China, Iran, South Africa and Trinidad and Tobago.

from large-interest to small-interest countries would be highly unattractive politically⁷

In his seminal work on Climate Clubs, Nordhaus (2015) illustrates how the Climate Club would work, where all members of the Club would agree to impose a minimum domestic carbon tax of \$25/ton.⁸ This is subject to the same critique of the carbon taxes/permits above. Thus, the likely result would either be that many nations would simply not join the Club, or that to get them to join the price would have to be set far below what is desirable. Nordhaus (2015) deals with the non-joiners' strategy by having members impose tariffs on the cheap-riding non-joiners. Whether such an arrangement could work in practice, given concerns about retaliation from the non-joiners, violation of existing trade agreements and other challenges, has been carefully considered by Nordhaus and hotly debated by others. Addressing that debate would take us well beyond the concerns of this paper. But it is appropriate for us to point out that such an arrangement does not adequately recognize the bargaining power that small-interest beneficiaries have when it comes to the provision of a global public good.

The implication is that major beneficiary nations, like the United States and China, which have by far the largest GNIs of any nations, would find a uniform carbon tax, or joining the Climate Club, strongly in their interest. However, nations with much smaller GNIs, per our findings in Section 2.3, would likely benefit much less at the margin. Small-interest nations would correctly point out that the strategic situation tilts in their favor. They could feel entitled to impose a much lower carbon tax than the United States and China, the type of result to be expected in a Nash equilibrium.

Given these challenges, we developed an alternative solution that recognizes the bargaining forces that arise because nations differ significantly in the benefits they receive at the margin from a global public good.

⁷That being said, the Paris Agreement's Article 6 has provisions for "internationally transferred mitigation outcomes", or ITMOs, which provide parties with the mechanism to trade emission reduction credits for funds.

⁸Nordhaus (2015) allows individual nations to impose a higher tax. But different taxes in different nations would sacrifice efficiency.

2.5 Cheap-Riding Efficient Equilibrium (CREE)

We propose a new solution, the Cheap-Riding Efficient equilibrium (CREE), that both achieves efficiency and respects cheap-riding incentives. This solution takes the Nash equilibrium as the starting point and then follows Pareto-improving paths from there to efficiency. First, we set up a simple global public good provision game. Second, we demonstrate that the Nash equilibrium respects cheap-riding incentive but is inefficient. Third, we demonstrate that Lindahl equilibrium does not respect cheap-riding incentives, despite being efficient. We then construct CREE by combining the best of the two: the Nash way of recognizing the cheap-riding incentive and the Lindahl way of achieving efficiency. We also note that replacing the latter with the Nash Bargaining way of obtaining efficiency also achieves our purpose. Thus, we see CREE as a *class* of solutions, which share the Nash equilibrium as the starting point but differ in the path they take to reach the Pareto frontier.

2.5.1 Model Setup

Index the players in the model by $i \in \mathcal{N} = \{1, \dots, n\}$, and let the contribution of each player to the public good be $m_i \geq 0$. For simplicity, side payments are ruled out. The public good is simply the sum of all individual contributions, $M = \sum_i m_i$. Denote the sum of contributions by players other than i by $M_{-i} = \sum_{j \neq i} m_j$. Player i gets benefit $V_i(M)$, tangible and intangible, from the public good.

We allow not only for cash contributions, but also for in-kind contributions (the norm with international public goods). Hence we allow players to receive private benefits from their own contributions, quite apart from the public good. Thus, for example, a nation's armed forces contribute to the deterrent level of a military alliance. However, they are also available to assist in the case of a natural disaster. A nation's efforts to curb GHGs by developing clean energy technologies would simultaneously advance its high tech capabilities. In addition, a nation may value the respect that other nations pay to it when it contributes to the public good, or may receive a warm glow. Represent this private benefit as $B_i(m_i)$.⁹

⁹A more complex model would include an additional argument in the benefit function, which would include

The cost to player i of providing m_i is $K_i(m_i)$. Note, because we are dealing with a situation where contributions are in kind, we might expect the marginal cost of contribution to increase sharply as one contributes more; that is, $K''(\cdot)$ can be not only positive, but significantly so. Player i 's net payoff is thus $U_i(M, m_i) = V_i(M) - K_i(m_i) + B_i(m_i)$.

For notational simplicity, in what follows we will just use $C_i(m_i)$ to represent the net private cost, $K_i(m_i) - B_i(m_i)$, so that the utility functions can be written as $U_i(M, m_i) = V_i(M) - C_i(m_i)$. Throughout, we assume that utility functions are common knowledge.

2.5.2 Nash Equilibrium

A Nash equilibrium is an allocation $(m_i^N)_{i \in \mathcal{N}}$ (where the superscript N stands for Nash) such that each player's choice is a best response to what the others do. That is, for each $i \in \mathcal{N}$:

$$m_i^N \in \arg \max_{m_i \geq 0} U_i(m_i + M_{-i}^N, m_i), \quad (\text{Nash})$$

where $M_{-i}^N = \sum_{j \neq i} m_j^N$. There exists a unique Nash equilibrium if for each i , $V_i' > 0, C_i' > 0, V_i'' < 0, C_i'' > 0$, as shown by applying Proposition 3.1 in [Cornes and Hartley \(2007\)](#).

We now examine cheap riding in the Nash equilibrium. Posit for now that nations differ only in size (for example, population, or national income), the “exploitation of the great by the small” in [Olson \(1965\)](#) refers to the result that large nations will contribute more than in proportion to their size at the Nash equilibrium. Taking GNP as the measure of nation size, [Olson and Zeckhauser \(1966\)](#) provide an intuitive proof as follows: suppose by way of contradiction that the large nation, which has twice the GNP of the small nation, contributed twice as much at the Nash equilibrium. The Nash equilibrium requires that the marginal rate of substitution of the private good for the public good (MRS) in each nation equal the marginal cost, which, in our setup, means that $V_1'(M)/C_1'(m_1) = V_2'(M)/C_2'(m_2) = 1$. Let 1 be the large nation and 2 the small. This requirement is not met at the proportional contribution plan, where the large nation contributes twice as much. Indeed, at this contribution plan, it

a nation's historical emissions, green orientation, etc. Our finding that historical emissions help to explain the nations' pledges indicates that such factors, which produce intangible benefits, can play an important role.

is reasonable that $V_1'(M) = 2V_2'(M)$ because there is twice as much national income to be protected in 1 as in 2, and that $C_1'(m_1) = C_2'(m_2)$ at $m_1 = 2m_2$ because there is twice as much national income that can be taxed in 1 as in 2. Then, the MRS of the large nation is twice that of the small. This implies that the large nation would want to contribute even more and/or the small nation even less, until they reach the equilibrium requirement.

These insights can be generalized as follows. Take a proportional contribution plan as given. The proportion can be based on a single salient measure, such as the GNP in the defense context, or can be based on a combination of multiple salient measures, such as GNP, population, vulnerability, and green orientation combined in the climate change control context. At the given proportional contribution plan, calculate the MRS of each nation. Unless all the MRSs are equal, nations have incentive to deviate from this proportional contribution plan; nations with larger MRS, call them “larger-interest” nations, would like to contribute more, and nations with smaller MRS, call them “smaller-interest” nations, would like to contribute less. Therefore, the smaller-interest nations effectively ride cheaply on the larger-interest nations by contributing disproportionately less. Formally, if the contribution plan assigns a proportion s_1 to Nation 1 and s_2 to Nation 2, and the MRSs are such that $V_1'(M)/C_1'(m_1) > V_2'(M)/C_2'(m_2)$ for $m_1/m_2 = s_1/s_2$, then Nation 1 would deviate from the contribution plan by increasing m_1 or Nation 2 would deviate by decreasing m_2 , leading to a disproportionate plan where $m_1/m_2 > s_1/s_2$.¹⁰

As a numerical example, we posit that $V_i(M) = a_i \lambda \log M$, $C_i(m_i) = m_i^2/a_i$, where λ and a_i are parameters. We can think of a_i as population, so that the benefit of the global public good M is proportional to the population and the private cost is inversely proportional. Note that the per capita incomes are the same in the two nations; indeed, the marginal cost of private contributions, $C_i'(m_i)$, are the same at the same per capita contributions. The

¹⁰We thank a referee for pointing out that Warr (1983) looks at the effect of varying incomes on private contributions in a Nash equilibrium: a player will contribute less as his or her income is transferred away. Our formulation of the model is different from Warr’s in that we do not treat income explicitly, and that we summarize the difference among nations by the marginal rate of substitution, which can depend on many things in addition income.

parameter λ is simply a common scalar on the benefit. Then:

$$MRS_i(m_i, M) = \frac{V'_i(M)}{C'_i(m_i)} = \frac{a_i \lambda / M}{2m_i / a_i} = \frac{a_i^2 \lambda M}{m_i \cdot 2}.$$

Let $a_1 = 1, a_2 = 1/4$. Take the proportional contribution plan as one in which Nation 1 contributes four times as much as Nation 2 (the former has four times as many people). The ratio of the MRSs is:

$$MRS_1(m_1, M) / MRS_2(m_2, M) = \frac{a_1^2}{m_1} / \frac{a_2^2}{m_2} = \frac{1}{4} / \frac{1/16}{1} = 4.$$

The large nation's MRS is four times as big as the small's at a four-to-one contribution; it has four times as large a marginal benefit and the same marginal cost as the small. Our calculation of the Nash equilibrium shows that the large nation contributes 1.372 at the Nash equilibrium, which is sixteen times the small nation's contribution of 0.086, although the former is only four times the size of the latter.

This Nash equilibrium outcome is also far from Pareto optimal. The large nation gets a net payoff of -0.375 and the small nation a net payoff of 0.347. But if the large nation contributed 20% more and the small nation gave twice its original contribution, those net payoffs would rise to -0.319 and 0.480, a major Pareto improvement, though this outcome also is far from Pareto optimal.

2.5.3 Lindahl Equilibrium

[Lindahl \(1958\)](#) conceived of a provision scheme where each player reported how much of the public good s/he wants depending on the share s/he would be required to pay. The cost shares would be determined such that each player desired the same amount of the public good. Formally, a Lindahl equilibrium established individualized public good prices (or cost shares) $(p_i^L)_{i \in \mathcal{N}}$ (where the superscript L stands for Lindahl), a private good price which we normalize

to 1, and an allocation $(M^L, (m_i^L)_{i \in \mathcal{N}})$ such that for each $i \in \mathcal{N}$:

$$(m_i^L, M^L) \in \arg \max_{m_i, M \geq 0} U_i(M, m_i),$$

subject to: $p_i^L M \leq m_i,$ (Lindahl)

and the market clears: $M^L = \sum_i m_i^L$. There exists a unique Lindahl equilibrium if for each i , $V_i' > 0, C_i' > 0, V_i'' < 0, C_i'' > 0$, by applying Proposition 1 in [Buchholz et al. \(2008\)](#).

Despite achieving efficiency, the Lindahl equilibrium suffers a grave defect: it fails to recognize the incentive to ride cheaply. Thus, in our numerical example, the Lindahl equilibrium has the nations contributing 1.414 and 0.354 respectively, which gives them net payoffs of 0.279 and 0.070 respectively. Here, however, the small nation has a simple threat to make to the large nation: “I will not participate in the Lindahl equilibrium. You can do so, or if you choose we can revert to the Nash equilibrium.” The threat is credible, since the small nation is better off at the Nash equilibrium than at the Lindahl equilibrium.

In many asymmetric situations, of course, no nation will be worse off at the Lindahl equilibrium than at the Nash equilibrium.¹¹ Nevertheless, the fact that the Lindahl equilibrium simply ignores the fact that small-interest nations do relatively much better at the Nash than at the Lindahl equilibrium is critical. To get agreement, any solution must recognize this bargaining advantage possessed by small-interest nations, for the Nash equilibrium is the fallback solution if an agreement is not reached. This implies that small-interest nations are likely to balk at the Lindahl equilibrium. For some parameter values, as we just showed, there will be a smaller-interest player who is strictly worse off at the Lindahl equilibrium, who would simply hold out for the Nash equilibrium. This player would be in a favored bargaining position just because of its small interest. It could be otherwise identical to the other players in terms of per capita income, exposure to threats due to insufficient provision of the public good, and costs of provision.

Appendix [B.2](#) offers an alternative, intuitive way of thinking about the Lindahl equilibrium,

¹¹This would typically be the case if nations were perfectly symmetric. [Shitovitz and Spiegel \(1998\)](#) and [Buchholz et al. \(2006\)](#) study the conditions on the income distribution that makes all players prefer the Lindahl equilibrium to the Nash equilibrium.

the Supply-Demand Arrangement (SDA). We start by asking each player how much s/he is willing to supply if it will be matched by a total of contributions from others that is n times as large. By varying n in $(1, \infty)$, we obtain that player's supply curve, with the horizontal axis representing n and the vertical representing the supply. We can also derive that player's demand curve by multiplying the supply curve by n . The Supply-Demand Arrangement just solves for a vector of matching ratios, where supply equals demand for each player.

2.5.4 CREE

There are, of course, an infinite number of Pareto optimal outcomes. Our approach employs a method that takes account of differential incentives to ride cheaply, and at the same time provides an intuitively appealing mechanism to get to alternative Pareto optimal outcomes. Our proposed mechanism, Cheap-Riding Efficient equilibrium (CREE), starts by establishing the Nash equilibrium as a base point. The question then is how to proceed from there in an intuitively appealing manner that leaves no player worse off (so that no player will have the incentive to fall back to the Nash equilibrium), while achieving efficiency. Here too, there are an infinite number of possibilities. The basic question is how the players should share the surplus above the Nash equilibrium on the path to the Pareto frontier.

We propose two alternative approaches to sharing the surplus over the Nash equilibrium. A CREE with our first approach uses the thinking of the Lindahl equilibrium to define the further path from the Nash equilibrium to the Pareto frontier. We believe that the prominence of the Lindahl equilibrium in the public goods context, plus its coincidence with the natural Supply-Demand Arrangement interpretation (see Appendix [B.2](#) for details), gives this CREE a particularly strong claim as a focal point.¹² A CREE with the second approach maximizes the product of each player's surplus over the Nash equilibrium, the outcome with the Nash Bargaining solution. This arrangement offers two advantages: it is the most widely applied

¹²The CREE-Lindahl degenerates to a Lindahl equilibrium if we take the Nash equilibrium allocation to be the starting point for the Lindahl process. They will also be equivalent if the players have identical preferences (e.g., have the same GNIs if it is nations, and that is the sole driving factor), hence are symmetrically situated, and we focus on the focal point solution at which all players contribute the same amount.

bargaining solution, and it derives from an appealing set of axioms. We emphasize the need to take the Nash equilibrium as the disagreement point for the Nash Bargaining formulation to apply to our global public goods provision problem.

It is worth noting that while a CREE is not a Nash equilibrium in individual contributions,¹³ it is in individual participation. Consider the following game. Countries are deciding whether to join the CREE agreement, which specifies the individual contributions. If at least one of them does not join, there will be no CREE agreement and countries will revert to the Nash equilibrium. Otherwise, the CREE agreement is effective. Then, for each country it is a dominant strategy to join the CREE agreement, and thus CREE is a Nash equilibrium in individual participation. To the best of our knowledge, CREE is probably the simplest Pareto efficient mechanism in global public good provision games that achieves full participation, without any need to distinguish between signatories and non-signatories as is required by many other mechanisms in the literature.

The Lindahl path. In a CREE-Lindahl, the cost sharing scheme is such that each player will want the same **additional** amount of the public good on top of the Nash amount. Formally, a *CREE-Lindahl* consists of individualized public good prices (or cost shares) $(p_i^{CRE})_{i \in \mathcal{N}}$ (where the superscript *CRE* stands for Cheap-Riding Efficient), a private good price which we normalize to 1, and an allocation $(M^{CRE} + M^N, (m_i^{CRE} + m_i^N)_{i \in \mathcal{N}})$ such that for each $i \in \mathcal{N}$:

$$(m_i^{CRE}, M^{CRE}) \in \arg \max_{m_i, M \geq 0} U_i(M + M^N, m_i + m_i^N),$$

$$\text{subject to: } p_i^{CRE} M \leq m_i, \quad (CRE)$$

and the market clears:

$$M^{CRE} = \sum_i m_i^{CRE},$$

where $(m_i^N)_{i \in \mathcal{N}}$ is a Nash equilibrium allocation and $M^N = \sum_{i \in \mathcal{N}} m_i^N$.

The following proposition establishes the existence and uniqueness of the Cheap-Riding

¹³In practice, as with many international agreements, reputational concerns or fear of sanctions make it more likely that nations will adhere to the CREE agreement after they sign to become part of it.

Efficient equilibrium. There exists a unique CREE-Lindahl if for each i , $V_i' > 0, C_i' > 0, V_i'' < 0, C_i'' > 0$. Take the total contribution from the unique Nash equilibrium as M^N , with individual contributions m_i^N . Then, we can re-write the utility functions as $\bar{U}_i(\cdot, \cdot) = U_i(\cdot + M^N, \cdot + m_i^N)$, so that Problem (CRE) is in effect the same as Problem (Lindahl) with utility functions $\bar{U}_i(\cdot, \cdot)$, in which there exists a unique Lindahl equilibrium by applying [Buchholz et al. \(2008\)](#). Therefore, there exists a unique CREE-Lindahl.

In our numerical example, the (unique) CREE-Lindahl has the nations contribute (0.252, 0.123) in addition to Nash contributions. This results in the total contributions (1.624, 0.209), and net payoffs (-0.215, 0.432), respectively. This outcome is Pareto optimal. The next proposition shows that CREE-Lindahl generally achieves Pareto optimality.

Under the conditions in [Proposition 2.5.4](#), the CREE-Lindahl allocation is Pareto optimal. [Bergstrom \(1973\)](#) establishes Pareto optimality of the Lindahl equilibrium.¹⁴ We adapt that proof here. Suppose by way of contradiction that there is a CREE-Lindahl allocation $(M^{CRE} + M^N, (m_i^{CRE} + m_i^N)_{i \in \mathcal{N}})$ that is not Pareto optimal. Then there is an alternative allocation $(\tilde{M}^{CRE} + M^N, (\tilde{m}_i^{CRE} + m_i^N)_{i \in \mathcal{N}})$ such that Player i does strictly better and the other players do weakly better. Because the utility functions are strictly monotone, they represent locally non-satiated preferences. Hence, Player i 's constraint must be violated and the other players' constraints weakly violated at the alternative allocation, that is $p_i^{CRE} \tilde{M} > \tilde{m}_i$, and $p_j^{CRE} \tilde{M} \geq \tilde{m}_j$ for $j \neq i$. Adding up these inequalities gives $(\sum_{i \in \mathcal{N}} p_i^{CRE}) \tilde{M} > \sum_{i \in \mathcal{N}} \tilde{m}_i$. To ensure market clearance, $\sum_{i \in \mathcal{N}} p_i^{CRE} = 1$. Therefore, $\tilde{M} > \sum_{i \in \mathcal{N}} \tilde{m}_i$, a contradiction.

The next proposition shows that the CREE-Lindahl takes care of the cheap riding incentives, in the sense that each player prefers the CREE-Lindahl allocation to the Nash equilibrium allocation, and usually strictly so. Intuitively, when asked how much of the public good above the Nash outcome is desired at a given cost-sharing rule, any individual player can always choose to desire nothing (and thus to contribute nothing) on top of the Nash outcome. Hence, they cannot be worse off at CREE-Lindahl than at the Nash equilibrium. Furthermore, for a player who would contribute a positive amount at the Nash equilibrium and partially

¹⁴He requires local-non-satiation for the public good, as is guaranteed in our context.

contribute in CREE-Lindahl, s/he would strictly prefer to participate in CREE-Lindahl. Indeed, since such a player would already have equated marginal benefit with marginal cost at the Nash equilibrium, contributing a little bit more as specified in CREE-Lindahl would only be marginally more expensive, but would bring non-marginal benefit due to the non-marginal increase in the ‘bang’ in terms of the level of the public good. Moreover, the path from the Nash equilibrium to the Pareto frontier is not ad hoc; it has the appealing justification that led to the initial identification of the Lindahl equilibrium.

Under the conditions in Proposition [2.5.4](#), each player prefers the CREE-Lindahl allocation to the Nash equilibrium allocation. The preference is strict for any player i with $m_i^N > 0$ and $m_i^{CRE} < M^{CRE}$. That each player (weakly) prefers the CREE-Lindahl allocation is obvious by noting that for each player i , the Nash bundle (M^N, m_i^N) corresponds to setting $M = 0, m_i = 0$ in the individual maximization problem in Problem (CRE), which trivially satisfies the constraints therein. In other words, $M = 0, m_i = 0$ is a candidate solution to the individual maximization problem, and therefore the value of the objective function at that candidate solution, or the utility at the Nash equilibrium, is no greater than the optimized value at $M = M^{CRE}, m_i = m_i^{CRE}$.

Now we show that player i strictly prefers the CREE-Lindahl allocation whenever $m_i^N > 0$ and $m_i^{CRE} < M^{CRE}$, by establishing the impossibility of $(0, 0)$ as the solution to the individual maximization problem in Problem (CRE). First, note that $m_i^N > 0$ implies that the individual maximization problem in Problem (Nash) has an interior solution, which implies in turn that the first order condition holds. That is, $V_i'(M^N) = C_i'(m_i^N)$. Now, in the individual maximization problem in Problem (CRE), since the individual constraint obviously binds, we substitute m_i/p_i^{CRE} for M in the objective function and effectively make the problem an unconstrained maximization problem. That problem will have a corner solution if and only if the first order derivative of the objective function with respect to m_i is no greater than zero at $M = 0, m_i = 0$. The first order derivative is $1/p_i^{CRE} V_i'(m_i/p_i^{CRE} + M^N) - C_i'(m_i + m_i^N)$. Evaluating that at $M = 0, m_i = 0$ gives $1/p_i^{CRE} V_i'(M^N) - C_i'(m_i^N) = (1/p_i^{CRE} - 1)V_i'(M^N) > 0$, because of the Nash first order condition and the fact that $p_i^{CRE} = m_i^{CRE}/M^{CRE} < 1$.

The Nash Bargaining path. The Nash Bargaining solution also provides an intuitively appealing way to proceed from the Nash equilibrium to the Pareto frontier while respecting cheap riding incentives. Proposed by [Nash \(1950\)](#), this formulation enjoys strong axiomatic support. Loosely speaking, a Nash Bargaining solution is a vector of payoffs that maximizes the product across all players of the gains over some disagreement point. In our context, the disagreement point is necessarily the Nash equilibrium.¹⁵

Formally, a *CREE-Nash Bargaining* is an allocation $(m_i^{CRE} + m_i^N)_{i \in \mathcal{N}}$ such that:

$$(m_i^{CRE})_{i \in \mathcal{N}} \in \arg \max_{(m_i)_{i \in \mathcal{N}}} \prod_i [U_i(\sum_i m_i + M^N, m_i + m_i^N) - U_i(M^N, m_i^N)],$$

subject to: $U_i(\sum_i m_i + M^N, m_i + m_i^N) \geq U_i(M^N, m_i^N), \quad \forall i \in \mathcal{N},$

where $(m_i^N)_{i \in \mathcal{N}}$ is a Nash equilibrium allocation and $M^N = \sum_{i \in \mathcal{N}} m_i^N$.

It is obvious that the CREE-Nash Bargaining is unique, Pareto optimal, and individually rational, under the conditions for existence and uniqueness of the Lindahl equilibrium.

Interestingly, in our numerical example, though the parameter values were not chosen for this purpose, the CREE-Lindahl allocation is remarkably close to the CREE-Nash Bargaining allocation. In fact, they differ by less than 0.4%: the (unique) CREE-Nash Bargaining has the nations contribute (0.251, 0.123) in addition to Nash contributions, resulting in the total contributions (1.623, 0.209), and net payoffs (-0.213, 0.431), respectively. Future work should determine what degree of closeness applies for other utility functions and other parameter values.

¹⁵From the literature, it does not appear that the disagreement point in a Nash Bargaining formulation should automatically be a Nash equilibrium, either in our context specifically or in general. [Nash \(1950\)](#) does not make this point. In teaching Nash Bargaining, the disagreement value is typically taken to be zero or left unspecified. In applied work, the choice of the disagreement point depends on the specific context. (See, for example, [Horn and Wolinsky \(1988\)](#).) We emphasize that the Nash equilibrium (rather than zero payoffs or any other consequence of disagreement) is necessarily the disagreement point in our global public goods provision context. This is because the Nash equilibrium, importantly, recognizes the cheap riding incentive of small players, the jumping off point of this paper.

2.5.5 Summary of the Numerical Example


Table 2.3 and Figure 2.1 summarize our results. To recap, the Nash equilibrium reflects the incentive to ride cheaply, but it is far from Pareto optimal. The Lindahl equilibrium, which is identical to the Supply-Demand Arrangement, achieves one of the Pareto optimal allocations, but it does not reflect the cheap-riding incentive. Thus, small-interest players are likely to balk at this solution. Indeed, for some parameter values, there will be one player (and possibly more) who is strictly worse off at the Lindahl equilibrium. Such a player would simply hold out for the Nash equilibrium. The CREE-Nash Bargaining and the CREE-Lindahl, however, both achieve Pareto optimality and both respect the cheap-riding incentive, while enjoying intuitive appeal.

Table 2.3: *Allocations at various solutions in our example.*

	Large's contribution	Small's contribution	Total contribution	Large's payoff	Small's payoff
Nash equilibrium	1.372	0.086	1.458	-0.375	0.347
Lindahl equilibrium/SDA	1.414	0.354	1.768	0.279	0.070
CREE-Nash Bargaining	1.623	0.209	1.832	-0.213	0.431
CREE-Lindahl	1.624	0.209	1.833	-0.215	0.432

Notes. The large nation is four times as large as the small nation; the utility functions are $U_1 = 4 \log(M) - m_1^2$ and $U_2 = \log(M) - 4m_2^2$, respectively.

It should be noted that our CREE solutions achieve Pareto efficiency, subject to the constraint that each nation produces its own reductions. Thus, at a CREE formulated with in-kind contributions, such as GHG reductions, it is not possible to reallocate the contributions to increase the welfare of one nation without hurting another. However, they do not achieve *production* efficiency unless one of three conditions holds: 1. The marginal costs of reductions are the same across nations in the relevant range. 2. Nations can buy offsets from other nations to help fulfill their reduction goals. 3. The contributions are made in cash, after which the cash is used to purchase the cheapest available reductions wherever available.



/Users/Cuicui/Dropbox/research/work_with_richard/CREE/supporting_material/untitled

Figure 2.1: *Allocations at various solutions in our example.*

Notes. The CREEs include the CREE-Nash Bargaining and the CREE-Lindahl, which differ by less than 0.4% in our example. The large nation is four times as large as the small nation; the utility functions are $U_1 = 4 \log(M) - m_1^2$ and $U_2 = \log(M) - 4m_2^2$, respectively.

2.6 Conclusion

The sum of voluntarily pledged GHG reduction pledges from individual nations in the Paris Agreement is woefully below what will be required to hold global warming by 2100 below 1.5-2 degrees Celsius as compared to pre-industrial times. This outcome reflects nations' powerful cheap-riding incentives, particularly nations that have lesser interest because of smaller GNIs or more modest historical emissions of GHGs. Hence, the voluntary pledge approach to the provision of global public goods is unlikely to come close to producing an efficient level of total contributions.

This analysis shows that the solution to global climate change must recognize the differential bargaining power of nations, where those receiving greater benefits, tangible and intangible, must contribute more than in proportion to their benefit levels. For many important public goods, including those provided by nonprofit organizations and collections of nations, there is no central authority to both provide the good and levy the exactions to pay for it. Negotiating to efficiency is conceivable when players are symmetrically situated. However, achieving an efficient level of provision with players whose circumstances differ substantially, as is the case with global public goods, encounters the challenge of nations' substantially different strengths of interest, implying greatly differing incentives to ride cheaply.

Our CREE solution explicitly takes account of strength of preference to define a significant starting point, the Nash equilibrium. It then moves by general agreement to a much greater level of reductions. Two principles for structuring such movements are considered, the Lindahl equilibrium and the Nash Bargaining solution. Each enables the players to reach the Pareto frontier. Hence, the mechanism respects uneven cheap-riding incentives, yet still achieves Pareto optimality. Interestingly, in the numerical example considered, the Lindahl and Nash Bargaining paths produce extremely similar results.

There would be many complications in the implementation of CREE to control climate change. Equity is not served when we rely on strength of preference. Side-payments may be necessary to improve production efficiency. Additional mechanisms might be needed to make sure nations stick to the CREE plan, once they have agreed to it. But the overarching lesson is

that it is necessary to have an agreement that incorporates two competing objectives, catering to the unequal bargaining power of those with intense and moderate preferences, and having all players contribute proportional to their preferences. CREE strikes the appropriate balance, and achieves an outcome where all benefit significantly, and none would be better off at the Nash outcome. It also follows comprehensible, logical and justifiable principles.

As citizens of the two nations that lead the world in level of GNI and historical GHG emissions, we recognize that our work here (inadvertently) reveals the weak bargaining positions of our homelands.¹⁶ It seems inevitable that they will have to bear a disproportionate share of the burden if there is to be an effective agreement to arrest climate change through significant reductions in GHG emissions. The leaders of our homelands, one current and the other former, seem to have grasped this when, as early as in November 2014 and then again months before the Paris Conference, they issued two U.S.-China Joint Announcements on Climate Change outlining their ambitions and commitments. Implicitly, these announcements simultaneously recognized the inevitable inadequacies of any agreement that might result should they, the biggest-interest players, insist on proportional burden sharing. This hinted at the ultimate potential for “agreed riding”, a forceful, albeit unbalanced, agreement that respects the bargaining strength of smaller-interest nations.

¹⁶The European Union, which consists of 28 countries, made a unified pledge. Together, the Union’s GNI puts it in first place, slightly ahead of the United States. Its historical emissions would come second, moderately ahead of China.

Chapter 3

Partial Identification in Repeated Games with Imperfect Public Monitoring

3.1 Introduction

Empirical industrial organization has structurally estimated many game theoretic models. Static games have been used to estimate auctions (e.g., [Athey and Levin \(2001\)](#)), entry (e.g., [Bresnahan and Reiss \(1990\)](#); [Berry \(1992\)](#); [Ciliberto and Tamer \(2009\)](#)), oligopolistic competition (e.g., [Genesove and Mullin \(1998\)](#)), and so on. Dynamic games have been used to estimate dynamic entry and exit, capacity accumulation (e.g., [Ryan \(2012\)](#)), R&D (e.g., [Gowrisankaran and Town \(1997\)](#)), learning by doing (e.g., [Benkard \(2004\)](#)), and so on. A common assumption of those game theoretic models in structural estimation is perfect monitoring; firms' actions are perfectly observed by each other.

Games with imperfect monitoring, however, are important in industrial organization. A classic example is collusion. Rarely do firms perfectly observe each other's pricing or production behavior. The seminal paper of [Green and Porter \(1984\)](#) models collusion as a repeated Cournot competition model, where firms do not observe each other's production. They observe only

the market price. The market price is a noisy signal of firms' production choices; a low market price could mean that some firms over-produce, or that the demand is low. The model aptly captures quantity-setting collusive behavior in the real world, such as the operations of the Organization of the Petroleum Exporting Countries. Subsequent work models collusion as repeated Bertrand competition (e.g., [Harrington and Skrzypacz \(2007\)](#)), which adequately applies to price-setting collusion cases such as the Joint Executive Committee ([Porter, 1983](#)) and the vitamins cartel ([de Roos, 2001](#)). More broadly than prices and quantities, in any repeated interaction among firms, not all aspects of a firm's strategy are perfectly observed; does Samsung perfectly know how much Apple invests in research and development? Does Pfizer perfectly know how much Novartis spends in engaging with physicians?

In this paper, I introduce repeated games with imperfect public monitoring to structural industrial organization, and explore partial identification of model primitives. I use incentive compatibility constraints from the repeated game, in which continuation payoffs come from limit (sequential) equilibrium payoff sets, to deliver bounds on the structural parameters. When the Folk Theorem with imperfect public monitoring holds, the limit sequential equilibrium payoff set is the feasible and individually rational payoff set ([Fudenberg et al., 1994](#)). Otherwise, under the assumption that firms play a perfect public equilibrium (PPE), the limit PPE payoff set can be computed using [Fudenberg and Levine \(1994\)](#). The assumption that firms play a PPE is without loss of generality if firms play pure strategies, as any pure-strategy sequential equilibrium is equivalent to a PPE.

This approach has two merits. First, structural assumptions are minimal. I allow firms' strategies to depend in arbitrary ways on the history, rather than only on a state variable as in dynamic games; I allow firms in different markets to play different equilibria, instead of restricting them to playing same equilibrium as in structural analysis of dynamic games (see, for example, [Bajari et al. \(2007\)](#)); I allow the game to have multiple equilibria, without the need to worry about the equilibrium multiplicity problem; I allow firms to play mixed strategies, without the restriction to pure strategies. Of course, the cost of imposing minimal structural assumptions is the failure to point-identify structural parameters. And yet it would

be interesting to see how much identifying power those minimal structural assumptions can have. Second, computation is light. It is a trivial exercise to derive the feasible and individually rational payoff set, a linear programming problem to computing the limit PPE payoff set, and a straightforward nonlinear programming problem to solve for the identified set.

The main idea of partial identification here is as follows. First, I decompose the repeated-game payoffs into the static payoffs and the continuation payoffs. This reduces the repeated game into a static game, conditional on the public history. Second, for all action profiles observed in that reduced game, I lay out the incentive compatibility constraints for the firms to play each of those. Third, I examine whether a given value of the structural parameters satisfies the incentive compatibility constraints with some continuation payoffs from the limit equilibrium payoff set under an assumed discount factor. The limit equilibrium payoff set bounds all continuation payoffs under trivial regularity conditions regardless of the public history. There are two ways where the structural parameters enter the incentive compatibility constraints: via the static payoffs, and the limit equilibrium payoff set.

The identification technique in this paper applies to empirical settings where firms observe a noisy public signal of others' behavior, actions and signals are discretizable, and the environment is stable. The first requirement justifies imperfect public monitoring as the information structure. The second facilitates the computation of the limit equilibrium payoff set. The third permits a representation by repeated games, without a payoff-relevant state variable. In settings where the researcher does not know the signal distribution, or does not perfectly observe firms' actions (despite court-mandated disclosure or surveillance by market intelligence firms), I propose a simulation-based approach to identification. In those cases, the identification result will be a distribution of identified sets. Potential empirical settings include repeated oligopolistic competition in quantities, prices, or other strategic margins such as sales efforts. An important example of structural parameters is the marginal cost, which is helpful in evaluating firm conduct and detecting collusion.

The most relevant paper to this is [Abito \(2015\)](#), who also studies partial identification in repeated games. There are two main differences. First, [Abito \(2015\)](#) focuses on games with

perfect monitoring, which is a special case of imperfect public monitoring. Second, I bound the continuation payoffs by the limit PPE payoff set when the Folk Theorem fails, which is a strict subset of the feasible and individually rational payoff set, while [Abito \(2015\)](#) simply uses the latter by invoking the Folk Theorem. Indeed, in [Section 3.3](#), I present two examples, one where the Folk Theorem obtains and the other where it does not. In the latter case, computing the limit PPE payoff set, rather than simply taking the feasible and individually rational payoff set, yields more informative bounds on the structural parameters.

[Fudenberg et al. \(1994\)](#) establishes sufficient conditions for the Folk Theorem in repeated games with imperfect public monitoring to obtain. When the sufficient conditions fail, the literature characterizing the limit PPE payoff set in repeated games with imperfect public monitoring starts with [Fudenberg and Levine \(1994\)](#). They propose a linear programming procedure based on the recursive structure of the payoff set discovered by [Dilip Abreu, David Pearce \(1990\)](#). This paper uses the [Fudenberg et al. \(1994\)](#); [Fudenberg and Levine \(1994\)](#) to bound the continuation payoffs. Subsequent literature proposes procedures to derive the limit PPE payoff set when the mild regularity condition in [Fudenberg and Levine \(1994\)](#) fails ([Fudenberg et al., 2007](#)), and the limit equilibrium payoff sets in more general games. For example, [Hörner et al. \(2011a, 2014\)](#) study the limit PPE payoff set for stochastic games with imperfect public monitoring. [Hörner et al. \(2011b\)](#) extends [Hörner et al. \(2011a\)](#) to measurable state, action, and signal spaces. [Yamamoto \(2017\)](#) studies the limit equilibrium payoff set for stochastic games with hidden states.

The rest of this paper proceeds as follows. [Section 3.2](#) sets up the model and illustrates the main idea for identification. In [Section 3.3](#), I first provide two examples of the model, one where the Folk Theorem obtains and the other where it does not. I then compute the identified sets, assuming that the researcher perfectly knows the signal distribution and the firms' actions. [Section 3.4](#) examines the implications of relaxing those two assumptions by simulation exercises. [Section 3.5](#) discusses future work.

3.2 Model

3.2.1 Set-up

Time is discrete, $t = 1, 2, \dots, \infty$. There are two firms, $i = 1, 2$. In each period t , each firm i simultaneously chooses action a_{it} from a finite set A_i . Firms do not observe each other's action, but rather a public signal y_t from a finite set Y . The (stationary) distribution of the public signal, denoted by $\pi(\cdot|a_t)$, depends on the profile of actions $a_t = (a_{1t}, a_{2t})$.

Denote by $r_i(a_{it}, y_t)$ firm i 's payoff to action a_{it} after the public signal realizes as y_t . firm i 's static payoff to an action profile a_t , expected before the public signals realizes, is:

$$u_i(a_t) = \sum_{y_t \in Y} \Pr(y_t|a_t) r_i(a_{it}, y_t)$$

Thus, one firm's static payoff depends on the other's action only through the distribution of the public signal.

Firms are interested in maximizing the average discounted value of static payoffs using the common discount factor $\delta \in (0, 1)$. Define the *public history at the beginning of period t* as:

$$h_t^{pub} = (y_1, y_2, \dots, y_{t-1}).$$

Note that the public history does not include the current public signal, which has not realized by the beginning of period t . The *private history for firm i at the beginning of period t* is:

$$h_{it} = (a_{i1}, a_{i2}, \dots, a_{i,t-1}).$$

A *strategy for firm i* , σ_i , is a sequence of maps mapping the public and i 's private histories to i 's the set of probability distributions on A_i , denoted by ΔA_i .

The Folk Theorem for repeated games with imperfect public monitoring, or Theorem 6.2 in [Fudenberg et al. \(1994\)](#), says that any normalized (by $1 - \delta$) repeated-game payoff vector in any smooth subset in the interior of the feasible and individually rational payoff set can be supported as a (sequential) equilibrium payoff vector for large enough discount factors. A payoff vector is feasible if it lies in the convex hull of the static payoff vectors. A payoff vector

is individually rational if it gives each firm at least its min-max static payoff.

When the Folk Theorem fails, some payoff vector in some smooth set in the interior of the feasible and individually rational payoff set cannot be a sequential equilibrium payoff vector. In that case, while characterizing the limit sequential equilibrium payoff set remains an open question, the literature has characterized (and computed) the limit equilibrium payoff set of a more restrictive equilibrium concept, the perfect public equilibrium (PPE). In such an equilibrium, firms' strategies depend only on what firms commonly observe, the public history. A *public strategy for firm i* , σ_i^p , is a sequence of maps mapping the public history to ΔA_i . A *perfect public equilibrium* is a profile of public strategies such that the ongoing play, after any public history, is a Nash equilibrium. That is, for each i , each t , each public history h_t^{pub} , and each alternative strategy σ_i' , we have:

$$U_i(\sigma_i^y, \sigma_j^y | h_t^{pub}) \geq U_i((\sigma_i)'\sigma_j^y | h_t^{pub})$$

where

$$U_i(\sigma^p | h_t^{pub}) = \sum_{\tau=t}^{\infty} \delta^{\tau-t} E_{\sigma^p} [u_i(a_\tau) | h_t^{pub}]$$

and the expectation is taken over random public signals. Note that a PPE is a sequential equilibrium, which in turn is a Nash equilibrium.

The restriction to public strategies is without loss of generality when firms play only pure strategies; any pure-strategy Nash equilibrium is equivalent to a Nash equilibrium in public (and pure) strategies. Furthermore, the collusive equilibria constructed in many papers, starting with [Green and Porter \(1984\)](#) and [Abreu et al. \(1986\)](#), are perfect public equilibria.

3.2.2 Idea of Identification

Suppose that we have data on the public history h_T^{pub} and the private histories (h_{1T}, h_{2T}) for some T . Note that this implicitly assumes that we (as econometricians) observe the private histories perfectly. Assume also for now that we know the public signal distribution $\Pr(\cdot | a)$. In [Section 3.4](#), we will explore the implications of unknown signal distributions and mis-measured private histories.

By inter-temporal incentive compatibility, we have in each period $t \leq T$, for each firm i , for all $a'_{it} \in A_i$, given the public and private histories $\bar{h}_t \equiv (h_t^{pub}, h_{1t}, h_{2t})$:

$$\begin{aligned} & u_i(a_{it}, a_j(\bar{h}_t)) + \frac{\delta}{1-\delta} \sum_{y_t \in Y} \Pr(y_t | a_{it}, a_j(\bar{h}_t)) w_i(\bar{h}_t, y_t) \\ & \geq u_i(a'_{it}, a_j(\bar{h}_t)) + \frac{\delta}{1-\delta} \sum_{y_t \in Y} \Pr(y_t | a'_{it}, a_j(\bar{h}_t)) w_i(\bar{h}_t, y_t) \end{aligned}$$

where $a_j(\bar{h}_t)$ is Firm j 's (possibly mixed) action prescribed by the sequential equilibrium given \bar{h}_t , and w_i is firm i 's continuation payoff at the beginning of $t+1$. Note that we have normalized the continuation payoffs to be on the same scale as static payoffs by the scalar $\frac{1}{1-\delta}$.

To illustrate the idea for identification, assume that $A_1 = A_2 = \{a^L, a^H\}$, and that we observe in data that only (a^H, a^H) and (a^L, a^L) are played. This means that those action profiles are incentive compatible for both firms given the respective public histories. If (a^H, a^H) is played given history \bar{h}_t , and (a^L, a^L) given $\bar{h}_{t'}$, we will have the following constraints:

$$\begin{aligned} & u_1(a^H, a^H) + \frac{\delta}{1-\delta} \sum_{y_t \in Y} \Pr(y_t | a^H, a^H) w_1(\bar{h}_t, y_t) \\ & \geq u_1(a^L, a^H) + \frac{\delta}{1-\delta} \sum_{y_t \in Y} \Pr(y_t | a^L, a^H) w_1(\bar{h}_t, y_t) \end{aligned} \tag{3.1}$$

$$\begin{aligned} & u_2(a^H, a^H) + \frac{\delta}{1-\delta} \sum_{y_t \in Y} \Pr(y_t | a^H, a^H) w_2(\bar{h}_t, y_t) \\ & \geq u_2(a^H, a^L) + \frac{\delta}{1-\delta} \sum_{y_t \in Y} \Pr(y_t | a^H, a^L) w_2(\bar{h}_t, y_t) \end{aligned} \tag{3.2}$$

and

$$\begin{aligned} & u_1(a^L, a^L) + \frac{\delta}{1-\delta} \sum_{y_{t'} \in Y} \Pr(y_{t'} | a^L, a^L) w_1(\bar{h}_{t'}, y_{t'}) \\ & \geq u_1(a^H, a^L) + \frac{\delta}{1-\delta} \sum_{y_{t'} \in Y} \Pr(y_{t'} | a^H, a^L) w_1(\bar{h}_{t'}, y_{t'}) \end{aligned} \tag{3.3}$$

$$\begin{aligned} & u_2(a^L, a^L) + \frac{\delta}{1-\delta} \sum_{y_{t'} \in Y} \Pr(y_{t'} | a^L, a^L) w_2(\bar{h}_{t'}, y_{t'}) \\ & \geq u_2(a^L, a^H) + \frac{\delta}{1-\delta} \sum_{y_{t'} \in Y} \Pr(y_{t'} | a^L, a^H) w_2(\bar{h}_{t'}, y_{t'}) \end{aligned} \tag{3.4}$$

Because we do not know the equilibrium strategies, neither do we know the continuation

payoffs, w_i 's. However, the continuation payoffs always come from the limit equilibrium payoff set, regardless of the history. We then leverage this bound in the constraints (3.1)-(3.4) to obtain bounds on the structural parameters θ .

If we have data from multiple markets, we do not need to restrict the plays across different markets to come from the same equilibrium. This is because we do not require equilibrium strategies for identification. This is in contrast to [Bajari et al. \(2007\)](#) and [Pakes et al. \(2007\)](#), both requiring the restriction that different markets share the same equilibrium.

3.3 Examples

In this section, I present two examples of the model described in the previous section to illustrate how to compute the bound on the continuation payoffs and, subsequently, the identified set.

When the Folk Theorem holds, we use the feasible and individually rational payoff set (as a function of the structural parameter θ), $V^*(\theta)$, to bound the continuation payoffs in *any* sequential equilibrium. When the Folk Theorem fails, we use the limit PPE payoff set, $Q(\theta) \equiv \lim_{\delta \rightarrow 1} E(\theta, \delta)$, to bound the continuation payoffs in *any* PPE, where $E(\theta, \delta)$ be the PPE payoff set. As described in the previous section, the assumption that firms play a PPE is without loss of generality if firms play pure strategies; any sequential equilibrium in pure strategies is equivalent to a PPE. Note that we focus on the “limit” set, instead of the PPE payoff set at a given discount factor δ , because there is no general result on characterizing the latter, and by Theorem 3.1 in [Fudenberg and Levine \(1994\)](#), the latter is a subset of $Q(\theta)$.

3.3.1 Example 1: The Folk Theorem Obtains

Two firms produce substitute goods. They face an exogenous, fixed market price of 1. Each firm can exert costly sales effort, c , to affect the demand for its product stochastically. Let 1 denote the exertion of sales effort and 0 otherwise. Each public signal is a profile of realized demands, from the set $\{LL, LH, HL, HH\}$, where we let $L = 0, H = 1$. Assume that the marginal production cost is zero.

One example of this setup is two pharmaceutical companies producing substitute medications. The market price of the medications is fixed by the insurance companies. Each firm can go to physicians to conduct promotional activities, which are not observed by the opponent firm. Sales, or demands, however, are publicly observed. The marginal cost of producing the medications is almost zero.

If neither firm promotes, low demands for both will be quite likely, but high demand for either or both will be possible and symmetric. If only one firm promotes, then high demand for that firm and low demand for the other will be quite likely. If both promote, then high demands for both will be quite likely. The signal distribution is as follows:

$$\begin{aligned}
\Pr(LL|0,0) &= 0.4, & \Pr(LH|0,0) &= 0.2, & \Pr(HL|0,0) &= 0.2, & \Pr(HH|0,0) &= 0.2, \\
\Pr(LL|1,0) &= 0.3, & \Pr(LH|1,0) &= 0.05, & \Pr(HL|1,0) &= 0.4, & \Pr(HH|1,0) &= 0.25, \\
\Pr(LL|0,1) &= 0.3, & \Pr(LH|0,1) &= 0.4, & \Pr(HL|0,1) &= 0.05, & \Pr(HH|0,1) &= 0.25, \\
\Pr(LL|1,1) &= 0.2, & \Pr(LH|1,1) &= 0.25, & \Pr(HL|1,1) &= 0.25, & \Pr(HH|1,1) &= 0.3,
\end{aligned}$$

which yields the following static payoffs:

$$\begin{aligned}
u_1(0,0) &= u_2(0,0) = 0.4, \\
u_1(1,0) &= u_2(0,1) = 0.65 - c, \\
u_1(0,1) &= u_2(1,0) = 0.3, \\
u_1(1,1) &= u_2(1,1) = 0.55 - c.
\end{aligned}$$

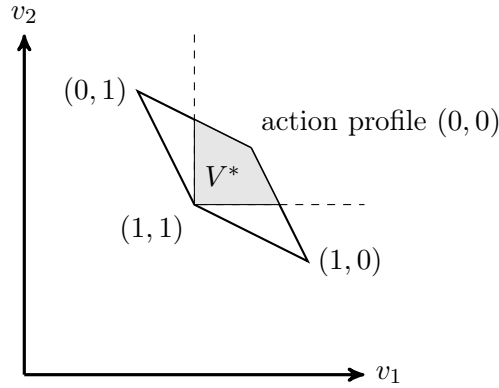
where c is the structure parameter to be estimated.

Table 3.1: *The stage game where firms produce substitute goods.*

		Firm 2	
		0	1
Firm 1	0	0.4, 0.4	0.3, 0.65 - c
	1	0.65 - c, 0.3	0.55 - c, 0.55 - c

Promotional effort is very costly. Thus, it is Pareto efficient for neither firm to promote.

Figure 3.1: *The feasible and individually rational payoff set, V^* , where firms produce substitute goods.*



This implies that $0.4 > 0.55 - c$, or $c > 0.15$. However, whatever the opponent firm's behavior, either firm has a strong incentive to promote to reap demand; the profile where both firms promote is the unique static Nash equilibrium. Thus, $0.4 < 0.65 - c$ and $0.3 < 0.55 - c$, which imply $c < 0.25$. We henceforth assume $c \in (0.15, 0.25)$.

Figure 3.1 plots the feasible and individually rational payoff set, V^* ; the dashed lines correspond to minimax payoffs for firms, both obtained at $(1,1)$. This set V^* is an upper bound on the sequential equilibrium payoffs at any discount factor; when the Folk Theorem obtains, it is the limit sequential equilibrium payoff set.

Does the Folk Theorem obtain? If it does, then we can simply take the feasible and individually rational payoff set V^* as the bound on the continuation payoffs. Theorem 6.2 in Fudenberg et al. (1994) gives sufficient conditions for Folk Theorem to obtain: if

1. any pure action profile has individual full rank;
2. either of the following holds:
 - (a) every pure-action, Pareto-efficient profile is pairwise-identifiable for all pairs of players,
 - (b) for all pairs of players, there exists a profile of (possibly mixed) actions that has pairwise full rank for that pair;

and

3. the feasible and individually rational payoff set has the same dimensionality as the number of players,

then any interior feasible and individually rational static payoff vector can arise as a perfect public equilibrium for large enough discount factors.

Let's now verify if the sufficient conditions above are satisfied. For Condition 1, there are four pure action profiles: $(0, 0)$, $(0, 1)$, $(1, 0)$, $(1, 1)$. All pure action profiles have individual full rank. For example, for $(0, 0)$, we have for Firm 1:

$$\Pi_1(0) = \begin{bmatrix} 0.4 & 0.2 & 0.2 & 0.2 \\ 0.3 & 0.05 & 0.4 & 0.25 \end{bmatrix}$$

whose rank is two, equal to the number of actions that Firm 1 has. This means that whatever the other firm's pure action, a firm's choice of promotion versus no promotion can be statistically distinguished.

For Condition 2(a), given the Pareto-efficient profile is $(0, 0)$, we have:

$$\Pi(0, 0) = \begin{bmatrix} \Pi_1(0) \\ \Pi_2(0) \end{bmatrix} = \begin{bmatrix} 0.4 & 0.2 & 0.2 & 0.2 \\ 0.3 & 0.05 & 0.4 & 0.25 \\ 0.4 & 0.2 & 0.2 & 0.2 \\ 0.3 & 0.4 & 0.05 & 0.25 \end{bmatrix}$$

whose rank is three, equal to the sum of ranks of $\Pi_1(q^L)$ and $\Pi_2(q^L)$ minus one. Therefore, Condition 2(a) holds. Intuitively, from $(0, 0)$, one firm's deviation to promotion can be statistically distinguished from the other's firm's deviation, because the deviations produce different probability distributions over the demand profile.

For Condition 3, we know from Figure [3.1](#) that the dimension of V^* is two, equal to the number of firms. Therefore, the Folk Theorem obtains.

3.3.2 Example 2: The Folk Theorem Does Not Obtain

When the Folk Theorem may not obtain, we could still use V^* to bound the sequential equilibrium payoffs. The point of computing $Q(c)$, which may be a strict subset of V^* , is to derive tighter bounds on the structural parameter. Below we lay out another example, where the Folk Theorem may not obtain. We shall then use [Fudenberg and Levine \(1994\)](#) to compute $Q(c)$, possibly a strict subject of V^* .

This example has the two firms above produce complement goods. Thus, what previously was two demands now collapses into one demand; each public signal a realized demand, from the set $\{L, M, H\}$, where we let $L = 1, M = 2, H = 3$.

One example of this setup is two pharmaceutical companies producing complement medications. Each firm can go to physicians to conduct promotional activities, and physicians prescribe the two medications jointly.

If neither firm promotes, low demand will be quite likely. If only one firm promotes, both firms will benefit in that high demand for both drugs will be more likely. If both promote, then high demands for both will be quite likely. The signal distribution is as follows:

$$\begin{aligned} \Pr(L|0,0) &= 0.6, & \Pr(M|0,0) &= 0.3, & \Pr(H|0,0) &= 0.1, \\ \Pr(L|1,0) &= 0.4, & \Pr(M|1,0) &= 0.2, & \Pr(H|1,0) &= 0.4, \\ \Pr(L|0,1) &= 0.4, & \Pr(M|0,1) &= 0.2, & \Pr(H|0,1) &= 0.4, \\ \Pr(L|1,1) &= 0.2, & \Pr(M|1,1) &= 0.1, & \Pr(H|1,1) &= 0.7, \end{aligned}$$

which yields the following static payoffs:

$$\begin{aligned} u_1(0,0) &= u_2(0,0) = 0.6 * 1 + 0.3 * 2 + 0.1 * 3 = 1.5 \\ u_1(1,0) &= u_2(0,1) = 0.4 * 1 + 0.2 * 2 + 0.4 * 3 - c = 2 - c, \\ u_1(0,1) &= u_2(1,0) = 0.4 * 1 + 0.2 * 2 + 0.4 * 3 = 2, \\ u_1(1,1) &= u_2(1,1) = 0.2 * 1 + 0.1 * 2 + 0.7 * 3 - c = 2.5 - c. \end{aligned}$$

Table 3.2: *The stage game where firms produce complement goods.*

		Firm 2	
		0	1
Firm 1	0	1.5, 1.5	2, 2 - c
	1	2 - c, 2	2.5 - c, 2.5 - c

Promotional efforts are a public good, rather than a defection as is in the previous example. Thus, it is Pareto efficient for both firms to promote. This implies that $2.5 - c > 1.5$, or $c < 1$. However, whatever the opponent firm's behavior, either firm has a strong incentive not to promote, or "shirk", to enjoy the positive spillover or not spending effort; the profile where neither firm promotes is the unique static Nash equilibrium. Thus, $2.5 - c < 2$ and $2 - c < 1.5$, which imply $c > 0.5$. We henceforth assume $c \in (0.5, 1)$.


We shall see that the Folk Theorem may not obtain, because neither Condition 2(a) nor 2(b) is met. For Condition 2(a), given the Pareto-efficient profile is $(1, 1)$, we have:

$$\Pi(1, 1) = \begin{bmatrix} \Pi_1(1) \\ \Pi_2(1) \end{bmatrix} = \begin{bmatrix} 0.4 & 0.2 & 0.4 \\ 0.2 & 0.1 & 0.7 \\ 0.4 & 0.2 & 0.4 \\ 0.2 & 0.1 & 0.7 \end{bmatrix}$$

whose rank is two, less than three, the sum of ranks of $\Pi_1(q^L)$ and $\Pi_2(q^L)$ minus one. Therefore, Condition 2(a) does not hold. Intuitively, from $(1, 1)$, one firm's deviation to no promotion cannot be statistically distinguished from the other's firm's deviation, because the deviations produce the same probability distribution over the demand.

Now we check Condition 2(b). For a general mixed-action profile $((x, 1 - x), (y, 1 - y))$

Figure 3.2: *The feasible and individually rational payoff set, V^* , when firms produce complement goods.*



where x is the probability Firm 1 puts on no promotion and y Firm 2:

$$\begin{aligned} \Pi((x, 1-x), (y, 1-y)) &= \begin{bmatrix} \Pi_1((y, 1-y)) \\ \Pi_2((x, 1-x)) \end{bmatrix} \\ &= \begin{bmatrix} 0.6y + 0.4(1-y) & 0.3y + 0.2(1-y) & \dots \\ 0.4y + 0.2(1-y) & 0.1y + 0.1(1-y) & \dots \\ 0.6x + 0.4(1-x) & 0.3x + 0.2(1-x) & \dots \\ 0.4x + 0.2(1-x) & 0.2x + 0.1(1-x) & \dots \end{bmatrix} \end{aligned}$$

whose rank is two, because the first column is always twice the second column, and after eliminating one, the number of columns is two. This is smaller than three, the sum of the number of actions that each firm has minus one. Hence, Condition 2(b) fails.

We now use [Fudenberg and Levine \(1994\)](#) to compute $Q(c)$, possibly a strict subset of V^* . To start, [Figure 3.2](#) plots the feasible and individually rational payoff set, V^* .

This game shares a feature with the partnership game in [Fudenberg and Levine \(1994\)](#),

which, by their Lemma 4.1, implies that it suffices to consider only the pure action profiles when computing $Q(c)$. In particular, whether firm i prefers promotion or not depends only on the specification of the continuation payoffs and is independent of the other firm's action. Indeed, given whatever the other firm's action, a firm always saves a static payoff of $(1 - \delta)(c - 0.5)$ (recall that $c > 0.5$) by moving from promotion to no promotion, but in doing so it unilaterally increases the probability of lower demands and incurs a cost of $\delta(0.2w_i(L) + 0.1w_i(M) - 0.3w_i(H))$. This means that firm i would not promote if and only if

$$(1 - \delta)(c - 0.5) > \delta(0.2w_i(L) + 0.1w_i(M) - 0.3w_i(H)).$$


We consider four classes of directions:

1. For directions $(-1, 0)$ and $(0, -1)$, Lemma 3.2 in [Fudenberg and Levine \(1994\)](#) shows that the maximum score in each direction is 1.5, obtained at the minimax profile $(0, 0)$.
2. For directions (λ_1, λ_2) where $\lambda_1, \lambda_2 < 0$, the maximum score cannot be smaller than $1.5(\lambda_1 + \lambda_2)$, obtained at the action profile $(0, 0)$. Indeed, that profile is a static Nash equilibrium, and therefore the repeated play of that equilibrium is a perfect public equilibrium.

The directions considered so far place no restrictions on $Q(c)$ beyond $Q(c) \subseteq V^*$. We now consider the directions (λ_1, λ_2) where $\lambda_2 > 0 > \lambda_1$ (and by symmetry, the direction where $\lambda_2 < 0 < \lambda_1$), and where $\lambda_1, \lambda_2 \geq 0$.

3. For directions (λ_1, λ_2) where $\lambda_2 > 0 > \lambda_1$, the maximum score is $(2 - c)\lambda_1 + 2\lambda_2$, enforced by the action profile $(1, 0)$ and the continuation payoffs on the line passing through $(2 - c, 2)$ and perpendicular to the direction (λ_1, λ_2) . One example of the continuation

Figure 3.3: The limit perfect public equilibrium payoff set, $Q(c)$, when firms produce complement goods and have relatively low promotion cost, $c < 0.6$. The graph is generated using $c = 0.55$.



payoffs is:

$$w_1(L) = 1.9 - c, \quad w_2(L) = 2 + \frac{\lambda_1}{\lambda_2}(2 - c - w_1(L))$$

$$w_1(M) = 2 - c, \quad w_2(M) = 2$$


$$w_1(H) = \frac{1}{0.3} \left[-\frac{1-\lambda}{\lambda}(c - 0.5) + 0.2(1.9 - c) + 0.1(2 - c) \right]$$

$$w_2(H) = 2 + \frac{\lambda_1}{\lambda_2}(2 - c - w_1(H))$$

4. For directions (λ_1, λ_2) where $\lambda_1, \lambda_2 \geq 0$, we need only consider action profiles $(1, 1)$ (as $(0, 0)$ only yields 0) and $(1, 0)$ (as $(0, 1)$ is symmetric). Using the technique in [Fudenberg and Levine \(1994\)](#), the highest score achieved by the action profile $(1, 1)$ is $(\lambda_1 + \lambda_2)[2.5 - c - (2c - 1)] = (\lambda_1 + \lambda_2)[3.5 - 3c]$, and that by the action profile $(1, 0)$ is $(2 - c)\lambda_1 + 2\lambda_2$. Therefore, the profile $(1, 1)$ generates a higher score in the first quadrant if and only if $3.5 - 3c > 2 - 0.5c$, that is $c < 0.6$.

Figures [3.3](#) and [3.4](#) plot $Q(c)$ with $c < 0.6$ and $c \geq 0.6$, respectively.

Figure 3.4: The limit perfect public equilibrium payoff set, $Q(c)$, when firms produce complement goods and have relatively low promotion cost, $c \geq 0.6$. The graph is generated using $c = 0.7$.



3.3.3 Discussion

Whether or not the Folk Theorem holds is generically independent of the structural parameter. Indeed, the key is the distribution of the public signal, which we can read off from any (simulated) dataset. Of course, if the value of the structural parameter leads to (the non-generic case of) a V^* with a smaller dimension than the number of firms, then the Folk Theorem may not obtain. Then, we can use [Fudenberg et al. \(2007\)](#) to derive a tighter bound on the PPE payoffs than $Q(c)$.

Another case where V^* has a smaller dimension than the number of firms is when we restrict the strategies to be symmetric. Then, the continuation payoffs following a given public history are restricted to be the same for both firms. As a result, the Folk Theorem may not obtain. Indeed, [Fudenberg et al. \(1994\)](#) show that this restriction causes the [Green and Porter \(1984\)](#) and [Abreu et al. \(1986\)](#) models to be bounded away from efficiency. For the identification results to go through in this paper, we do not need to impose this symmetry assumption. But if researchers do decide to impose this assumption for intuitive reasons in some contexts, as before we can use [Fudenberg et al. \(2007\)](#) to derive a tighter bound on the

PPE payoffs than $Q(c)$.

We have allowed firms to use mixed strategies in this paper. If, as with standard practice in empirical industrial organization, we allow only pure strategies, then we need only to check pure action profiles when we compute $Q(c)$. Thus, the pure strategy restriction replaces the linear structure restriction used in the analysis above in justifying our checking only pure action profiles.

3.3.4 The Identified Set

To illustrate the identification programming with the two examples above, for now we assume that the researcher knows the signal distribution and observes the private histories without measurement errors.

The identification programming first finds the minimum and the maximum of the structural parameter c , by searching over the continuation payoffs in the limit equilibrium payoff set that respects the incentive compatibility constraints implied by observed play. The structural parameter may be subject to some a priori constraint too; for example, in the first example, for “no promotion” to be Pareto efficient and “both promotion” to be the unique Nash equilibrium, $c \in (0.15, 0.25)$. This is a nonlinear optimization problem. Then, having obtained the extreme points of the structural parameter, we argue that any value in between is achievable.

Identification in Example 1. We first elaborate on the first step of the identification programming using Example 1. We present the nonlinear optimization problem, and supply three ways to approach the problem: manual inspection, Matlab, and AMPL. Suppose we see only two action profiles being played, $(0, 0)$ and $(1, 1)$. The incentive compatibility constraints

are:

$$\begin{aligned}
u_1(0,0) + \frac{\delta}{1-\delta} \sum_{y \in Y} \Pr(y_t|0,0)w_1(\bar{h}_t, y) &\geq u_1(1,0) + \frac{\delta}{1-\delta} \sum_{y \in Y} \Pr(y_t|1,0)w_1(\bar{h}_t, y), \\
u_2(0,0) + \frac{\delta}{1-\delta} \sum_{y_t \in Y} \Pr(y_t|0,0)w_2(\bar{h}_t, y) &\geq u_2(0,1) + \frac{\delta}{1-\delta} \sum_{y_t \in Y} \Pr(y_t|0,1)w_2(\bar{h}_t, y), \\
u_1(1,1) + \frac{\delta}{1-\delta} \sum_{y \in Y} \Pr(y|1,1)w_1(\bar{h}_{t'}, y) &\geq u_1(0,1) + \frac{\delta}{1-\delta} \sum_{y \in Y} \Pr(y|0,1)w_1(\bar{h}_{t'}, y), \\
u_2(1,1) + \frac{\delta}{1-\delta} \sum_{y \in Y} \Pr(y|1,1)w_2(\bar{h}_{t'}, y) &\geq u_2(1,0) + \frac{\delta}{1-\delta} \sum_{y \in Y} \Pr(y|1,0)w_2(\bar{h}_{t'}, y).
\end{aligned}$$

The first two say that at a history \bar{h}_t where $(0,0)$ is observed, given the static payoffs u and the continuation payoffs $w(\bar{h}_t, \cdot)$, neither firm wants to unilaterally deviate to promotion. The second two say that at a history $\bar{h}_{t'}$ where $(1,1)$ is observed, given the static payoffs u and the continuation payoffs $w(\bar{h}_{t'}, \cdot)$, neither firm wants to unilaterally deviate to no promotion. Rearranging, we have:

$$u_1(1,0) - u_1(0,0) \leq \frac{\delta}{1-\delta} \sum_{y \in Y} [\Pr(y|0,0) - \Pr(y|1,0)]w_1(\bar{h}_t, y), \quad (3.5)$$

$$u_2(0,1) - u_2(0,0) \leq \frac{\delta}{1-\delta} \sum_{y_t \in Y} [\Pr(y_t|0,0) - \Pr(y_t|0,1)]w_2(\bar{h}_t, y), \quad (3.6)$$

$$u_1(1,1) - u_1(0,1) \geq \frac{\delta}{1-\delta} \sum_{y \in Y} [\Pr(y|0,1) - \Pr(y|1,1)]w_1(\bar{h}_{t'}, y), \quad (3.7)$$

$$u_2(1,1) - u_2(1,0) \geq \frac{\delta}{1-\delta} \sum_{y \in Y} [\Pr(y|1,0) - \Pr(y|1,1)]w_2(\bar{h}_{t'}, y). \quad (3.8)$$

The first two say that the gain from unilaterally deviating to promotion must be dwarfed by the expected long-term loss through the change in the signal distribution (and in turn the future plays). The second two say that the loss from unilaterally deviating to no promotion must overwhelm the expected long-term gain through the change in the signal distribution (and in turn the future plays).

The continuation payoffs w 's come from the limit equilibrium payoff set regardless of history. From Figure [3.1](#), we know that the limit equilibrium payoff set, which is the feasible

and individually rational payoff set, contains (w_1, w_2) such that:

$$w_i \geq 0.55 - c, \quad i = 1, 2, \quad (3.9)$$

$$w_2 \leq w_1 \frac{0.25 - c}{-0.1} + 0.4 - 0.4 \frac{0.25 - c}{-0.1}, \quad (3.10)$$

$$w_2 \leq w_1 \frac{-0.1}{0.25 - c} + 0.4 - 0.4 \frac{-0.1}{0.25 - c}. \quad (3.11)$$

The a priori constraint on c is:

$$c \in [0.15, 0.25] \quad (3.12)$$

The nonlinear optimization problem is to optimize c subject to the constraints (3.5)-(3.12).

For a manual approach to this nonlinear problem, we could ask if $c = 0.154$ is between the lower and upper bounds on c , given $\delta = 0.98$. To select w , note that since $c = 0.154$ is relatively low, the static gains from deviation to promotion is relatively high. Therefore, we need an even higher long-term punishment to prevent this deviation. Since deviation reduces the probability of observing L relative to H at the deviating firm, to make the loss big, we need a high continuation payoff to L and a low payoff to H . An intuitively appealing choice for w is:

$$w(LL) = (0.4, 0.4)$$

$$w(LH) = (\bar{v}, 0.55 - c)$$

$$w(HL) = (0.55 - c, \bar{v})$$

$$w(HH) = (0.55 - c, 0.55 - c)$$

where \bar{v} is such that $w(LH)$ and $w(HL)$ are on the boundary of $Q(c)$:

$$\bar{v} = (0.15 - c) \frac{c - 0.25}{0.1} + 0.4.$$

Therefore, constraints (3.9)-(3.11) are satisfied.

But this choice of w 's does not satisfy all incentive compatibility constraints. Substituting the w 's into constraints (3.5), we have $0.0960 \leq 0.0772$, a contradiction. Therefore, unless we can find some other w 's that meet all the constraints, $c = 0.154$ is not in the identified set

given $\delta = 0.98$.

We implement this nonlinear optimization problem using Matlab's `fmincon`, and AMPL's `knitro`. Both confirm that $c = 0.154$ is not in the identified set; indeed, the identified set given $\delta = 0.98$ is $[0.1549, 0.25]$. AMPL's `knitro` performs better than Matlab's `fmincon`, which tends to be trapped in local optima as the minimizer is sensitive to the a priori constraint on c .¹

The identified set $[0.1549, 0.25]$ is slightly narrower than the a priori constraint $[0.15, 0.25]$. This is because we remain agnostic about firms' equilibrium strategy, but uses only the assumption that firms are playing *some* sequential equilibrium. Indeed, when we choose the continuation payoffs w 's, we know that firms can use *some* strategy to achieve them in the limit, but we do not seek to figure out what those strategies are.

The identified set will be narrower, if we assume a lower discount factor (because when firms care less about the future, it is harder to vary the discounted future payoffs to support a given c), observe more action profiles (which introduces more incentive compatibility constraints), or the folk theorem fails (which leads to a smaller limit equilibrium payoff set than the feasible and individually rational payoff set). We will explore the role of measurement errors and unknown signal distributions in the next section.

We have thus far shown that the minimum and the maximum of c are supported by continuation payoffs that meet all the constraints. The second step of the identification programming is to show that any c strictly in between, that is any $c \in (0.1549, 0.25)$, can be supported as well. We give an informal argument. On one hand, increasing c reduces the static gain from unilaterally deviating to no promotion, $(0.25 - c)$. This makes the collusive profile $(0, 0)$ easier to support. On the other hand, increasing c does not affect the attainability of the defection profile $(1, 1)$, because the latter is the unique Nash equilibrium for any c in its a priori range $(0.15, 0.25)$. Therefore, if a small c such as $c = 0.1549$ can be supported, so can larger c 's.

Identification in Example 2. Suppose we see only two action profiles being played, $(0, 0)$

¹Matlab's `fmincon` returns 0.1549 as the minimizer only when c is constrained between 0.15 and 0.22.

Table 3.3: Identified set of c at different discount factors, when only action profiles $(0, 0)$ and $(1, 1)$ are observed, with the known signal distribution and no measurement error.

Discount factor	Lower bound	Upper bound
0.99	0.5	0.9685
0.98	0.5	0.9401
0.95	0.5	0.8701
0.9	0.5	0.7872
0.8	0.5	0.6875
0.5	0.5	0.5957

and $(1, 1)$. Recall that Example 2 is a partnership game, so $(1, 1)$ is the Pareto efficient profile, while $(0, 0)$ is the inefficient, unique Nash equilibrium. A larger c reduces the incentive to work, so if $(1, 1)$ is observed, c cannot be too large.

Because different ranges of c lead to different $Q(c)$, the nonlinear optimization problem consists of two sub-problems. The first sub-problem looks for the minimum and maximum of c in the a priori range of $[0.5, 0.6]$ by choosing w 's from the shaded area in Figure 3.3 that satisfy incentive compatibility constraints. The second sub-problem looks for the minimum and maximum of c in the a priori range of $[0.6, 1.0]$ by choosing w 's from the shaded area in Figure 3.4 that satisfy incentive compatibility constraints.

Table 3.3 reports the identified set at different discount factors δ . When firms care less about the future, it is harder to rationalize a high c when we see both firms exert efforts.

3.4 Simulation

This section explores the implications of unobserved signal distributions and mis-measured private histories for the identified set.

I generate the data as follows. The game is given by Example 1, with $c = 0.238$ and $\delta = 0.99$. I construct a sequential equilibrium in the style of Green and Porter (1984). In this equilibrium, both firms start out with no promotion. They stick to no promotion until they observe signals other than LL , following which they promote for 15 periods, and then

switch back to no promotion. According to this equilibrium, we generate a T -period play of this game, $(h_T^{pub}, h_{1,T}, h_{2,T})$, where h_T^{pub} is the public history of signals, $h_{i,T}$ is firm i 's private history of actions, and $T = 10000$. Each element in $h_{i,T}$ may be subject to measurement errors; with a known small probability $m = 0.01$, each action is mis-measured as the other action, independently across periods and firms. The data with measurement errors is $(h_T^{pub}, \tilde{h}_{1,T}, \tilde{h}_{2,T})$.

3.4.1 Unknown Signal Distribution Only

Suppose that we do not know the true signal distribution, but do observe the true play, $(h_T^{pub}, h_{1,T}, h_{2,T})$.

There are three cases to consider. First, with infinite data ($T = \infty$) and all action profiles on the equilibrium path, we can simply derive the signal distribution from the conditional probabilities of signals under each action profile. The derived signal distribution is the true signal distribution.

Second, with infinite data but some action profiles off the equilibrium path, we cannot derive the signal distribution conditional on the action profile(s) not observed. For example, in the equilibrium constructed above, only two action profiles are observed on the equilibrium path, one where both promote, and the other where neither promotes. Then, even with infinite data, we cannot derive the signal distribution conditional on the action profiles where one and only one firm promotes. In such cases, we need to rely on contextual knowledge to obtain or assume those missing probabilities or scenarios of them.

Third, with finite data, even when all action profiles have realized, the conditional frequencies of signals only approximate the true signal distribution. We would then derive the sampling distribution of those signal frequencies, and simulate the sampling distribution. For each sampled signal distribution, we obtain an identified set. The identification result would then be a distribution of identified sets.

Below I derive the sampling distribution of the conditional frequencies. Conditional on a

particular action profile a that has realized.² let X_i denote a draw of the public signal:

$$X_i = \begin{bmatrix} 1(y_i = LL|a) \\ 1(y_i = HL|a) \\ 1(y_i = LH|a), \end{bmatrix}$$

where the remaining signal, HH , is omitted. For notational convenience, I henceforth omit the conditioning on a . Let there be n such draws. Denote $E(X_i) = \beta$, so each row of β is the probability that the public signal takes the corresponding value, except HH , for which the probability is implied by the other three. Denote $\frac{1}{n} \sum_{i=1}^n X_i = b$, the conditional frequency estimator. We would like to get the asymptotic distribution of:

$$\begin{aligned} \sqrt{n}(b - \beta) &= \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n X_i - \beta \right) \\ &= \frac{1}{\sqrt{n}} \sum_{i=1}^n (X_i - \beta) \end{aligned}$$

Let $G_i = X_i - \beta$. Since $E(G_i) = E(X_i - \beta) = \beta - \beta = 0$, and

$$\begin{aligned} Cov(G_i) &= \begin{bmatrix} Cov(1(y_i = LL), 1(y_i = LL)) & Cov(1(y_i = LL), 1(y_i = HL)) & \dots & \dots \\ Cov(1(y_i = HL), 1(y_i = LL)) & Cov(1(y_i = HL), 1(y_i = HL)) & \dots & \dots \\ \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots \end{bmatrix} \\ &= \begin{bmatrix} \beta_1(1 - \beta_1) & -\beta_1\beta_2 & -\beta_1\beta_3 \\ -\beta_1\beta_2 & \beta_2(1 - \beta_2) & -\beta_2\beta_3 \\ -\beta_1\beta_3 & -\beta_2\beta_3 & \beta_3(1 - \beta_3) \end{bmatrix}, \end{aligned}$$

and G_i 's are independently distributed, by the Central Limit Theorem:

$$\sqrt{n}(b - \beta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n G_i \xrightarrow{d} \mathcal{N}(0, Cov(G_i)).$$

By the Law of Large Numbers and the Slutsky Theorem, a consistent estimator of $Cov(G_i)$ is

²For those action profiles that have not realized in the data, I use the true public signal distribution associated with them.

the frequency analog based on b . So we can randomly draw from $\mathcal{N}(0, \hat{Cov}(G_i))$, divide the draws by \sqrt{n} , and subtract those from b to obtain draws of the true signal probabilities, β .³

The signal distribution affects three inputs to the identification programming:

- whether the Folk Theorem holds, through the sufficient conditions in [Fudenberg and Levine \(1994\)](#). Since the Folk Theorem holds under generic signal distributions, if it holds under the true signal distribution (as it does in Example 1), so will it under the estimated signal distribution unless the latter happens to be knife-edge. In any case, using V^* even when the Folk Theorem fails is still valid despite being conservative. For these reasons, I use the V^* to bound the continuation payoffs in this simulation exercise.
- the feasible and individually rational payoff set V^* , through the static payoff matrix. To preserve the “prisoners’ dilemma” feature of Example 1, I put a priori constraints on c so that $(0, 0)$ achieves Pareto efficiency, $(1, 1)$ is the unique static Nash equilibrium and the min-max profile.
- the incentive compatibility constraints, through the probabilities of public signals as they relate to continuation payoffs, and through the static payoffs.

Table [3.4](#) shows the frequency that the true promotion cost, $c = 0.238$, is within the lower and upper bounds identified, under various discount factors and numbers of signal distributions drawn. The coverage frequency increases, and the bounds widen, with the discount factor. The identification result is fairly insensitive to the number of signal distributions drawn.

3.4.2 Mis-measured Private Histories Only

Suppose that we know the true signal distribution, but observe the mis-measured play, $(h_T^{pub}, \tilde{h}_{1,T}, \tilde{h}_{2,T})$.

As a motivation for measurement errors in the private histories, it is intuitively unappealing that the researcher perfectly knows the chosen actions that firms themselves imperfectly observe.

³We need to make sure the implied probability of the fourth signal, HH , is positive. This is true in all simulations I have implemented.

Table 3.4: Identification result for $c = 0.238$ at various discount factors and number of signal distributions drawn.

Discount Factor	# of Signal Distributions Drawn	Coverage Frequency of $c = 0.238$	Median of Lower Bounds	Median of Upper Bounds
0.99 (true)	100	0.61	0.1572	0.2406
	500	0.58	0.1586	0.2395
0.95	100	0.53	0.1763	0.2398
	500	0.54	0.1771	0.2388
0.90	100	0.30	0.2124	0.2384
	500	0.32	0.2067	0.2384

The true discount factor is 0.99, and the length of play in the data is 10000 periods.

Even if the researcher could obtain, after a lag, possibly more accurate information on the chosen actions, due to disclosure mandated by courts, or surveillance by market intelligence firms, that information may still contain measurement errors. Second, measurement errors are a natural way to introduces randomness into the data, while the usual way of achieving this, by private payoff shocks, leads to a different monitoring structure, the private monitoring. A general characterization of the limit equilibrium payoff set in repeated games with private monitoring is an open question.

From the mis-measured private histories, we could simulate many paths of private histories, repeat the identification exercise for each path, and obtain a distribution of identified sets. However, the distribution most likely degenerates to a point mass on one identified set. The reason is that our identification comes from the *existence* of action profiles in the data, and simulated private histories most likely have the same existence profile as the mis-measured private history. Thus, when the mis-measured and simulated private histories share the existence profiles, even though one can compute the likelihood of the simulated private histories by looking at how close the implied signal distribution is to the true, known signal distribution, that likelihood adds no additional information to the identified set.

The distribution of the identified set is a point mass on $[0.1525, 0.2500]$ with a discount factor of 0.99, $[0.1621, 0.2500]$ with a discount factor of 0.95, and $[0.1733, 0.2500]$ with a discount factor of 0.90, all covering the true $c = 0.238$ with probability one.

3.4.3 Both Unobserved Signal Distribution and Mis-measured Private Histories

Suppose that we neither know the true signal distribution nor observe the true play.

The simulation-based identification programming goes as follows. First, I simulate the private histories from the measurement error distribution. Second, for each private history drawn, I derive the signal frequencies, and draw signal distributions from the sampling distribution of the derived signal frequencies. Third, for each private history and each signal distribution drawn, I obtain an identified set. The identification result is a distribution of identified sets.

This procedure works well when, absent measurement errors, all action profiles are on the equilibrium path. This ensures that the signal frequencies conditional on any action profile are not an artifact of mere measurement errors, but contain useful information of the true signal distribution. Thus, in Example 1, where only $(0, 0)$ and $(1, 1)$ are on the equilibrium path when measurement errors are absent, any observation of $(0, 1)$ or $(1, 0)$ in the mis-measured private histories can only be due to measurement errors, and therefore the signal frequencies conditional on $(0, 1)$ and $(1, 0)$ are induced by measurement errors and have little relation to the true signal distribution.

Indeed, while the true signal distribution in Example 1, conditional on $(1, 0)$ or $(0, 1)$, is:

$$\begin{aligned} \Pr(LL|1, 0) &= 0.3, & \Pr(LH|1, 0) &= 0.05, & \Pr(HL|1, 0) &= 0.4, & \Pr(HH|1, 0) &= 0.25, \\ \Pr(LL|0, 1) &= 0.3, & \Pr(LH|0, 1) &= 0.4, & \Pr(HL|0, 1) &= 0.05, & \Pr(HH|0, 1) &= 0.25, \end{aligned}$$

the signal frequencies derived from the mis-measured private histories are:

$$\begin{aligned} \hat{\Pr}(LL|1, 0) &= 0.237, & \hat{\Pr}(LH|1, 0) &= 0.204, & \hat{\Pr}(HL|1, 0) &= 0.290, & \hat{\Pr}(HH|1, 0) &= 0.269, \\ \hat{\Pr}(LL|0, 1) &= 0.223, & \hat{\Pr}(LH|0, 1) &= 0.214, & \hat{\Pr}(HL|0, 1) &= 0.205, & \hat{\Pr}(HH|0, 1) &= 0.351. \end{aligned}$$

Under these signal frequencies, the desired structure of the static payoffs (e.g. Pareto efficiency of $(0, 0)$, the unique static Nash equilibrium and the minimax profile being $(1, 1)$) cannot be satisfied, leading to an empty identified set for the promotional cost c .

We now generate the data using an alternative equilibrium where all action profiles are on the equilibrium path. In this equilibrium, both firms start out with no promotion. They stick to no promotion until they observe signals other than LL . There are three punishment phases. If the signal is LH , play $(1, 0)$ until signal turns HL for the first time, at which point they revert to $(0, 0)$. If the signal is HL , play $(0, 1)$ until signal turns LH for the first time, at which point they revert to $(0, 0)$. If the signal is HH , play $(1, 1)$ for T period, and then revert to $(0, 0)$.

Let V_0, V_1, V_2, V_3 denote the value of being in the collusive phase and the punishment phases. We have:

$$V_0 = u(0, 0) + \delta(\Pr(LL|0, 0)V_0 + \Pr(LH|0, 0)V_1 + \Pr(HL|0, 0)V_2 + \Pr(HH|0, 0)V_3)$$

$$V_1 = u(1, 0) + \delta[(1 - \Pr(HL|1, 0))V_1 + \Pr(HL|1, 0)V_0]$$

$$V_2 = u(0, 1) + \delta[(1 - \Pr(LH|0, 1))V_2 + \Pr(LH|0, 1)V_0]$$

$$V_3 = u(1, 1) + \delta u(1, 1) + \dots + \delta^{T-1}u(1, 1) + \delta^T V_0$$

The incentive compatibility constraints for either firm are:

$$V_0 \geq u(1, 0) + \delta(\Pr(LL|1, 0)V_0 + \Pr(LH|1, 0)V_1 + \Pr(HL|1, 0)V_2 + \Pr(HH|1, 0)V_3)$$

$$V_1 \geq u_{00} + \delta[(1 - \Pr(HL|0, 0))V_1 + \Pr(HL|0, 0)V_0]$$

$$V_2 \geq u(1, 1) + \delta[(1 - \Pr(LH|1, 1))V_2 + \Pr(LH|1, 1)V_0]$$

Those constraints ensure that neither firm wants to deviate in the collusive phase and the first two punishment phases, respectively. No firm wants to deviate in the third punishment phase because the associated punishment profile, $(1, 1)$, is a static Nash equilibrium.

The parameter values that sustain such an equilibrium are $c = 0.20, T = 3, \delta = 0.99$, and the signal probability conditional on $(1, 1)$ changed to:

$$\Pr(LL|1, 1) = 0.24, \quad \Pr(LH|1, 1) = 0.23, \quad \Pr(HL|1, 1) = 0.23, \quad \Pr(HH|1, 1) = 0.3$$

Table [3.5](#) shows the frequency that the true promotion cost, $c = 0.20$, is within the lower

Table 3.5: *Identification result for $c = 0.20$ at various discount factors.*

Discount Factor	Coverage Frequency of $c = 0.20$	Median of Lower Bounds	Median of Upper Bounds
0.99 (true)	0.81	0.1298	0.2091
0.95	0.65	0.1623	0.2079
0.90	0.31	0.1936	0.2057

The number of private histories drawn is 10, and the number of signal distributions drawn per private history is 100. The true discount factor is 0.99, and the length of play in the data is 10000 periods.

and upper bounds identified, under various discount factors. Again, the coverage frequency increases, and the bounds widen, with the discount factor.

3.5 Future Work

I plan to apply the technique developed in this paper to an empirical setting to study firm conduct and detect collusion. Future work also includes extending partial identification to dynamic (as opposed to repeated) games, including those with imperfect monitoring (Hörner et al., 2011a,b, 2014) and those with hidden states (Yamamoto, 2017).

References

- Abito, J. M. (2015). How Much Can We Identify from Repeated Games?
- Abreu, D., Pearce, D., and Stacchetti, E. (1986). Optimal Cartel Equilibria with Imperfect Monitoring. *Journal of Economic Theory*, 39:251–269.
- Abul Naga, R. H. and Jones, P. (2012). Better to Give Than to Receive? Altruistic Provision of a Global Public Good. *Economics Letters*, 115(3):480–483.
- Aldy, J. E., Chen, C., and Pizer, W. A. (2017). Evaluating the Paris Agreement Mitigation Pledges with Statistical Forecasting Models. *Harvard Kennedy School Working Paper*.
- Arce M., D. G. and Sandler, T. (2001). Transnational Public Goods: Strategies and Institutions. *European Journal of Political Economy*, 17:493–516.
- Athey, S. and Levin, J. (2001). Information and Competition in U.S. Forest Service Timber Auctions. *Journal of Political Economy*, 109(2):375–417.
- Auffhammer, M. and Steinhauser, R. (2012). Forecasting the Path of U.S. CO2 Emissions Using State Level Information. *The Review of Economics and Statistics*, 94(1):172–185.
- 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.
- Barbu, A.-D. (2015). Trends and Projections in Europe 2015 — Tracking Progress Towards Europe’s Climate and Energy Targets. *EEA Report No 4/2015*.
- Barrett, S. (1990). The Problem of Global Environmental Protection. *Oxford Review of Economic Policy*, 6(1):68–79.
- Barrett, S. (1994). Self-Enforcing International Environmental Agreements. *Oxford Economic Papers*, 46:878–894.
- Benkard, C. L. (2004). A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft. *Review of Economic Studies*, 71:581–611.

- Bergstrom, T. (1973). On Efficient Provision of Social Overhead Goods. *Yearbook of East European Economics*, 4.
- 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.
- Berry, S. T. (1992). Estimation of a Model of Entry in the Airline Industry. *Econometrica*, 60(4):889–917.
- Boadway, R. and Hayashi, M. (1999). Country Size and the Voluntary Provision of International Public Goods. *European Journal of Political Economy*, 15(4):619–638.
- Boadway, R., Song, Z., and Tremblay, J. F. (2011). The Efficiency of Voluntary Pollution Abatement When Countries Can Commit. *European Journal of Political Economy*, 27(2):352–368.
- Bray, R. L. (2017a). Markov Decision Processes with Exogenous Variables.
- Bray, R. L. (2017b). Strong Convergence and Dynamic Economic Models.
- Bresnahan, T. F. and Reiss, P. C. (1990). Entry in Monopoly Markets. *Review of Economic Studies*, 57(4):531–553.
- Buchholz, W., Cornes, R., and Peters, W. (2006). Lindahl vs. Cournot Nash: The Role of the Income Distribution. *Public Finance Analysis*, 62:28–49.
- Buchholz, W., Cornes, R., and Peters, W. (2008). Existence, Uniqueness and Some Comparative Statics for Ratio and Lindahl Equilibria. *Journal of Economics*, 95(2):167–177.
- Buchholz, W., Cornes, R., and Rübbelke, D. (2011). Interior Matching Equilibria in a Public Good Economy: An Aggregative Game Approach. *Journal of Public Economics*, 95:639–645.
- Buchholz, W., Cornes, R., and Rübbelke, D. (2014). Potentially Harmful International Cooperation on Global Public Good Provision. *Economica*, 81(322):205–223.
- Byrd, R. H., Nocedal, J., and Waltz, R. A. (2006). Knitro: An Integrated Package for Nonlinear Optimization.
- Callander, S. and Harstad, B. (2015). Experimentation in Federal Systems. *Quarterly Journal of Economics*, 130(2):951–1002.
- Carlson, C., Burtraw, D., Cropper, M., and Palmer, K. (2000). Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade? *Journal of Political Economy*, 108(6):1292–1326.
- Cason, T. N. (1993). Seller Incentive Properties of EPA’s Emission Trading Auction. *Journal of Environmental Economics and Management*, 25(2):177–195.
- Cason, T. N. (1995). An Experimental Investigation of the Seller Incentives in the EPA’s Emission Trading Auction. *American Economics Review*, 85(4):905–922.

- Cason, T. N. and Plott, C. R. (1996). EPA's New Emissions Trading Mechanism: A Laboratory Evaluation. *Journal of Environmental Economics and Management*, 30:133–160.
- Center for International Earth Science Information Network (CIESIN)/Columbia University (2013). Low Elevation Coastal Zone: Urban-Rural Population and Land Area Estimates, Version 2.
- 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.
- Ciliberto, F. and Tamer, E. (2009). Market Structure and Multiple Equilibria in Airline Markets. *Econometrica*, 77(6):1791–1828.
- Cornes, R. and Hartley, R. (2007). Aggregative Public Good Games. *Journal of Public Economic Theory*, 9(2):201–219.
- Cornes, R. and Sandler, T. (1984). The Theory of Public Goods: Non-Nash Behaviour. *Journal of Public Economics*, 23(3):367–379.
- 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.
- de Roos, N. (2001). Collusion with a Competitive Fringe: An Application to Vitamin C. (202):1–30.
- Dilip Abreu, David Pearce, E. S. (1990). Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica*, 58(5):1041–1063.
- 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.
- Energy Research Institute (2009). *2050 China Energy and CO2 Emissions Report (CEACER)*. Beijing: Science Press, 2009.
- European Commission (2017). Paris Agreement.
- Falkinger, J., Hackl, F., and Pruckner, G. J. (1996). A Fair Mechanism for Efficient Reduction of Global CO-Emissions. *Public Finance Analysis*, 53:308–331.
- 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. *The American Economic Review*, 100(3):837–869.
- Fowle, M., Reguant, M., and Ryan, S. P. (2016). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy*, 124(1):249–302.
- Fudenberg, D. and Levine, D. K. (1994). Efficiency and Observability with Long-Run and Short-Run Players. *Journal of Economic Theory*, 62:103–135.
- Fudenberg, D., Levine, D. K., and Maskin, E. (1994). The Folk Theorem with Imperfect Public Information. *Econometrica*, 62(5):997–1039.
- Fudenberg, D., Levine, D. K., and Takahashi, S. (2007). Perfect Public Equilibrium When Players Are Patient. *Games and Economic Behavior*, 61(1):27–49.
- Genesove, D. and Mullin, W. P. (1998). Testing Static Oligopoly Models: Conduct and Cost in the Sugar Industry, 1890-1914. *The RAND Journal of Economics*, 29(2):355–377.
- Gerber, A. and Wichardt, P. C. (2009). Providing Public Goods in the Absence of Strong Institutions. *Journal of Public Economics*, 93:429–439.
- Goldfarb, A. and Xiao, M. (2011). Who Thinks about the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets. *The American Economic Review*, 101(7):3130–3161.
- Gowrisankaran, G. and Town, R. J. (1997). Dynamic Equilibrium in the Hospital Industry. *Journal of Economics & Management Strategy*, 6(1):45–74.
- Green, E. J. and Porter, R. H. (1984). Noncooperative Collusion under Imperfect Price Information. *Econometrica*, 52(1):87–100.
- Griffin, J. M. and Xiong, W. (1997). The Incentive to Cheat: An Empirical Analysis of OPEC. *The Journal of Law & Economics*, 40(2):289–316.
- Grubb, M., Sha, F., Spencer, T., Hughes, N., Zhang, Z., and Agnolucci, P. (2015). A review of Chinese CO2 Emission Projections to 2030: The Role of Economic Structure and Policy. *Climate Policy*, 15:7–39.
- Hahn, R. W. and May, C. A. (1994). The Behavior of the Allowance Market: Theory and Evidence.
- Hansen, L. P. (2007). Beliefs, Doubts and Learning : Valuing Macroeconomic Risk. *American Economic Review*, 97(2):1–30.
- Harrington, J. E. and Skrzypacz, A. (2007). Collusion under Monitoring of Sales. *RAND Journal of Economics*, 38(2):314–331.
- Heitzig, J., Lessmann, K., and Zou, Y. (2011). Self-Enforcing Strategies to Deter Free-Riding in the Climate Change Mitigation Game and Other Repeated Public Good Games. *Proceedings of the National Academy of Sciences*, 108(38):15739–15744.

- Hennigan, W. and Bennett, B. (2016). U.S. to Urge Partners to Do More to Fight Islamic State amid Complaints from Pentagon. *Los Angeles Times*.
- Hohne, N., Blok, K., Jeffery, L., , and Rocha, M. (2015). Paris Agreement: Near-Term Actions Do Not Match Long Term Purpose, But State Is Set to Ramp up Climate Action. *Climate Action Tracker Statement*.
- Holtz-Eakin, D. and Selden, T. M. (1995). Stoking the fires? CO 2 Emissions and Economic Growth. *Journal of public economics*, 57:85–101.
- Horn, H. and Wolinsky, A. (1988). Bilateral Monopolies and Incentives for Merger. *RAND Journal of Economics*, 19(3):408–419.
- Hörner, J., Sugaya, T., Takahashi, S., and Vieille, N. (2011a). Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \rightarrow 1$ and a Folk Theorem. *Econometrica*, 79(4):1277–1318.
- Hörner, J., Takahashi, S., and Vieille, N. (2011b). Recursive Methods in Discounted Stochastic Games II: Infinite State Space. pages 1–8.
- Hörner, J., Takahashi, S., and Vieille, N. (2014). On the Limit Perfect Public Equilibrium Payoff Set in Repeated and Stochastic Games. *Games and Economic Behavior*, 85(1):70–83.
- Hortaçsu, A., Puller, S. L., Luco, F., and Zhu, D. (2016). Strategic Ability and Productive Efficiency in Electricity Markets.
- Isaac, R. M. and Walker, J. M. (1988). Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism. *Quarterly Journal of Economics*, 103:179–199.
- Joskow, P. L. (1972). *A Behavioral Theory of Public Utility Regulation*. PhD thesis.
- Kahn, M. E. and Matsusaka, J. (1997). Demand for Environmental Goods: Evidence from Voting Patterns on California Initiatives. *Journal of Law and Economics*, 40:137–174.
- 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.
- Kolstad, C. D. (2007). Systematic Uncertainty in Self-Enforcing International Environmental Agreements. *Journal of Environmental Economics and Management*, 53(1):68–79.
- 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*.
- Lindahl, E. (1958). Just Taxation: A Positive Solution. In *Classics in the Theory of Public Finance*.
- 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.
- Nash, J. F. (1950). The Bargaining Problem. *Econometrica*, 18(2):155–162.
- Nordhaus, W. (2015). Climate Clubs: Overcoming Free-riding in International Climate Policy. *The American Economic Review*, 105(4):1339–1370.
- Olson, M. (1965). *The Logic of Collective Action: Public Goods and the Theory of Groups*. Schocken Books, New York.
- Olson, M. J. and Zeckhauser, R. (1966). An Economic Theory of Alliances. *The Review of Economics and Statistics*, 48(3):266–279.
- Pakes, A., Ostrovsky, M., and Berry, S. (2007). Simple Estimators for the Parameters of Discrete Dynamic Games (With Entry/Exit Examples). *RAND Journal of Economics*, 38(2):373–399.
- Porter, R. H. (1983). A Study of Cartel Stability: The Joint Executive Committee, 1880-1886. *The Bell Journal of Economics*, 14(2):301–314.
- 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 Strategies, Costs, and Emission Allowances.
- Ryan, S. P. (2012). The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica*, 80(3):1019–1061.
- Schelling, T. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge.
- 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.
- Schmalensee, R., Stoker, T. M., and Judson, R. A. (1998). World Carbon Dioxide Emissions: 1950-2050. *The Review of Economics and Statistics*, 80(1):15–27.
- Shitovitz, B. and Spiegel, M. (1998). Cournot Nash and Lindahl Equilibria in Pure Public Good Economies. *Journal of Economic Theory*, 83:1–18.
- Sugden, R. (1985). Consistent Conjectures and Voluntary Contributions to Public Goods: Why the Conventional Theory Does Not Work. *Journal of Public Economics*, 27(1):117–124.
- United States Department of Agriculture (2015). The International Macroeconomic Data Set.
- United States Department of State (2015). The United States Climate Action Report 2014.
- U.S. Environmental Protection Agency (1994). National Air Pollutant Emission Trends.

- Warr, P. G. (1983). The Private Provision of a Public Good is Independent of the Distribution of Income. *Economics Letters*, 13:207–211.
- World Bank (2016). World Development Indicators.
- World Resources Institute (2015). CAIT Climate Data Explorers.
- World Values Survey Association (2016). *World Values Survey Wave 6 (2010-2014), Official Aggregate v.20150418*. Aggregate File Producer: Asep/JDS, Madrid Spain.
- Yamamoto, Y. (2017). Stochastic Games with Hidden States.

Appendix A

Appendix to Chapter 1

A.1 Data Compilation

A.1.1 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. I multiple units share a scrubber, I allocate the operating cost based on each unit's sulfur input (that is, heat input multiplied by sulfur content), and the installed cost based on the capacity of the generator associated with each unit.

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

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

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

A.1.2 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.1.3 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.

A.2 Additional Data Patterns

This section presents two more data patterns that inform the structural model, in addition to those already provided in Section [1.3](#).

Volumes and vintages of allowance transactions. While utilities can trade allowances of any vintage, almost all transactions during the period of analysis concern the current-year, the prior-year, and the next-year vintages. Figure [A.1](#) 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 [1.4](#) thus focuses on allowance trading in those vintages.

Figure [A.2](#) plots the distribution of current- and prior-vintage net allowance purchases for utility-years during the compliance periods up to 2003. Figure [A.3](#) plots the counterpart for next-vintage net allowance purchases since one year before the compliance periods up to 2002.




Figure A.1: *Distribution of volume share of current-, prior-, and next-vintage allowance transactions in any-vintage transactions during the compliance periods, utility-year, 1995-2003.*

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 [A.4](#) plots the distributions of the delivered prices of coal and natural gas reported to Form FERC-423 during the period of analysis. The lower quartile gas price always exceeded the upper quartile coal price, in most years by a lot. The added cost of dispatching coal from its sulfur emissions, given the relatively low allowance price, is unlikely to switch the dispatch order of coal and gas. The boom of shale gas production did not start until around 2007, which is four years after the end of the period of analysis; Figure [A.5](#) shows that the massive drop in the natural gas price did not start until 2008.

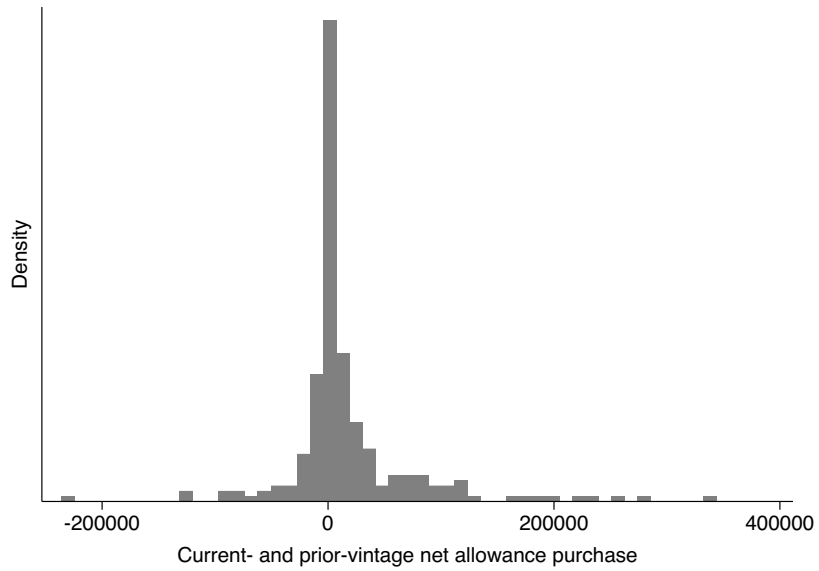


Figure A.2: *Distribution of current- and prior-vintage net allowance purchases during the compliance periods, utility-year, 1995-2003.*

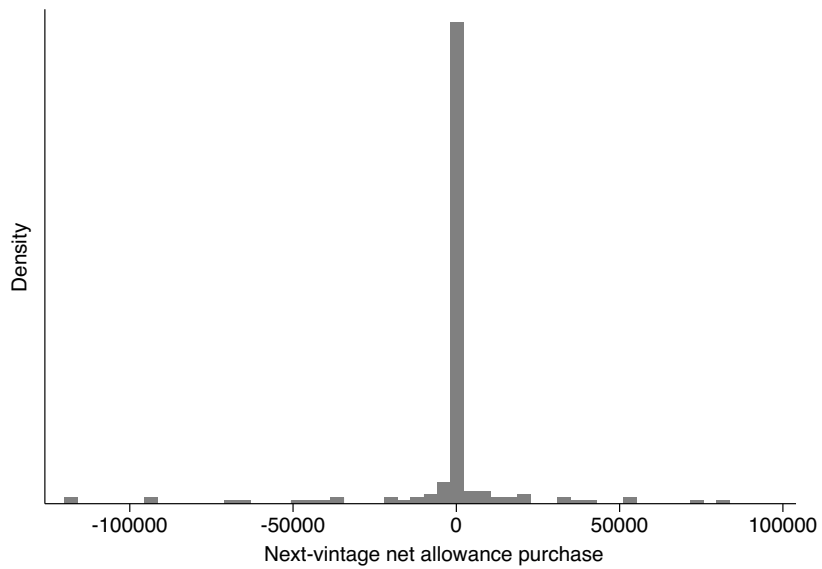


Figure A.3: *Distribution of next-vintage net allowance purchases since one year before the compliance periods, utility-year, 1994-2002.*

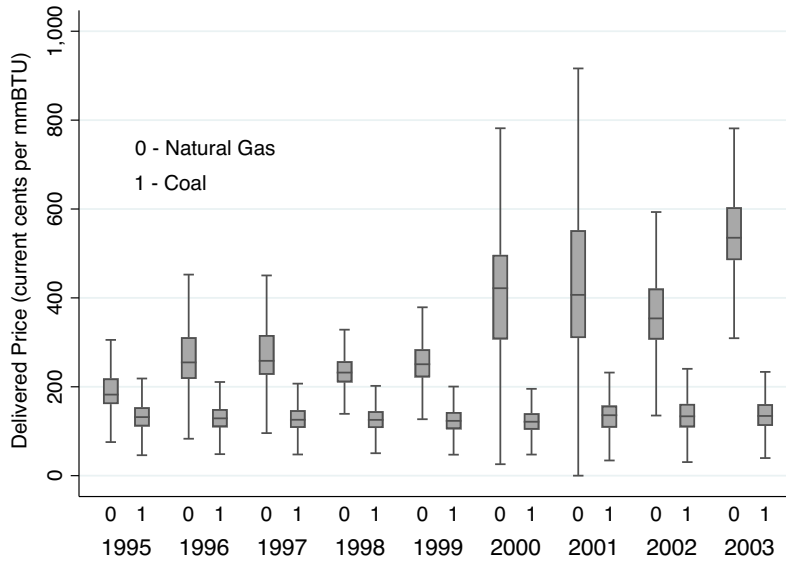


Figure A.4: *Delivered price of coal and natural gas, 1995-2003, in current U.S. cents per mmBTU. Calculated from the FERC-423 data.*




Figure A.5: *Natural gas price, 1990-2013, in 2010 U.S. dollars per thousand cubic feet. From [Davis \(2015\)](#).*

A.3 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 [1.4](#) are sufficient for optimality.

The Phase II value function is:

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

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, $\frac{\partial V_i(W_i, P, H_i)}{\partial W_i}$ 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.

A.4 Inner Loop Details

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

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

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

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

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

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

$$T_0(s) = 1,$$

$$T_1(s) = s,$$

$$T_d(s) = 2sT_{d-1}(s) - T_{d-2}(s), \quad d \geq 2.$$

To obtain $\text{coef}(\cdot, \cdot, \cdot)$ 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:

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

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 (A.1) 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:

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

where N_w is the degree of Gauss-Hermite quadrature, $\omega(w)$ is the the Gauss-Hermite weight at degree w , and $n(w)$ is the the Gauss-Hermite node at degree w . The parameters $(\sigma^P, \gamma_0^P, \gamma_1^P)$ are the standard deviation (b_5), intercept ($b_1 + b_2(s - 1)$, or $b_1 + b_2 \times 2000$, or b'_1), and slope of the allowance price transition ($b_3 + b_4(s - 1)$, or $b_3 + b_4 \times 2000$, or b'_3) from Equation (1.1) or (1.2). The parameters $(\sigma^H, \gamma_0^H, \gamma_1^H)$ are the standard deviation, intercept, and slope of the heat input transition from Equation (1.8).

Accelerating dynamic programming using Relative and Endogenous Value Function Iterations. Both Relative and Endogenous Value Function Iterations methods leverage

the fact that what matters for behavior is the shape rather than the level of the value function. Thus, we only need to check the iterative difference in the shape, not the level, for convergence. The shape converges at least as fast as does the level, and in many cases much faster. See [Bray \(2017b,a\)](#) for formal results.

The two methods differ in how they define the shape. The Relative Value Function Iterations method uses the shape covering all states. Then, the relative value function is the full value function shifted by a constant. The Exogenous Value Function Iterations method looks at the shape specific to each exogenous state; each shape covers all endogenous states at each exogenous state. To see why the collection of exogenous-state-specific shapes is sufficient for behavior, suppose that at a particular exogenous state, the payoffs at all endogenous states are shifted by the same constant. Since the firm controls the endogenous state but not the exogenous, the relative attractiveness of choices does not change. I implement the Relative Value Function Iterations method by normalizing the value function iterate by the value at the first approximation node, and the Exogenous Value Function Iterations method by normalizing the value function iterate at each exogenous state by the value at that exogenous state and the first approximation node of the endogenous state.¹

Table [A.1](#) 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

¹[Bray \(2017b,a\)](#) 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 A.1: 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} .

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

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

$$V^{full}(s) = \pi(x^*(s), s) + \beta E[V^{full}(s')|x^*(s), s], \quad (\text{A.2})$$

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 (A.2) now becomes:

$$V^{(k)}(s) + L^{(k)} = \pi(x^*(s), s) + \beta E[V^{(k)}(s')|x^*(s), s] + \beta L^{(k)}. \quad (\text{A.3})$$

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

$$V^{(k)}(s) = \pi(x^*(s), s) + \beta E[V^{(k-1)}(s')|x^*(s), s] \quad (\text{A.4})$$

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 (A.4), and comparing with Equation (A.3), 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.

Appendix B

Appendix to Chapter 2

B.1 Technical Details of Section 2.3

We collected all the INDCs that had been submitted to the UNFCCC submission portal¹ by December 5, 2015, a total of 158 INDCs. Despite the cut off date for our data collection effort being earlier than the Paris Conference, we only miss 3 INDCs. Of the 158 INDCs, 23% are percentage reductions from historical emissions levels, 44% are percentage reductions from BAU emission forecasts, a couple involve reductions on a per capita or per dollar of GDP basis, and the rest do not include any specific numbers in their submitted pledge. We focus on unconditional reductions, as opposed to conditional reductions, to be consistent with the voluntary nature of our global public good provision model. Some nations did not submit their INDCs.

We use carbon emissions despite the fact that many pledges are in terms of GHG emissions. The reason is that the historical data on individual nations' GHG emissions are very limited. We hence trade off the match with the pledges for the accuracy of forecasts, and assume that the reduction in carbon emissions will be proportional to that in GHG emissions.

To convert the reduction goals in the INDCs to an absolute amount of reduction, we need to know the BAU emissions. For the three big emitters, China, the U.S., and the EU-28, there

¹Available at <http://www4.unfccc.int/submissions/indc/Submission%20Pages/submissions.aspx>

are existing analyses that assess the fine details of the emission determinants. We simply take their BAU emission forecasts. [Energy Research Institute \(2009\)](#) predicts that China’s BAU emissions will be around 12,500 million tons in 2030.² The [United States Department of State \(2015\)](#) synthesizes multiple data sources and predicts the U.S. BAU carbon emissions to be 5,705 million tons in 2025. [Barbu \(2015\)](#) predicts that the EU-28’s BAU GHG emissions in 2030 will be 27% lower than the 1990 level.

For the rest, we use nation-wise auto-regression models to forecast their BAU emissions in 2030, drawing upon [Aldy et al. \(2017\)](#). We choose nation-wise auto-regression models because we find that they achieve a much smaller mean squared forecasting error of the aggregate carbon emissions on the last five years of available data, than other major carbon forecasting models, including [Holtz-Eakin and Selden \(1995\)](#), [Schmalensee et al. \(1998\)](#), and [Auffhammer and Steinhauser \(2012\)](#). Specifically, we regress the current log per capita carbon emission on the previous-year log per capita carbon emission, log population density, log per capita GDP, and a linear time trend, using data up to 2012, the last year for which the carbon emissions data are available from the CAIT Climate Data Explorer ([World Resources Institute \(2015\)](#)). We employ the estimated coefficients to forecast each nation’s 2030 carbon emission by relying on the population and GDP forecasts in [United States Department of Agriculture \(2015\)](#), and then iteratively use the prior year’s carbon emission forecasts. We then calculate the absolute amount of reduction from the 2030 BAU emissions that the pledges represent.³ The carbon emissions that we use do not include land use, land-use change, and forestry activities emissions, which is consistent with the practice of most pledges. For nations that submitted pledges with reduction goals, the reduction is calculated by taking the difference between our BAU emission estimates and the emission target indicated or implied in each individual INDC.

To nations that submitted pledges without a reduction goal, we assign a more negative

²For readers concerned about strategic over-reporting by the Chinese government, that forecast is within the range of 8,000-18,000 million tons of energy-related BAU emissions generated by 12 engineering or general-equilibrium modeling platforms reviewed in [Grubb et al. \(2015\)](#). It is also below what our auto-regression model would have predicted.

³Calculations are available from the authors.

normalized reduction (hence one that is less likely to bind) than that of the most non-binding pledge we have estimated; we treat those nations as less generous than those which state explicit reduction goals in their INDCs. For nations that did not submit a pledge at all, we assign an even more negative normalized reduction; we treat them as the least willing to abate.

We also collect data on nations' climate vulnerability, pollution level, and measures of environmental concerns. The vulnerability measures include the percentages of urban and rural population living in coastal areas where elevation is below 10 meters ([Center for International Earth Science Information Network \(CIESIN\)/Columbia University \(2013\)](#)), historical annual average temperature, and the percentage of population subject to drought, flood, and extreme temperature events. The pollution measures include energy use from fossil fuels, NO_x emission, and PM 2.5 concentration. We construct the per capita GNI by dividing GNI by population. These data, except otherwise noted, are all from [World Bank \(2016\)](#). We identify environmental concern measures by the percentage of subjects in [World Values Survey Association \(2016\)](#) who respond positively to environment-related questions on active membership in environmental organizations, importance of looking after the environment, protection of environment over economic growth, participation in environmental demonstration for the past two years, and confidence in environmental organizations. Data are available upon request.

B.2 Supply-Demand Arrangement

We provide an alternative, intuitive way of thinking about the Lindahl equilibrium, the Supply-Demand Arrangement. Each individual has an amount s /he would like to contribute, if his or her supply will be matched by a total of contributions from others n times as large. This matching ratio can also be thought of as the inverse of the price of the total contributions from others s /he demands (the price of his or her contribution is normalized to 1). The higher the matching ratio, the lower the price of total contributions from others. Under a Supply-Demand Arrangement, the vector of the prices that each individual faces is such that for each individual, the demand is exactly fulfilled by the total supply of all others.

Formally, a Supply-Demand Arrangement consists of individualized prices for the public

good provided by others, $(p_{-i}^{SD})_{i \in \mathcal{N}}$ (the superscript SD stands for Supply-Demand), a private good price, which we normalize to 1, and an allocation $(m_i^{SD})_{i \in \mathcal{N}}$. That allocation is such that for each $i \in \mathcal{N}$:

$$(m_i^{SD}, M_{-i}^{SD}) \in \arg \max_{m_i, M_{-i}} U_i(m_i + M_{-i}, m_i),$$

subject to: $p_{-i}^{SD} M_{-i} \leq m_i,$ (SD)

and the market clears:

$$M_{-i}^{SD} = \sum_{j \neq i} m_j^{SD}, \text{ for each } i.$$

The next proposition establishes the equivalence between the Lindahl equilibrium and the Supply-Demand Arrangement. A Supply-Demand Arrangement allocation is identical to a Lindahl equilibrium allocation. Given a supply-demand arrangement $((p_i^{SD})_{i \in \mathcal{N}}, (m_i^{SD})_{i \in \mathcal{N}})$, for each i , let

$$M = m_i^{SD} + M_{-i}^{SD}, \quad p_i = \frac{p_{-i}^{SD}}{1 + p_{-i}^{SD}}.$$

Then for each i , (m_i^{SD}, M) solves i 's Problem (*Lindahl*) with price p_i . The identity from Lindahl to Supply-Demand Arrangement can be similarly established.