

Three Essays in Environmental Economics and Applied Econometrics

by

Erik Paul Johnson

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Public Policy and Economics)
in The University of Michigan
2011

Doctoral Committee:

Professor Paul N. Courant, Co-Chair

Assistant Professor Ryan M. Kellogg, Co-Chair

Professor Michael R. Moore

Assistant Professor Meredith L. Fowlie, University of California, Berkeley

© Erik Paul Johnson 2011
All Rights Reserved

ACKNOWLEDGEMENTS

I am deeply grateful to the members of my dissertation committee – Paul Courant, Ryan Kellogg, Meredith Fowlie, and Michael Moore – for their teaching, mentorship, and patience throughout my time in graduate school. Jessica Goldberg and Elias Walsh have been indispensable friends throughout graduate school and have always made time to talk with me about research.

Additionally, I have received valuable suggestions about my research from Zoe McClaren, Emily Beam, Caroline Theoharides, Josh Hyman, and seminar participants at the University of Michigan. I am also grateful to the entire Gerald R. Ford School of Public Policy, especially Mary Corcoran, for creating a rich, interdisciplinary community of scholars who all care deeply about policy analysis.

My wife, Jen Sarafin, has helped and supported me every step of the way with words of encouragement and in every other imaginable way, especially this year. My parents, Steven and Candace Johnson, have been tremendously supportive. My sister, Sonja Nodland, has offered useful and timely advice throughout my time in graduate school, and she proved along the way that it is indeed possible to finish graduate school.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
LIST OF TABLES	v
LIST OF FIGURES	vii
LIST OF ABBREVIATIONS	ix
ABSTRACT	x
CHAPTER	
I. The Price Elasticity of Supply of Renewable Electricity Generation: Evidence from State Renewable Portfolio Standards	1
1.1 Introduction	1
1.2 RPS Policy Background	6
1.2.1 Restructured Electricity Markets and Renewable Energy Credits	7
1.2.2 Variation in State RPSs	9
1.3 Model	11
1.4 Empirical Strategy	15
1.4.1 Estimating Equations	16
1.4.2 Which States get RPSs?	17
1.4.3 Key Variables	21
1.5 Data	24
1.5.1 Generation Capacity Data	24
1.5.2 Electricity Sales and Production Data	25
1.5.3 Price Data	25
1.5.4 Policy Variables	26
1.5.5 Data Restrictions	28
1.5.6 Aggregation	29
1.6 Results	30
1.7 Policy Implications for RPSs as a CO ₂ Abatement Tool	32

1.8	Conclusion	37
II. Cost Effectiveness of Residential Radon Remediation with Household Mobility		
2.1	Introduction	51
2.2	Model	54
2.2.1	Effect of Radon on Health	54
2.2.2	Remediation Technology	56
2.2.3	Behavioral Model	57
2.3	Data	66
2.4	Results	66
2.4.1	Simulations without Children	67
2.4.2	Simulations with Children	70
2.5	Discussion of Policy Implications	71
2.5.1	Alternative Scenarios	73
2.5.2	Information Requirements	75
2.6	Conclusions	76
III. Monte Carlo Simulations of the Nested Fixed-Point Algorithm 98		
3.1	Introduction	98
3.2	Rust's Model	101
3.2.1	The Simple Model	101
3.2.2	The Relaxed Model	103
3.3	Data	105
3.4	Rust's Results and Replication	107
3.5	Asymptotic Results and Simulation Procedure	108
3.6	Simulation Results	110
3.6.1	"Small" Sample Results	110
3.6.2	Distributional Results	111
3.7	Conclusion	113
BIBLIOGRAPHY		128

LIST OF TABLES

Table

1.1	Environmental Impact of Common Fuels and Frequency of Inclusion in Renewable Portfolio Standards	39
1.2	Summary Statistics	40
1.3	First Stage Regression Estimates	41
1.4	Second Stage Regression Estimates	42
1.5	Cost of CO ₂ Abatement From an RPS	42
2.1	Constants used in the BEIR VI model of radon related lung cancer mortality	78
2.2	Percent of Population that Moved Residences by Age	78
2.3	Parameters of Economic Model	79
2.4	Capitalization and Prevalence of Remediation Equipment	80
2.5	Peak Age and Radon Levels for Remediation Decisions	81
2.6	Capitalization and Prevalence of Remediation Equipment when some Households have Children	82
2.7	Peak Age and Radon Levels for Remediation Decisions	83
3.1	Summary of Replacement Data (Buses where at least 1 replacement occurred)	115
3.2	Censored Data (Subsample of buses for which no replacements occurred)	116

3.3	Structural Estimates from Rust 1987 Fixed Point Dimension = 90, $\beta = 0.9999$	117
3.4	Summary Statistics of Parameter Estimates Model Assumptions Satis- fied	118
3.5	Summary Statistics of Parameter Estimates 13,000 Observations . .	119

LIST OF FIGURES

Figure

1.1	Statutory Renewable Requirements for Selected States	43
1.2	Timing of Renewable Portfolio Standard Adoption	44
1.3	Timing of Renewable Portfolio Standard Adoption by State	45
1.4	Selected Wholesale Electricity Markets	46
1.5	Capacity for Renewable Generation by Fuel	47
1.6	Electricity Capacity Market without RPS	48
1.7	Electricity Capacity Market with RPS	49
1.8	State REC Prices	50
2.1	Decision Tree for Households	84
2.2	Probability of Remediation Equipment Being Installed in a House	85
2.3	Probability of Dying before reaching the next age	86
2.4	Density of Ages in the United States, 2006-2008	87
2.5	Density of Radon Concentrations in an Unremediated Housing Stock	88
2.6	Policy Function for a Typical Household's Remediation Decision with a Life-Year Valuation of \$300,000	89
2.7	Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 with No Children	90

2.8	Policy Function for a Household’s Remediation Decision with a Life-Year Valuation of \$300,000 with 1 Child	91
2.9	Policy Function for a Household’s Remediation Decision with a Life-Year Valuation of \$300,000 with 2 Children	92
2.10	Policy Function for a Household’s Remediation Decision with a Life-Year Valuation of \$300,000 with 3 Children	93
2.11	Policy Function for a Household’s Remediation Decision with a Life-Year Valuation of \$300,000 with 4 Children	94
2.12	Policy Function for a Household’s Remediation Decision with a Life-Year Valuation of \$300,000 with 5 Children	95
2.13	Typical Policy Function for a Household’s Remediation Decision with a Life-Year Valuation of \$300,000 if There is No Capitalization of Remediation Equipment	96
2.14	Policy Function for a Household’s Remediation Decision with a Life-Year Valuation of \$300,000 if the Household Will Never Move	97
3.1	Simulations of ”Small” Data Sets: Replacement Cost Parameter	120
3.2	Simulations of ”Small” Data Sets: Cost Function Parameter	121
3.3	Simulations of ”Small” Data Sets: $P(x_{t+1} - x_t = 0)$	122
3.4	Simulations from ”Small” Data Sets: $P(x_{t+1} - x_t = 1)$	123
3.5	Simulations from Different Unobservable Distributions: Replacement Cost Parameter	124
3.6	Simulations from Different Unobservable Distributions: Cost Function Parameter	125
3.7	Simulations from Different Unobservable Distributions: $P(x_{t+1} - x_t = 0)$	126
3.8	Simulations from Different Unobservable Distributions: $P(x_{t+1} - x_t = 1)$	127

LIST OF ABBREVIATIONS

ASD	Active Slab Depressurization
BEIR	Biological Effects of Ionizing Radiation
DSIRE	Database of State Incentives for Renewables & Efficiency
EIA	Energy Information Administration
EPA	Environmental Protection Agency
ERCOT	Electric Reliability Council of Texas
ISO	Independent System Operator
MMBTU	Million British Thermal Units
MW	Megawatt
MWh	Megawatt hour
NFXP	Nested Fixed Point Algorithm
pCi/l	picoCurries per liter
PJM	Pennsylvania, New Jersey, and Maryland wholesale electricity market
PV	Photovoltaic
REC	Renewable Energy Credit or Renewable Electricity Credit
RGGI	Regional Greenhouse Gas Initiative
RPS	Renewable Portfolio Standard

ABSTRACT

Three Essays in Environmental Economics and Applied Econometrics

by

Erik Paul Johnson

Co-Chairs: Paul N. Courant and Ryan M. Kellogg

In the first chapter of this dissertation, I estimate the long-run price elasticity of supply of renewable electricity generating capacity using the variation in state mandates for renewable electricity production. These mandates require electricity providers to purchase a predetermined fraction of their electricity from renewable sources and typically increase annually in each state. Using the variation that these mandates induce in generating capacity, I use an instrumental variables strategy to estimate the price elasticity of supply. My estimate of 2.7 translates into a cost of carbon dioxide abatement of at least \$12 per ton of carbon dioxide. This cost of abatement is six times more expensive than the cost of carbon dioxide abatement from the Regional Greenhouse Gas Initiative.

The second chapter, written jointly with Paul Courant, David Mendez, and Kenneth Warner, conducts a benefit-cost analysis on the Environmental Protection Agency's guidelines for residential radon remediation. We use an agent based approach that adds important dimensions of heterogeneity to the analysis. This allows us to estimate the likely capitalization of the capital cost of remediation technology into housing prices. We find that most households are better off by not paying the annual

cost of remediation and that only the least mobile households with smokers in high radon concentration houses would undertake the capital cost of remediation. Since only a small fraction of the population values radon remediation, our model suggests that approximately 10% of the capital cost is capitalized into the resale value of the house.

The third chapter examines the finite sample and distributional properties of the nested fixed point algorithm. Starting from the basic setup described in John Rust's 1987 paper, I simulate data sets with varying sizes and distributional assumptions on the unobserved component of the model. I find that even in sample sizes of up to 8,000 observations, the nested fixed point algorithm can display finite sample bias and variances substantially larger than the theoretical asymptotic variance. This is also true with departures from distributional assumptions, with the mean square error increasing by a factor of 10 for some distributions of unobserved variables.

CHAPTER I

The Price Elasticity of Supply of Renewable Electricity Generation: Evidence from State Renewable Portfolio Standards

1.1 Introduction

Renewable energy has become a prominent policy issue at both the state and federal levels. Many states have adopted policies aimed at promoting the growth of renewable electricity within their state to decrease carbon dioxide (CO₂) emissions, most prominently through a renewable portfolio standard (RPS). An RPS is a mandate that retail electricity providers purchase a specified fraction of their electricity sales from renewable sources. A typical RPS is passed by a state legislature a few years before the first year retail providers are required to meet the standard to allow new renewable capacity to be built. RPSs usually begin with a requirement that approximately one percent of electricity be produced by renewable sources and incrementally increases over a 15-25 year period. For example, Massachusetts's RPS requires retail providers to demonstrate initially that one percent of their electricity sales come from renewable generation with the amount of required renewable electricity increasing by between one-half and one percentage points in every subsequent year. The end goal for Massachusetts's RPS occurs in 2020, when 15% of electricity

sales must come from renewable sources. (See Figure 1.1 for example timelines.) If a retail provider fails to meet its requirement in a given year, it must pay a penalty proportional to the the difference between the target and the amount of renewable electricity it purchased.

In 1997, three states had renewable portfolio standards (Iowa, Massachusetts, and Nevada) whereas by the end of 2009, 35 states had passed an RPS into law. (Figure 1.2 displays the number of states adopting RPSs in a given year and Figure 1.3 displays which states have passed RPSs.) Since the electricity sector accounts for 42% of CO₂ emissions nationally, RPSs may have the ability to substantially decrease CO₂ emissions. However, there has been little quantitative examination of the effectiveness or the cost of CO₂ abatement from RPSs, particularly accounting for heterogeneity in state policies.

This paper estimates the long-run price elasticity of supply of renewable generation capacity by using state RPS implementation schedules as an instrument for changes in the price received by renewable electricity generators. The price elasticity is an important parameter for policy makers since many states have introduced aggressive RPSs to increase the share of renewable electricity sold in their state, but policy makers are unlikely to have empirically based estimates of the cost of these policies.¹ I find that the price elasticity of renewable electricity capacity is approximately 2.7. Using my estimates of the long-run supply price elasticity, I calculate the cost of exclusively using an RPS to decrease the carbon dioxide emissions in the northeastern US. This elasticity suggests that the cost of abating an equivalent amount of CO₂ from an RPS in the northeastern US is between six and fourteen times larger than the costs of CO₂ abatement under a regional cap-and-trade program (the Regional Greenhouse Gas Initiative). I estimate the marginal cost of abatement for a 1%

¹There are some cost estimates of a federal RPS in the literature, for instance see Palmer and Burtraw [32], but these estimates come from simulation models of the electricity sector rather than empirically estimating the response to policies.

reduction in CO₂ emissions to be between \$12 and \$35 per ton of CO₂ compared with a price near \$2 for the cap-and-trade program.

To identify the long-run supply price elasticity of renewable generating capacity, I use variation from the prespecified RPS implementation schedules. The incremental changes in demand for renewable generation from the implementation schedules create an exogenous change in the demand for renewable electricity. These changes provide me with an instrument for the price renewable generators receive for electricity, allowing me to consistently estimate the elasticity of supply.

In order to correctly measure the changes in demand for renewable capacity due to RPSs, I develop a measure of the strength of the incentives created by a particular state's RPS. This measure is different than what has been used in most previous work on RPSs. Menz and Vachon [24] and Adelaja and Hailu [4] use cross-sectional data to examine the effectiveness of RPSs in promoting the development of wind generators.² However, both of these papers treat all RPSs the same by estimating the effect of RPSs on new capacity using a simple indicator for a state having an RPS. Both papers find that RPSs are correlated with a greater presence of wind generators in that state, but they cannot establish any causal link due to their cross-sectional approaches. In fact, Lyon and Yin [23] suggest that a large wind potential in a state increases the probability of that state passing an RPS, which suggests the causality may go the other direction. Powers and Yin [33] do account for much of the heterogeneity in policies and adopt a measure of the RPS requirement similar to this paper's measure. By using their preferred method of incorporating this heterogeneity, they find a significant impact of RPSs on the share of renewable generation. In another related paper, Kneifel [20] uses panel data on state renewable capacity and attempts

²There are also a few qualitative assessments of RPS policies. Wisser, Porter and Grace [43] examine many of the policy design issues associated with RPSs and identify broad principles that could be considered best practices. Langniss and Wisser [22] also do a qualitative assessment of the Texas RPS and suggest that it has likely been an effective driver of renewable generation development in Texas.

to discern which of the variety of renewable electricity policies are most effective at increasing in-state renewable electricity capacity.

The papers mentioned above, with the exception of Powers and Yin, assume that all RPSs create identical incentives for wind generators regardless of how difficult the policies are to meet. This is clearly not a valid assumption, given the heterogeneity in the difference between the state RPS statutory requirements and the amount of new renewable capacity needed to meet the RPS. For instance, the first year that Pennsylvania's RPS was implemented, the state had more than enough renewable capacity to meet the requirement; whereas the first year that Delaware's RPS was implemented enough new renewable generation had to be built to power approximately 2% of the state's electricity demand. The difference between the statutory RPS requirement and incentives for new renewable generating capacity can be seen in Figure 1.1. The light blue bars display the statutory requirement, and the dark red bars show the percent of electricity that must be generated by new sources due to the the RPS.

Importantly, and in contrast to previous work, I aggregate each state RPS to the regional level weighted by the state's consumption, since this is the level at which most RPSs create incentives for wholesale generators. RPSs create incentives for all renewable generators in the region since RPSs can be met with renewable capacity anywhere in the wholesale market. The requirement effectively makes each state's RPS policy an incremental increase in the region's RPS requirements. Without acknowledging that state RPSs are actually regional policies, the previous estimates of the impact of RPSs on renewable generation are biased toward zero since the effective control group in the differences-in-differences estimation is contaminated by neighboring states' policies.

My price elasticity estimates help to inform estimates of the excess burden of CO₂ reductions from RPSs since they are not a first-best policy. In a recent paper, Holland, Hughes, and Knittel [17] show under general conditions that policies that

govern the rate of pollution, rather than the level (CO₂ emissions per megawatt hour rather than total CO₂ emissions) cannot be efficient. An efficient (first-best) policy can be described where the price is equal to the marginal cost plus the marginal damages from the externality, as in the case of a Pigouvian tax or a cap-and-trade program.

However, one reason state politicians may prefer an RPS to a cap-and-trade program, even though it is not a first-best policy, is that it is harder for firms to avoid the requirements of an RPS than a cap-and-trade program. RPSs are a regulation that is hard for firms to avoid since they apply to the electricity sold, not produced, in a particular location.³ There is a large literature examining the extent to which firms avoid environmental regulation by moving production to other jurisdictions, typically called leakage. (See Fowlie [14] for a discussion of these issues.) Since leakage is unlikely to be a problem for an RPS but may be under a cap-and-trade program, my estimates of the excess burden from an RPS can be interpreted as the excess cost to avoid leakage.

The remainder of the paper is organized as follows. In the next section I discuss the details of RPSs, electricity markets, and the dimensions on which there is heterogeneity in RPS policies. In Section 1.3 I develop a model to ground our thinking about renewable generating capacity investment. In Section 1.4, I discuss the empirical methodology I use and the key variables. Section 1.5 describes the data I use to examine RPSs. Section 1.6 discusses my results which is followed by a discussion of the policy implications of my estimates in Section 1.7. Section 1.8 concludes.

³To the extent that renewable electricity increase electricity prices and residents and business locate in a state based on electricity prices, there will be some leakage in these policies. However, for all but the most electricity intensive industries this is likely not to be a problem.

1.2 RPS Policy Background

Renewable portfolio standards have become increasingly common over the past twenty years. The first law resembling an RPS was passed in Iowa in 1983. The Iowa law required the state's two investor-owned utilities to install a combined 105 megawatts (MW) of new renewable generating capacity. After this law, very little legislation was passed at the state level relating to the fuel mix of electricity generators until 1997 when Massachusetts passed its RPS. This was done as part of the electric utility restructuring legislation whereby electric generating capacity was separated from retail operations of electric utilities.⁴

The RPS requires retail providers to purchase an increasing fraction of their electricity from renewable generators beginning in 2003. After Massachusetts, many other New England and Mid-Atlantic states followed suit.⁵ By December 2009, 35 states had passed an RPS. Figures 1.2 and 1.3 show the number of states that passed and RPS in a given year and the spatial distribution of when RPSs were passed.

Among the 35 states that have passed RPSs in 2009 some of those states do not have a restructured electricity market. In states that do not have a restructured electricity market, the state public utility commission tends to have substantial influence over the fuel mix of the vertically integrated utilities it regulates since new capacity projects must be approved by the public utility commission and utilities are guaranteed rate of return on their new capacity. This may make a state's RPS superfluous if the state has vertically integrated utilities. This is one reason I will focus only on states that have restructured electricity markets in this paper.

Though all RPSs are passed by state legislatures, nearly every RPS is actually a regional policy since most states simply require that the renewable electricity used

⁴In a restructured electricity market electricity generators are owned by separate firms from retail providers that sell electricity to end users. This model is in contrast to a vertically integrated utilities that both sell to consumers and produce electricity from generators they own. Section 1.2.1 will discuss this in more detail.

⁵This was usually done as part of restructuring legislation or shortly afterwards.

to meet the requirement be generated within the wholesale market and wholesale electricity markets contain multiple states. Figure 1.4 shows the New England, PJM and New York wholesale markets.^{6,7} Therefore I will aggregate all of the state RPSs in a region into a single regional requirement. I restrict my attention in this paper to states with a restructured electricity market and a transparent RPS compliance mechanism. This restriction allows me to observe the price renewable generators receive for their electricity. The rest of this section will discuss restructured electricity markets and important dimensions of variation and then how RPSs work in practice.

1.2.1 Restructured Electricity Markets and Renewable Energy Credits

States with restructured electricity markets have three main types of market actors: wholesale electricity generators, retail electricity providers, and end consumers. In restructured electricity markets, firms typically can only be a wholesale generator (and therefore own generating capacity) *or* a retail provider. This is in contrast to the market structure that was common prior to the 1990's where retail electricity providers were vertically integrated with wholesale generators, thus owning generating capacity and selling electricity to consumers.

The restructuring process broke up these vertically integrated firms into retail providers and wholesale generators and created a wholesale electricity market where generators and retail providers submit bids to sell and buy electricity. These markets are operated by a regional independent system operator that makes sure supply and demand in the electricity market balance in real time.

In addition to selling electricity into the wholesale market, generators that use

⁶The only states that require the renewable generation be located in the state are states that are a wholesale electricity market unto themselves, including Texas, California, New York, and Hawaii. Texas and Hawaii have their own electricity grids, while New York and California have their own Independent System Operators, thus making the state the natural unit of observation.

⁷The regions I will examine in this paper are the New England ISO, comprised of Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, and Vermont; the PJM control region, which includes Delaware, Maryland, New Jersey, Pennsylvania, Virginia, West Virginia, and parts of Illinois, Indiana, Michigan, and Ohio; and the Electric Reliability Council of Texas (ERCOT).

renewable sources also create a renewable energy credit with every megawatt hour (MWh) of electricity produced. Renewable energy credits (RECs) are a pure financial product (in most markets) that retail electricity providers are required to purchase to show compliance with a state's renewable portfolio standard.⁸ Typically a REC describes the attributes of the electricity that was produced such as the location of the generator, the fuel that was used, and the date that the electricity was produced. Using this information, retail providers can purchase RECs that qualify to meet a particular state's RPS. At the end of the year, retail providers retire the RECs that they have purchased to meet the RPS to the state regulator.

In order to ensure that retail providers comply with RPS requirements, nearly all states have set up a system of fines for retail providers that are short of their required number of RECs. These fines, generally called alternative compliance payments, effectively set a price ceiling in the market for RECs. If a retail provider has not purchased their required number of RECs, the alternative compliance payment specifies a dollar amount per megawatt hour that the retail provider must pay to the state. The level of the alternative compliance payment is usually determined by the state's public utility commission and is generally above the market price for RECs, giving retail providers an incentive to purchase RECs instead. Some states explicitly link the alternative compliance payment to a multiple of the market REC price, while others such as Massachusetts re-evaluate the penalty every few years to make sure the price is still above the market price for RECs. In many states, retail providers end up paying relatively few fines. For instance, in 2003, Massachusetts collected less than 1% of the RPS requirements through alternative compliance payments.

These RECs provide a second stream of revenue for renewable generators. Since

⁸California requires retail providers to enter into long term contracts with renewable electricity producers to purchase both the electricity and RECs to fulfill the state's RPS obligation. This requirement was lifted by the California Public Utility Commission in March 2010. California retail providers may purchase unbundled RECs to fulfill the RPS requirement in a limited amount in 2010 and 2011, with the market becoming completely unrestricted in 2012.

the average cost of renewable generation tends to be higher than that of fossil generation, the revenue from selling RECs encourages new renewable generating capacity to be built. The total price renewable generators receive for each megawatt hour of electricity is the price of the electricity plus the price of the REC.

1.2.2 Variation in State RPSs

Most states' RPSs have a final goal for their RPS by 2020 or 2025, usually between 10 and 30 percent of electricity sales, but the RPS is phased in over time. For instance, Massachusetts' end goal is for 15% of electricity sales to come from renewable sources by 2020, but interim requirements begin in 2003 at 1% of sales and increase by 0.5% or 1% every year until 2020. Other states have more aggressive schedules by increasing their renewable requirement by a larger amount every year while other states have large jumps in their requirements, such as California which has a requirement of 20 percent in 2019, and 33 percent in 2020. The light bars in Figure 1.1 shows a representative sample of RPS implementation schedules. These implementation schedules provide the variation across time that will allow estimation of the supply elasticity.

Another important dimension along which state policies differ is the treatment of how retail providers are required to comply with the policy mandate. In nearly all states, all retail electricity providers comply with an RPS policy by retiring renewable energy certificates (RECs). Some states require that the renewable electricity used to meet the RPS be produced within the state,⁹ while most other states just require that the generator that produced the electricity be part of the regional transmission organization so that, in theory, the electrons from the renewable electricity were in the same system. These geographic requirements attempt to get around the problem

⁹The states that require RPS generation to be within their borders are Hawaii, Iowa, North Carolina, New York, and Texas.

of creating a simple reshuffling of electricity purchases.¹⁰

State RPSs also vary as to which fuels that each state considers eligible to meet the RPS. While they generally include generation from sources such as wind turbines, solar (both photovoltaic cells and solar thermal sources), biomass (such as wood or wood waste), landfill gas, and small hydroelectric generation, some states also consider fuels such as municipal solid waste as renewable. In order to encourage specific types of fuels or to encourage a larger set of fuels, some states have created multi-tiered RPSs. Table 1.1 shows the fraction of states with an RPS that include each fuel type in the first tier of its RPS, the average carbon emissions per unit of heat input from using that fuel, and the average size of generators for each fuel. I will focus only on the first tier of each state's RPS since they provide stronger incentives to renewable generators (through higher prices for the RECs).

Each state must decide how to treat the renewable electricity generating capacity that exists when the RPS law is passed. Many states allow all existing renewable capacity for wind and solar generators to produce RECs that qualify to meet the state's RPS, but treat generators that use other fuels such as landfill gas, municipal solid waste, and especially hydroelectric facilities differently depending on when they were built and/or last modified in a substantive way. Most states consider hydropower installations that are smaller than 30 MW to qualify for the RPS. However, sometimes incremental additions to larger installations will not qualify for the RPS. Also, some

¹⁰Bushnell, Peterman, and Wolfram [9] show that local, consumer-based policies can be circumvented by a simple reshuffling of buying/selling pairs. For instance, consider an example where there are just two states with a common wholesale electricity market: state A has passed a 10% RPS and state B does not have an RPS. Moreover, assume that state B has enough renewable capacity to meet state A's RPS requirement, while state A does not have any renewable generating capacity. Before the RPS was enacted in state A, both states' retail electricity providers purchased all of their electricity from within their respective states. (Since there is a common wholesale electricity market, the electricity prices are equalized across states.) However, after the RPS is enacted, the retail electricity providers in state A switch to purchasing electricity from the renewable generators in state B and retail providers in state B switch to purchasing electricity from non-renewable generators in state A. Thus, state A's RPS has only resulted in reshuffling the buying/selling pairs and failed to increase the fraction of renewable generating capacity as a whole. As this example illustrates, this may be a major drawback of state level policies such as RPSs that interact with an interstate market. However, geographic requirements restrict the amount of reshuffling that is possible under an RPS.

states allow existing generators to meet only a fraction of the RPS requirement while new generators must meet the rest.

I have incorporated all of these important dimensions of heterogeneity into my empirical work and ensured that all of the facilities in my data that are eligible to meet an RPS requirement are actually eligible under the RPS rules of at least one state in the region. The next section will introduce a basic model of investment electricity generating capacity which will guide my empirical work.

1.3 Model

In this section I develop a model of investment in electricity generating capacity to illustrate the effect an RPS has on the incentives of electricity producers. The model provides the basic intuition of how policy can affect generating capacity decisions and motivates my empirical specifications.

Consider a representative firm deciding whether to invest in new generating capacity. For the power plant to be profitable, the revenue the generator produces over its lifetime must exceed its capital and operating costs:

$$\sum_{t=0}^T \beta^t \mathbb{E}[p_t q_t] \geq K_0 + \sum_{t=0}^T \beta^t (m_t + \mathbb{E}[f_t q_t]) \quad (1.1)$$

where p_t is the price of electricity at time t , q_t is the quantity of electricity the generator provides at time t , K_0 is the initial capital cost of the generator, m_t is the variable operating and maintenance costs associated with the generator at time t , f_t , is the fuel cost at time t , T is the number of years of the useful life of the new capacity, and β is a discount factor.¹¹ In equilibrium this condition holds with equality. This

¹¹This profitability condition abstracts away from any payments that generators receive for participating in ancillary service markets where generators may be paid to be on standby, ready to produce electricity if called upon by the market operator. Typically renewable generators are not eligible to participate in these markets due to the unpredictability of wind and solar generation. However, an increase in the amount of wind and solar generators participating in the electricity market will cause an increase in demand for standby capacity services from fossil generators.

equilibrium is graphically depicted in Figure 1.6, with the long-run supply curve for renewable generation (right axis) separated from the long-run supply curve for fossil generation (left axis). When demand is perfectly inelastic I can show these supply curves on the same figure with the length of the horizontal axis showing the total quantity of electricity demanded (inelastically).¹² The vertical line shows the fraction of demand that is met with fossil generating capacity versus renewable generating capacity. Since retail providers have no preference over the fuel used for electricity, the equilibrium fraction must be at a point where the price is equalized across types of generators.

A new generator will enter the market when there is sufficient excess quantity demanded over the life of the generator such that inequality 1.1 holds. (This is depicted graphically by lengthening the horizontal axis, necessarily increasing the market clearing price.) Two ways to induce generators to enter the market are to decrease the cost of the capital investment or to increase the price the generator will receive for its electricity over the new capacity's life span. These two levers have been used by the federal government to induce more renewable generators to enter the market in the form of the Investment Tax Credit and the Production Tax Credit, respectively.

Renewable portfolio standards also induce renewable generators to enter the market by shifting the fraction of demand that is met with renewable sources to the left within the figure. This is not a change in the total electricity demand but a change in the composition of production. This change causes a wedge in the price for electricity since renewable generators must receive a higher price for their electricity to build capacity. The price wedge resulting from the RPS is shown in Figure 1.7. The price renewable generators need to receive for their electricity to enter the market is p^r , but the equilibrium price given the number of generators in the market is p^e , so the

¹²This assumption is for convenience in the depiction of the wholesale market. The model does not need this assumption and is relaxed in the empirical work in the rest of the paper.

difference is made up by the price of the REC:

$$p^{REC} = p^r - p^e \quad (1.2)$$

Expanding equation 1.2 to show the total electricity price, the profitability condition in equation 1.1 becomes:

$$\sum_{t=0}^T \beta^t \mathbb{E}[(p_t^e + p_t^{REC})q_t] \geq K_0 + \sum_{t=0}^T \beta^t (m_t + \mathbb{E}[f_t q_t]) \quad (1.3)$$

for renewable generators. Renewable generators consider the price path of both the price of electricity and the price of RECs when considering entry decisions. So long as the price of RECs is expected to be greater than zero, renewable generators have an additional incentive to enter the market. Notice also that entry decisions depend on the flow of revenue to the generator over the life of the generator, not just the contemporaneous revenue.

The profitability condition implies that each generator has a critical (total) price at which it will enter the market. Therefore, as contemporaneous prices and expectations about future prices change we see generators entering the market consistent with the profitability condition.

We can derive a supply curve for renewable generators by aggregating each firm's decision about whether to enter the market. Each generator enters the market if their profitability condition holds. Thus, the total new generating capacity in the market at time t is:

$$Q_t = \sum_i \mathbb{I} \left[\sum_{t=0}^T \beta^t \mathbb{E}[(p_t^e + p_t^{REC})q_t] \geq K_i + \sum_{t=0}^T \beta^t (\mathbb{E}[m_{it} + f_t q_{it}]) \right] \quad (1.4)$$

where i indexes generators.¹³ Notice that there is a generator-specific capital cost,

¹³Note that the same condition holds for fossil generators, except their expectations over the price of RECs do not enter their profitability condition.

and each generator can expect a different amount of output. These two terms rationalize why we observe some renewable generators in existence in areas without a binding RPS. Consider a wind developer looking at potential locations to install a wind turbine. Not all locations are of equal value to the developer due to the fact that the wind blows at different speeds and different times at each location. Sites where the wind blows more frequently, all else equal, will be worth more to the developer since the turbine will create more electricity and has a marginal cost near zero. Thus, the best locations will be developed first with each subsequent wind turbine being placed in marginally inferior location, necessitating a marginally higher price for the electricity generated by that turbine to make it profitable. This suggests that existing renewable capacity satisfies the profitability condition in equation 1.3, but that as demand for renewable capacity increases, renewable generators will need to receive a higher price for their electricity. Thus, the upward slope of the supply curve is driven by heterogeneity in the value of locations and capital costs.

In order to aggregate across generators, I need to make several assumptions. The main assumption in the aggregation is that all generators have the same expectations over the trajectory of prices (electricity and RECs) over the life of each generator. With this assumption, I can rewrite the equation 1.4 as

$$Q_t = f(p_t^e, p_{t+1}^e, \dots, p_{t+T}^e, p_t^{REC}, p_{t+1}^{REC}, \dots, p_{t+T}^{REC}) + \delta X_t \quad (1.5)$$

where X_t is a set of variables capturing the other factors that effect a generator's entry decisions such as fuel costs.

This equation suggests that I can estimate the price elasticity of supply of renewable generation using the familiar log-log specification by regressing the log of quantity of new renewable capacity on the log of price and other factors that affect entry decisions. However, this presents the traditional problem of simultaneous equa-

tions bias since price and quantity are determined by the intersection of supply and demand. Since, price is an endogenous regressor, I need an instrument for the price that renewable generators receive to consistently estimate this equation.

1.4 Empirical Strategy

In order to estimate the long-run price elasticity of supply of renewable capacity, I use the implementation schedules of state RPSs as an instrument for the price that renewable generators receive for their electricity and then use the predicted change in price in a second-stage regression to estimate the price elasticity of renewable generation. RPS implementation schedules provide me with an exogenous change in the demand for renewable generating capacity that can instrument for the changes in price that renewable generators receive. RPS implementation schedules are typically written into the original RPS legislation and increase the RPS requirement each year that the RPS is in effect until the end goal is met. Because these schedules are incremental changes in demand that are not determined at the same time as the price, and therefore are not correlated with unobserved supply shocks, they are a good instrument for the total price that renewable electricity generators receive for their electricity.

This leads to a way to estimate the the long-run price elasticity of supply for the renewable generators. Each new RPS requirement increases the demand for RECs, increasing the wedge between the price that renewable and fossil generators receive for electricity. Importantly, the change in demand that I observe in the REC market comes from the RPS legislation, making the variation more plausibly exogenous. Thus, I can separate the change in total electricity price renewable electricity generators receive due to the RPS from other market forces to trace out the supply curve of renewable generating capacity and estimate the price elasticity of supply.

1.4.1 Estimating Equations

My model suggests two natural estimating equations to estimate the long-run price elasticity of supply. The first-stage equation estimates the price response to an exogenous change in the demand for RECs. The demand for RECs change in a predictable way due to the implementation schedule of each state's RPS. As derived in section 1.3, new renewable generating capacity should respond to the entire flow of payments over the life of the generator. To capture this variation I use a measure of changes in RPSs' stringency averaged over the next five years.¹⁴ This leads to a first stage equation of the following form:

$$\log(p_{it}^{total}) = \beta \log(RPS \text{ Requirement}_{it,t+5 \text{ years}}) + \delta X_{it} + \alpha_i + \gamma_t + \epsilon_{it} \quad (1.6)$$

where the α_i 's are region fixed effects, the γ_t 's are year and month fixed effects, and X_{it} 's are a set of controls for other policy variables that may effect the incentives of renewable generators. The region level fixed effects absorb time invariant differences across regions such as differential renewable generating potential, as Lyon and Yin [23] suggest may be important in the decision to adopt an RPS. The year and month fixed effects absorb differences across time that are constant across region. These are important since over our period of examination various federal tax incentives have taken effect (and occasionally not been renewed immediately) such as the Investment Tax Credit and the Production Tax credit that affect the financial desirability of building renewable generation. For a discussion of the history of these policies and their consequences see Metcalf [25], Wiser et al. [41], and the Joint Committee on Taxation [3]. I also include a group of other policy variables in both estimating

¹⁴I will discuss how this variable is constructed in Section 1.4.3. The results are not sensitive to the choice of a five year average. The results are similar for averages up to 10 years of the RPS requirement, though are noisier the longer the time period that is averaged. All of the specifications have also been estimated using the future requirements instead of an average. The results are identical in these specifications too.

equations, X_{it} , to control for other policies that affect the incentives for renewable electricity providers unrelated to RPSs. These variables allow for a more isolated estimate of the effect of *only* the RPS.¹⁵

As discussed in Section 1.3, generators should be making entry decisions based on the time path of prices, not just contemporaneous prices. In order for the identification strategy to identify the long-run price elasticity, generators must also forecast that today's REC prices will persist into the future, and therefore that the RPS, as written, will persist into the future.

It is likely that these conditions are met for a number of reasons. Firstly, all states allow RECs to be saved for use in future compliance periods. Typically, RECs can only be saved for between two and five years, but so long as there is not a large oversupply of RECs for a sustained period of time older dated RECs can be retired for compliance while the newer dated RECs are re-banked. Because of this banking feature of RECs, the contemporaneous price should contain all information and expectations at the future prices of RECs. Therefore we should observe generators responding to the contemporaneous price since it is also a signal about future prices.

The second stage equation that will give us an estimate of the price elasticity of supply takes the following form:

$$\log(\text{RenewableCap}_{it}) = \beta \widehat{\log(p_{it}^{\text{total}})} + \delta X_{it} + \alpha_i + \gamma_t + \epsilon_{it} \quad (1.7)$$

where the α_i 's are region fixed effects, the γ_t 's are year and month fixed effects, and X_{it} 's are a set of controls for other policy variables as in the first stage.

1.4.2 Which States get RPSs?

In order for the instrument of RPS stringency to be valid, it must not enter the supply equation except by entering through the demand equation. The implementa-

¹⁵These other policy variables are discussed in detail in Section 1.5.4.

tion schedule may indeed enter the supply equation if, for instance, states that have more renewable generating potential choose to adopt RPSs. This would cause me to overestimate the effects of RPSs since those states are also likely to develop more renewable generation capacity than other states even in the absence of an RPS.

Upon casual observation of the dates at which various states adopted their RPS policies, this doesn't seem to be a particularly large problem. (See Figure 1.3 for a description of when each state passed their RPS.) Some of the states with the largest renewable potential from both wind and solar are in the Plains states and the Southwest. While many of the states in the Southwest do indeed have RPSs they are not uniformly early policy adopters. Conversely, most of New England and the North Atlantic states have adopted RPS policies, some being among the first adopters but do not have a large renewable generating potential.

Moreover, many of the first adopters of RPS policies adopted their RPS as part of the electricity restructuring legislation. The electricity restructuring legislation in many of these states was a major piece of legislation that separated the retail and wholesale electricity markets, making the latter market a "deregulated" market, usually with plans to make the retail electricity market a competitive, unregulated market in the future. Most of the deregulation of the electricity markets were motivated by high retail electricity prices in the state and generally a group of states in a region deregulated the wholesale electricity market at similar times. There is very little reason to believe that the deregulation legislation is correlated with unobserved covariates that affect renewable electricity capacity.

Lyon and Yin [23] empirically examine which states get RPSs. Their findings suggest that wind potential in the state increases the probability of RPS adoption (though not potential in other fuels that are typically included in RPSs such as solar or biomass). This will not be a problem for me since I will be controlling for this variation through my region level fixed effects.

Lyon and Yin also find that high local pollution levels, as measured by the fraction of the population living in counties that are designated as “nonattainment” under the Clean Air Act, increase the probability of adoption of an RPS as well as some evidence that organized renewable energy lobbying groups increase the probability of adoption. In contrast to Rabe’s [34], [35] qualitative examination, they also find that a state’s unemployment rate decreases the likelihood that a state will adopt an RPS. Rabe [35] finds that states often emphasize the potential economic benefits of RPSs such as creating “green” jobs or gaining a competitive advantage as a first mover in renewable energy technology, but this does not seem to be a driving factor empirically as measured by the unemployment rate. These papers give me confidence that RPS adoption is likely to be uncorrelated with many of the unobservables that would invalidate the instrument. Moreover, since RPS policies affect neighboring states as well as the states in which they are passed, they are even less likely to be correlated with in-region unobservables.

Since my instrument for changes in demand for renewable capacity is not just the beginning of an RPS in the region, but also the implementation schedule that each RPS follows, the implementation schedule must be uncorrelated with in-region unobservable characteristics. Two main concerns come to mind when considering the exogeneity of the RPS implementation schedules. First, states that are early adopters of RPSs may have particularly aggressive implementation schedules, either due to a strong desire to promote renewable electricity generation or because they have a lot of renewable resources that can be exploited.

A second concern about the exogeneity of RPS implementation schedules is that states that have a lot of renewable generating capacity at the time the RPS is passed will have more aggressive implementation schedules. Since more aggressive implementation schedules likely lead to higher REC prices sooner, existing renewable generators clearly have a lot to gain by lobbying state legislatures for more stringent require-

ments. This lobbying may happen since more stringent requirements will lead to higher REC prices and this is a windfall profit for generators that find it profitable to operate in the absence of selling their RECs.

Both of these concerns can be addressed empirically by examining the state implementation schedules. Since most RPSs follow a nearly linear implementation schedule, I estimate the slope of the implementation schedule by regressing each schedule on a time trend. I then examine the correlation between these slopes and variables that address the concerns raised above about the endogeneity of implementation schedules.

To address the first concern that early adopter states have more aggressive implementation schedules, I regress the slope of the implementation schedule requirements (in MW of required new capacity) on the year that each state's RPS went into effect. The coefficient on the year the RPS went into effect is not statistically different from zero with a coefficient of -7.7 and a standard error of 17.7. The point estimate suggests that early adopters require an extra 8 MW of renewable capacity each year of an RPS but is clearly not statistically different from zero ($p = 0.67$). Moreover, an additional requirement of 8 MW per year is a relatively small difference given that the average increase in RPS requirements is 166 MW per year.

To address the second concern that states with a larger renewable sector before RPS passage will have a more aggressive implementation schedule, I regress the slope of the implementation schedule on the renewable capacity in that state at the time of RPS passage. The coefficient is not statistically different from zero with a coefficient of 0.18 and a standard error of 0.11. This suggests that the implementation schedules are not a function of the renewable interests already established in a particular state.

This may be because many states choose “round” numbers for both their end goal, such as “20% renewable electricity by 2020,” and a linear implementation schedule. Likely, these end goals are less amenable to manipulation by pressure groups and since the intervening years' requirements are essentially a linear interpolation back

through time, the implementation schedule is not changed much by pressure groups.

Another concern about my approach is that the RPS policies may spill over into other regions that do not have RPSs or less stringent RPSs. However, there are likely only very small spillover effects in my setup since the unit of observation is a regional electricity market. Many states publish a list of all of the approved generation facilities that are eligible to produce RECs that meet the state RPS. While occasionally there are power generators located in states not included in the wholesale power market, a vast majority of the approved generation facilities are indeed located in states in the wholesale power market.

1.4.3 Key Variables

The primary variable of interest in the first stage regression is a variable constructed to measure the stringency of a particular state's RPS. Most states, with the exception of Iowa and Texas, set their RPS goals out as a percentage of electricity sales, measured in megawatt hours (MWh). For instance, Michigan's RPS, passed in 2008 calls for 10% of each retail provider's electricity sales to come from renewable sources by the end of 2015 with a phase-in period beginning in 2012. The first challenge we face is converting an RPS goal stated in MWh¹⁶ to our capacity data in MW. One megawatt hour of electricity is created simply by a 1 MW facility producing at full capacity for one hour. This means, in theory any facility's nameplate capacity (in MW) can be converted into a yearly capacity in MWh by multiplying the nameplate capacity by 8760(= 24 × 365) hours.

However, generators do not run the entire year since they must be shut down for maintenance and may choose not to operate for any number of reasons, including bidding in a price that is higher than the market clearing price in a particular hour. Plants that are almost always producing are usually large coal and nuclear plants that

¹⁶RPS goals are usually stated as a fraction of electricity sales, measured in MWh. Thus, it is simple to convert percentage goals into MWh.

operate between 85%-90% of the time (capacity factor of 85%-90%). Other plants are built to only operate a small fraction (as little as 1% or less) of the time, when the demand (and hence price) for electricity is at its peak. These generators tend to use natural gas since they can bring themselves up to full capacity quickly. Wind generators typically have a capacity factor near 35% [45] since wind is an intermittent resource. For the purposes of our main analysis, I assume that all new plants have a capacity factor of 40% since most of the needed capacity to meet RPSs is expected to be wind but some of it will be met with fuels that can have a significantly higher capacity factor [42].¹⁷

In order to correctly measure the incentives of these policies I first need to construct a variable to measure the eligible megawatt hours of renewable generation for state i at time t .

$$\text{RPSCapMWh}_{it} = \sum_{f \in F} \text{RPSCapMW}_{itf} \times 8760 \times \text{AvgCF}_f \quad (1.8)$$

where f is a particular fuel and F is the set of all eligible fuels. RPSCapMW_{itf} is the sum of nameplate capacities of all generators in state i at time t for fuel f , and AvgCF_f is the average capacity factor for fuel f .¹⁸ I set AvgCF , the average potential capacity factor for a particular fuel, equal to 0.4 for wind and solar generators and 0.8 for all other generators.

Returning to the Michigan example, in order to figure out how hard this goal is to reach, I must consider how much eligible renewable capacity already exists in Michigan to meet the RPS. Some states have RPS implementation schedules such that during the first few years of the RPS, the whole requirement can be met with existing generating capacity.¹⁹ As mentioned above, each state treats existing capacity

¹⁷Figure 1.5 shows that it is indeed the case that most of the change in generating capacity over the last few years has been in wind capacity.

¹⁸Capacity factor is the fraction of hours in a year that a generator is producing electricity.

¹⁹This means it is possible for an RPS to create zero incentive in some or all years. For instance,

differently. Since Michigan allows existing renewable capacity to be eligible to meet the RPS, to compute the incentives created by the RPS, I subtract the eligible capacity at the time the RPS was passed from each year’s RPS requirement.

Using the amount of renewable generating capacity at the time the RPS was passed, I calculate the RPS stringency measure in megawatts as,

$$\text{Stringency}_{it} = \frac{\text{RPS Req.}_{it} \times \text{ConsumpMWh}_{it} - \text{RPSCapMWh}_{i0}}{8760} \times \frac{1}{\text{AvgCF}} \quad (1.9)$$

This stringency measure will be our key independent variable as it capture much of the heterogeneity across state policies and I expect the coefficient on it to be positive and significant.

After computing these variables on the state level, they are aggregated up to the regional level by weighting them by each state’s electricity consumption share in the region.²⁰ Moreover, since I am using price data in the second stage, I need to limit my sample to regions that have a robust wholesale electricity market and REC market. Thus, I will be focusing on three regions of the country: New England, the Mid-Atlantic states in PJM, and Texas.²¹

The final I need to construct is the complete price that renewable generators receive for the electricity they produce. As mentioned above, there are two revenue streams for renewable generators under an RPS: the revenue from each megawatt

Maine passed an RPS in September 1999 that had a final goal that was less than the existing eligible capacity within the state. Subsequently, in 2006 Maine passed another RPS that only new renewable facilities were eligible to meet the requirement. This is the RPS we consider for Maine in this paper.

²⁰The weights are computed using a state’s consumption share in 2003 to keep them across time. Changing the weights to the contemporaneous consumption share in each region does not change the results.

²¹I exclude California from the analysis because until recently, there was not a market for RECs since the California Public Utility Commission required retail providers to purchase both renewable electricity and its attributes (essentially RECs) together via bilateral (private) contracts. Therefore, there is not a market price for RECs to use in the second stage. I exclude Midwest states since there is not a developed market for RECs. I also exclude New York since the New York State Energy Research and Development Authority (NYSERDA) centrally procures the RECs for the entire state’s commitment through an annual bidding process. It is not clear that this processes elicits the same price due to possible market power on behalf of the NYSERDA.

hour of electricity they sell into the electricity market and the revenue they receive from each REC associated with each megawatt hour of electricity that retail providers retire at the end of each year to comply with the RPS. This means that the complete price for renewable generators is

$$p_t^{total} = p_t^{electricity} + p_t^{REC}. \quad (1.10)$$

1.5 Data

To estimate the empirical models, I use data on all existing electricity generators and production from the Energy Information Administration (EIA). I collect data on REC prices from public utility commission reports and electricity price data from independent system operators of wholesale electricity markets. Data on central policy variables are aggregated from the North Carolina State University's Database of State Incentives for Renewables & Efficiency.

1.5.1 Generation Capacity Data

The Energy Information Administration annually surveys all electricity generators to collect basic data for each electric generating unit in the United States. All generators that have a potential capacity of at least 1 megawatt, are connected to the electric power grid, and are able to deliver power are required to fill out form EIA-860. The data files include information about each generator including its capacity, all fuels used during that year, the year and month the generator began operation, the year and month of retirement, the city and state that the plant is located in, and basic information about the owner.

Though these data are reported annually to the EIA, they can be translated into monthly data on total generating capacity since the data report the first month of operation for each plant. I use the data reported in these surveys from 1999-2007.

In these data, each generator provides detailed data on the type of fuel used during that year for electricity generation, including up to six or more fuels that were used. I consider the first fuel listed as the generator's primary fuel.²² For the purposes of this paper, I aggregate these fuels into 17 categories. Most of the fuels that are lumped into the same category are different types of coal, petroleum products, and various waste products. None of these fuels are considered to be renewable in any of the RPSs and therefore this aggregation should not affect the results.

1.5.2 Electricity Sales and Production Data

In order to translate a RPS requirement that is usually in terms of percent of sales of electricity and to create each state's weight in the region I use data from the Energy Information Administration's state historical tables on electricity sales. These data report the total megawatt hours sold in the entire electric industry for a given state in a given year. I also use data on electricity generation aggregated to the state×year level by the EIA. These aggregate data are based on another survey the EIA conducts, EIA-906.

1.5.3 Price Data

Data on the wholesale price of electricity were collected from each Independent System Operator's (ISO) web site. ISOs publish data on the market clearing price of electricity for many locations in each region for every hour of the day. Where available, I use the published regional weighted average price for each hour and then average the price over each calendar month. Some ISOs do not publish a regional electricity price, instead only publishing data for each location in the ISO. Where this is the case, I take a simple arithmetic mean of the prices across all locations to form

²²Only one-third of plants report using two fuels, and less than 5 percent report using more than two fuels. Of the plants that list using two fuels, only 6 percent of generators that are categorized as using a renewable fuel list a non-renewable fuel as their second fuel, concentrated among generators that are categorized as biomass, landfill gas, and municipal solid waste.

an hourly regional price and then average this mean over the entire calendar month.

In order to compute the complete price that renewable electricity generators receive for their electricity, I need to add the price of renewable energy credits (RECs) to the price of the electricity. I have collected average annual prices for RECs in every state that allows RECs to be purchased separately from electricity. These prices are gathered from public utility commission documents or other agencies administering a state's RPS.

The raw price data for REC prices can be seen in Figure 1.8. The state REC prices exhibit distinctly regional variation confirming that the market for RECs is indeed regional. Much of the within-region variation is due to some temporary state-level policy uncertainty and small variations in eligible fuels, as well as small variations in which generators are certified in which state.

1.5.4 Policy Variables

The policy variables are constructed from information compiled at North Carolina State's Database of State Incentives for Renewables & Efficiency (DSIRE). DSIRE has cataloged all state incentives for renewable energy including the date they were enacted, when and if they were modified, as well as many details about each policy. Where necessary, this information was supplemented by consulting the actual state statutes.

The variables that were constructed include the date that a particular renewable energy policy was passed by the legislature, when the policy began to bind (if different), and the implementation schedules for RPSs. In addition, information for each RPS regarding what fuels are eligible to meet the requirements, and in some cases maximum capacities for eligible facilities, were taken from this database.

In addition to collecting data on state RPSs from DSIRE, I collect data about other policies that have been implemented in some states that could change the

incentives for renewable electricity generators. These policies include:

- **Net metering:** This type of legislation requires that electricity meters “run backward.” If a customer has installed generation equipment on site (usually a photovoltaic solar panel) that produces more electricity than a customer is currently consuming, the excess electricity is fed back onto the grid and the customer’s electricity bill is credited the retail electricity rate for each kilowatt hour. (See Borenstein [8] for an analysis of these policies.) This net metering may provide an additional incentive for electricity customers to invest in their own generating capacity and then sell the RECs from this generation.
- **Public Benefits Fund:** In many states with competitive wholesale electricity markets, retail electricity providers are required to levy a surcharge on all rate payers to remit to the state government. This money is often used for energy efficiency programs, to help finance renewable energy projects including transmission and distribution projects, and to assist low-income rate payers. Since these funds partly subsidize renewable generation, they are controlled for in the regressions.
- **Government Purchases of Green Electricity:** Some state governments have committed themselves to purchasing a share of their electricity from renewable sources. Since governments are large customers, this may (and is presumably hoped to) affect the amount of renewable capacity. Though both government purchases of green electricity and RPSs require retail providers to retire RECs in the amount of the green purchases, the RECs retired are not counted toward a retail provider’s RPS requirement.
- **Mandatory Green Power Option:** Some states have passed legislation that requires retail electricity providers to offer their customers an option to purchase green electricity. Retail providers are allowed to charge extra for providing this electricity. These customer purchases generally are explicitly forbidden from

counting toward the RPS requirement. These policies may, however, increase renewable capacity if a sufficient number of customers sign up for these programs.

Table 1.2 displays summary statistics for the policy variables listed above. The top panel displays summary statistics after aggregating state policies to the region level, and the bottom panel displays the RPS requirements during my sample period for individual states. These bottom statistics correspond to the values of the light blue bars in Figure 1.1.

Just over half of the region-months in my sample have an active RPS in the region, with a mean renewable requirement of 0.6% renewable generation and a maximum of 2.5%. The mean RPS requirement, conditional on an RPS being enforced, is just over 1% of electricity consumption coming from renewable generation. Taking a look at the state-level data, I observe just 20% of state-months in our sample with an active RPS, with an average requirement of 1.7% renewable generation, conditional on an operational RPS.

1.5.5 Data Restrictions

As discussed before, there are many dimensions of heterogeneity across state RPSs. In order to simplify my analysis, I will only consider the first tier of each state's RPS. Typically, if a state has multiple tiers to its RPS, the second, third, and fourth tiers allow a greater degree of flexibility for fuels that have higher carbon emissions per unit of heat input. Tiers two and below tend to include fuels such as municipal solid waste or large, existing hydroelectric facilities. (Some states with a single tiered RPS include these fuel types in the RPS.)

If a particular fuel counts for both tier one and tier two in a state, I attribute all of the capacity from facilities using that fuel to fulfilling the first tier of the RPS. Compliance RECs for the first tier are uniformly more expensive than compliance

RECs for other tiers (with the exception of states with a solar photovoltaic tier), so this assumption is likely consistent with firm incentives [44]. To the extent that not considering these other tiers of RPSs biases my results, the results should bias the elasticity toward zero. This is because I may be excluding some renewable facilities that may have been built in response to an RPS.

I also do not examine the solar photovoltaic (PV) tiers of state RPSs. Usually if a state has a specific tier for PV, it is the only fuel in that tier. These tiers usually have small requirements, since PV is an expensive way to produce electricity. Moreover, most PV installations are excluded from my data since only generators over 1 megawatt are required to report to the EIA. PV installations tend to be less than 0.1 megawatts, since many of these installations are on the roofs of residential or commercial buildings.

1.5.6 Aggregation

Aggregating state policies to a regional policy is relatively straightforward. Each state's RPS requirement is weighted by the fraction of electricity consumption that the state accounts for in the region. Thus, if a region consists of three states, state A consumes 50% of the electricity in the region and states B and C each consume 25% of regional electricity. If state A passes an RPS that requires 2% of the electricity sold in that state, the region then is assigned an RPS requirement of 1% ($= 2\% \times 50\%$). In the following year, state A's requirement increases to 3% and state B introduces a 1% requirement so the region's RPS requirement is then 1.75% ($= 3\% \times 50\% + 1\% \times 25\%$).

The other state level policy variables (public benefits funds, green power options, etc.) are aggregated in a similar fashion to this, except each variable is simply an indicator for each state, so the variables take on the cumulative fraction of electricity consumption in the region covered by those policies.²³

²³For simplicity, the weights used are calculated as the state's fraction of consumption in the region during 2003. This keeps the policy and RPS variables weakly monotonic across time. It is

1.6 Results

The results from the first stage regression are displayed in Table 1.3. Column 1 begins by simply regressing the logarithm of the total price for renewable electricity (electricity price + REC price) on the logarithm of the average effective RPS requirement for renewable capacity in that region over the following five years.²⁴ As mentioned above, the average requirement is used since it is correlated with future stream of payments over the lifetime of the generator.

We see that the measure of the stringency of an RPS is statistically and economically significant. Column 2 allows each type of control policy to have a one-off effect in the region once any state adopts it. Column 3 instead adds control variables that can take values between zero and one depending on the fraction of electricity consumption in the region that is covered by one of the policies. Column 4 allows for both a one-off effect in the region and an increasing effect over time as more states in the region adopt these policies. This flexible specification makes sense intuitively, since I would expect that the more expansive these policies are, the larger in magnitude the effect should be. The point estimates for the excluded instrument, the average stringency of the RPS over the next five years, are relatively stable across columns.

Examining the other coefficients in Column 4, the coefficients match my intuition about the direction of the effect. I expect a positive effect on REC prices from government purchases of green power since this increases the demand for green power without decreasing the demand for RECs. Most states do not allow green power purchased through government purchases, to count toward retail providers' REC fulfillment obligations; instead these purchases simply add buyers into the green electricity / REC market. I also see that public benefits funds tend to reduce the price

unlikely that generators can accurately predict the small variations in electricity consumption across regions for them to take these fluctuations into account. Changing the year used for the weights or using contemporaneous weights do not change the results.

²⁴All specifications are robust to the number of years over which the effective requirement variable is averaged.

of RECs. Again, this matches our intuition since often the money collected in public benefits funds is used to subsidize the construction of renewable generating facilities, thus reducing the price needed to make the facilities profitable.

The last row of each column shows the F-statistic of a test that the excluded instrument in the regression is zero. All four columns reject the null that both coefficients are zero at all usual levels of confidence. This gives confidence moving forward that the instruments are indeed relevant.

All standard errors in this table and the second stage regressions are estimated using Newey-West heteroskedasticity and auto-correlation robust standard errors. The number of lags included in the auto-correlation estimation was chosen using the procedure suggested by Newey and West [29].

Table 1.4 displays the results from the second stage regression that estimates the price elasticity of supply for renewable electricity generators. The variable of interest in this set of regressions is the first row, $\widehat{Log(Total Price)}$. This is the predicted price of RECs in the region given the shift in demand induced by the stringency of the state RPS estimated in the first stage. Though the elasticity estimates vary across specifications, the preferred estimate in Column 4 is between the other estimates, which allows the other policies to enter in multiple ways.

Column 1 shows a baseline specification without any additional controls, with Columns 2-4 progressing to a full set of flexible controls for other policies aimed at renewable generators. The preferred estimate in column 4 of the price elasticity is 2.714. Thus, for every 1% increase in the price of RECs, there will approximately a 2.7% increase in renewable generating capacity. In the next section I will use this estimate to bound the cost of focusing on reducing greenhouse gas emissions through only an RPS-style policy.

1.7 Policy Implications for RPSs as a CO₂ Abatement Tool

In this section, I use my estimates of the long-run supply elasticity of renewable generating capacity to estimate the cost of decreasing carbon dioxide emissions in states covered by the Regional Greenhouse Gas Initiative by pursuing carbon dioxide reductions exclusively through an RPS.

The Regional Greenhouse Gas Initiative (RGGI) is a cap-and-trade program established in the northeastern United States to reduce greenhouse gas emissions from electric power plants to 10 percent below (approximately) 2005 levels by 2018. There are currently ten states participating in RGGI, including all of the states in the New England wholesale electricity market, New York, and parts of the PJM wholesale electricity market.²⁵ In these states, RGGI regulates all fossil fuel fired electricity generators in the 10 states that have a capacity of 25 megawatts or more. Each quarter, new emissions permits are auctioned with approximately 70% of the auction proceeds being invested in energy efficiency and renewable generation projects.

The states in RGGI had a total of 184 million tons of carbon dioxide emissions from the electricity sector in 2005²⁶ [1]. Beginning in 2009 and continuing through 2014, carbon dioxide emissions are capped at the baseline level of 188 million tons. Beginning in 2015, the carbon dioxide cap is reduced by 2.5% annually until the final goal is met after 2018 when carbon dioxide emissions are reduced by 10% from the original cap. Under the RGGI cap-and-trade program, emission reductions are most likely to come from using a different mix of fuel to produce electricity (more natural gas and renewable sources, less coal) and energy efficiency investments. This suggests that the carbon price in the RGGI market provides a good cost estimate of reducing

²⁵The ten states currently participating in RGGI are: Connecticut, Delaware, Massachusetts, Maryland, Main, New Hampshire, New Jersey, New York, Rhode Island, and Vermont. The major states in the PJM wholesale electricity market that are not participating in RGGI are Ohio, Pennsylvania, Virginia, and West Virginia.

²⁶The baseline level of carbon dioxide emissions that the RGGI reductions are based on is 188 million tons of CO₂.

greenhouse gas emissions in the electricity sector through ways that differ from an RPS's exclusive reliance on switching production to renewable sources.

It is important to note when comparing the cost of CO₂ abatement from these two policy instruments that the each cost includes slightly different components. Firstly, because RGGI is taking place in states that already have renewable portfolio standards, the cost of carbon abatement from RGGI does not include the cost of CO₂ abatement from any emissions reductions that take place because of the RPSs. Thus, we would expect the cost of RGGI permits to be slightly lower than they would in a world where there is only a cap-and-trade program.

Secondly, the source of the CO₂ reductions is different under and RPS and RGGI. Since most of the emission reductions in RGGI come from switching fuel, they are only as permanent as RGGI is. In contrast, since RPSs create new generating capacity, the CO₂ reduction benefits will continue to accrue without an RPS in place so long as the operating costs of the generators are lower than the price for which the electricity is sold. For instance if all RPSs were repealed today, generators would likely suffer a large loss on their capital investment but would still find it profitable to operate if they have a low marginal cost of electricity production (and the capital has a sufficiently low scrap value.)

I will examine two different levels of carbon reduction produced by a northeastern RPS, a 2.5% reduction of 2005 CO₂ levels and a 10% reduction of 2005 CO₂ levels, to compare to the cost of carbon dioxide abatement through RGGI. In order to estimate the cost of carbon dioxide abatement under an RPS, in addition to knowing the price elasticity of supply or renewable generation that I estimated in the previous section, I need to make a few assumptions. Whenever possible, I will make assumptions that will make an RPS look as favorable (lowest cost of carbon dioxide abatement) as possible so my estimates will be a lower bound on the cost of CO₂ abatement under an RPS.

First, I need to make an assumption about what fossil fuel the new renewable capacity will be displacing. As can be seen in Table 1.1, coal is the fuel that emits the most amount of CO₂ per unit of heat input at 215 pounds of CO₂ per million British Thermal Units (MMBTU). Therefore, to make an RPS look as attractive as possible, I will assume that each megawatt hour of renewable generation produces no carbon dioxide and replaces a megawatt hour of coal production.²⁷ To the extent that renewable generation produces carbon dioxide or displaces generation other than coal, an RPS would have a higher cost of carbon abatement than I estimate.

Secondly, I need to assume a capacity factor (the fraction of the year that a generator produces electricity) for the new renewable generation built to meet the RPS. As discussed above, a capacity factor of 85% is in the upper range for fossil generation and 35% is relatively high for wind generation. I assume a capacity factor of 40% for all new renewable generation that is built for the RPS. This acknowledges that most of the renewable generation being built in response to RPSs are wind turbines, but some is likely to be from other sources with a higher potential capacity factor such as biomass and landfill gas generation.²⁸

Finally, I need to assume something about how the demand for electricity changes in the future. I will assume that electricity consumption does not change from the amount consumed in 2005. Likely, electricity consumption will grow between now and 2015 (the first year that the RGGI CO₂ cap is decreased).²⁹

In 2005, the total renewable generating capacity in RGGI states was 2,932 megawatts.

²⁷For renewable resources, such as wind, that are not completely predictable, there is some carbon emissions from using these sources since more generators need to be on standby in case the electrical output is less than expected. Compounding this, usually these standby generators need to ramp up their output quickly which creates higher than average emissions per MMBTU consumed. I abstract from both of these issues.

²⁸Another renewable source that is included in nearly all RPSs is solar photovoltaics (solar cells). However, these typically have a capacity factor near 15% and are currently too expensive to be deployed on a large enough scale to make a significant contribution to renewable capacity.

²⁹Electricity consumption has grown by approximately 2.75 million megawatt hours annually in the northeast between 1990 and 2008. Electricity consumption in the northeast in 2005 was 278 million megawatt hours.

If every megawatt hour of renewable generating capacity displaces a megawatt hour of coal generation, a 40 percent increase in renewable generating capacity would achieve a 2.5% decrease in CO₂ emissions in RGGI. Using my preferred elasticity estimate of 2.7, this means renewable generators would need a price increase of 15% in order to be profitable. Since the total price of electricity for renewable generators (electricity price plus REC price) averaged \$82 per megawatt hour, a 15% increase implies that renewable generators would need to receive approximately \$94 per megawatt hour to be profitable. This implies a marginal cost of CO₂ abatement of over \$12.

A more reasonable assumption is that each megawatt hour of renewable electricity produced replaces the carbon emissions of an “average” megawatt hour,³⁰ which requires a 68% increase in renewable generating capacity. The preferred elasticity estimate tells us that renewable generators would need a 25% increase in price to enter the market. This implies a marginal cost of CO₂ abatement of \$35 per ton of CO₂.

The final goal of RGGI is to reduce CO₂ emissions by 10 percent from their 2005 levels. In order to achieve a reduction of this size from an RPS, there would need to be a 163% increase in the renewable generating capacity in RGGI states. In order for my estimate of the cost of CO₂ abatement to be correct, the price elasticity has to be correct for the entire supply curve. While I am comfortable making the assumption that my price elasticity estimate is nearly correct for the smaller changes above, I am hesitant to believe my estimate is correct for this large of an increase in renewable capacity.

With this caveat in mind, I proceed to extrapolate the cost of CO₂ abatement from an RPS that reduces CO₂ emissions by 10% from their 2005 levels. In order to move that far up the long-run supply curve for renewable generation, the total price renewable generators would need to receive for their electricity is \$132, implying

³⁰Average is defined as the total tons of CO₂ produced by the electricity industry divided by the number of megawatt hours consumed in a year.

a marginal cost of CO₂ abatement of \$50. If instead of replacing coal generation, renewable capacity replaced the “average” megawatt hour of generation, the marginal cost of CO₂ abatement to decrease emissions by 10% is \$140. A summary of these results can be found in Table 1.5.

These estimates of the marginal cost of CO₂ abatement, ranging from \$10 to \$140 depending on how much carbon is abated and what fuel the renewable generation replaces, are substantially higher than the expected cost of carbon abatement under the RGGI cap-and-trade system. Currently the RGGI CO₂ emissions permits being traded and auctioned are for the years when CO₂ emissions are capped at a level just over 2005 CO₂ emissions. However, since these permits can be banked indefinitely into the future, they give us a window into the expected marginal cost of CO₂ abatement in the future. Currently the price of emissions permits are at approximately \$2 per ton of CO₂, and permits were trading at approximately \$3 per ton of CO₂ in early 2009 with an average price of \$2.50 per ton of CO₂ over two years of trading. Since permits purchased today can be used to comply with RGGI indefinitely into the future, the current price is indicative of future CO₂ abatement costs. Moreover, the price of contracts for CO₂ emissions permits in 2012 (the furthest ahead future contracts are traded at the moment) is similar to the current price of CO₂ permits further suggesting that prices are not expected to increase.³¹

These results suggest that within the electricity sector, an RPS is an expensive way to decrease carbon dioxide emissions, costing between six and seventeen times more to reduce CO₂ emissions by 2.5% than from a cap-and-trade program. Moreover, since both RGGI and an RPS focus just on the electricity sector, the marginal cost of CO₂

³¹RGGI has two mechanisms built into its structure to curb potential price volatility. If the average price of CO₂ permits is above \$7 for a 12-month period, more permits are released and generators are allowed to meet more of their obligations through offsets. If the average price of CO₂ permits is above \$10 for a 12-month period, a second mechanism is triggered and even more offsets can be used to meet CO₂ obligations. It is widely expected that neither of these trigger events will happen, suggesting that it is unlikely that the the marginal cost of CO₂ abatement is below \$10 in the electricity sector in the states in RGGI.

abatement in the economy as a whole is likely lower than either of these estimates since there may be cheaper ways to decrease CO₂ emissions in other sectors of the economy.

1.8 Conclusion

This paper estimates the long-run supply elasticity of renewable electricity generating capacity. The price elasticity is an important parameter for policy makers to know since many states have introduced aggressive RPSs to increase the share of renewable electricity sold in their states. Also, the US Congress has considered legislation on multiple occasions that would introduce a federal RPS. Since RPSs' main goal are to reduce carbon dioxide emissions, it is important to know the cost of the carbon abatement from these policies relative to other ways that could reduce carbon dioxide emissions.

In order to estimate this parameter, I use the policy variation in the the implementation schedule of renewable portfolio standards across states that have restructured electricity markets. Since most state RPSs can be met by renewable generation located anywhere in the wholesale electricity market, I aggregate individual state policies into region-level renewable portfolio standards. Each year, each state's RPS increases in its stringency, creating the variation that I use to estimate the the long-run supply elasticity. In my preferred specification, I estimate that a 1 percent increase in the total price received for renewable electricity (price of electricity plus the price of the renewable energy credit) results in a 2.7% increase in the supply of renewable generation.

Politicians appear to prefer using RPS policies to those of broader policies such as cap and trade or a carbon tax. Part of the attraction is likely that the costs of this method of carbon dioxide abatement are less transparent to voters. However, these policies still come with a cost. My estimates suggest that the cost of abating the last

ton of carbon dioxide from an RPS in the northeastern US to reduce emissions by 10 percent from their 2005 levels (approximately equal to a 6 percent RPS) would cost between \$50 and \$140 per ton of carbon dioxide, depending on the type of fossil generation that the renewable generation was replacing. My estimate for the cost of CO₂ abatement is more than 5 times more expensive than the maximum price of CO₂ under the regional cap-and-trade program for the electricity sector. Therefore, residents would be paying an extremely high premium for carbon dioxide abatement under RPSs, even though they appear to be more politically palatable policies.

Tables

Table 1.1: Environmental Impact of Common Fuels and Frequency of Inclusion in Renewable Portfolio Standards

	Pounds of CO ₂ per MMBtu	Percent of States that Consider Fuel as Renewable	Average Generator Size (MW)
Wind	0	1.00	27.3
Solar Photovoltaic	0	0.89	0
Biomass*	0	0.89	14.0
Solar Thermal	0	0.81	17.0
Geothermal	0	0.70	21.5
Hydropower	0	0.70	86.6, 6.0
Nuclear	0	0	999.9
Municipal Solid Waste	92	0.26	26.1
Landfill Gas	115	0.78	1.4
Natural Gas	123	0	62.3
Petroleum	166	0	18.1
Coal	215	0	222.3
Number of States with RPS		27	

CO₂ estimates taken from Energy Information Administration.

* Biofuels contain “biogenic” carbon and are not considered to add to atmospheric carbon levels. [30]

Table 1.2: Summary Statistics

	Mean	Standard Deviation	Min	Max	N
Active RPS in Region	0.556	0.498	0	1	293
RPS Requirement in Region	0.600	0.747	0	2.476	293
RPS Requirement when Active in Region	1.079	0.697	0	2.476	163
Mandatory Green Power Option in Region	0.041	0.199	0	1	293
Government Purchases of Green Power in Region	0.491	0.501	0	1	293
Public Benefits Fund in Region	0.724	0.448	0	1	293
Net Metering Laws in Region	0.724	0.448	0	1	293
Active RPS in State	0.200	0.400	0	1	1848
RPS Requirement in State	0.339	0.882	0	4.92	1848
RPS Requirement when Active in State	1.698	1.261	0	4.92	369

Table 1.3: First Stage Regression Estimates

Dependent Variable: Log(Total Renewable Electricity Price)				
	(1)	(2)	(3)	(4)
Log(Average Stringency $_{t,t+5 \text{ years}}$)	0.184**	0.268**	0.254**	0.298**
	(0.047)	(0.062)	(0.058)	(0.068)
II(Green Power Option)		0.074		0.162*
		(0.061)		(0.071)
II(Gov't Power Purchase)		0.181*		0.108
		(0.082)		(0.097)
II(Public Benefits Fund)		0.192*		0.844*
		(0.086)		(0.317)
Gov't Power Purchase (Frac. of Consumption)			0.121*	0.103
			(0.053)	(0.066)
Public Benefits Fund (Frac. of Consumption)			-0.081	-0.118
			(0.075)	(0.076)
Net Metering (Frac. of Consumption)			0.010	-0.031
			(0.035)	(0.038)
Observations	293	293	293	293
R^2	0.57	0.59	0.58	0.59
F-test that excluded instrument equal to zero	15.61	18.64	19.05	19.36

OLS estimates. Estimates include region, year, and month fixed effects as well as region specific trends. Standard errors robust to heteroskedasticity and autocorrelation using the Newey-West method. * $p < 0.05$, ** $p < 0.01$

Table 1.4: Second Stage Regression Estimates

Dependent Variable: Log(Renewable Generating Capacity)				
	(1)	(2)	(3)	(4)
$\widehat{\text{Log(Total Price)}}$	1.810** (0.520)	3.810** (0.872)	1.732** (0.412)	2.714** (0.611)
$\mathbb{I}(\text{Green Power Option})$		-0.269 (0.257)		-0.452 (0.234)
$\mathbb{I}(\text{Gov't Power Purchase})$		0.616* (0.232)		0.763* (0.238)
$\mathbb{I}(\text{Public Benefits Fund})$		0.072 (0.383)		-3.824** (0.816)
Gov't Power Purchase (Frac. of Consumption)			0.450** (0.090)	-0.144 (0.195)
Public Benefits Fund (Frac. of Consumption)			0.815** (0.141)	0.750** (0.202)
Net Metering (Frac. of Consumption)			0.424** (0.070)	0.344** (0.097)
Observations	293	293	293	293
First Stage F-statistic	15.61	18.64	19.05	19.36

OLS estimates. Estimates include region, year, and month fixed effects as well as region specific trends. Standard errors robust to heteroskedasticity and autocorrelation using the Newey-West method. * $p < 0.05$, ** $p < 0.01$

Table 1.5: Cost of CO₂ Abatement From an RPS

	Replace Coal	Replace Average Fuel
2.5% Reduction in CO ₂ Levels		
Percent Increase in Renewable Capacity	41%	68%
Cost of CO ₂ Abatement	\$12.46	\$34.82
10% Reduction in CO ₂ Levels		
Percent Increase in Renewable Capacity	163%	273%
Cost of CO ₂ Abatement	\$49.86	\$139.28

All CO₂ reductions are measured from the 2005 baseline levels, similarly to the Regional Greenhouse Gas Initiative

Figures

Figure 1.1: Statutory Renewable Requirements for Selected States

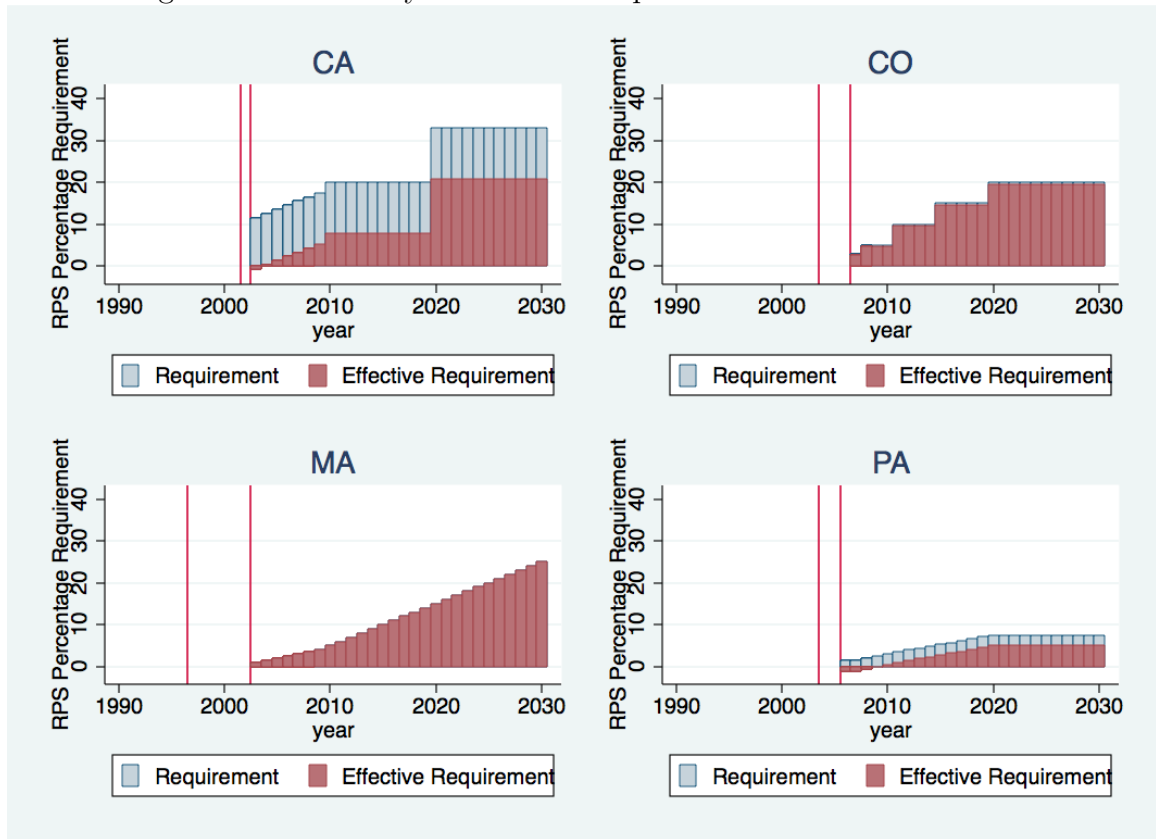


Figure 1.2: Timing of Renewable Portfolio Standard Adoption

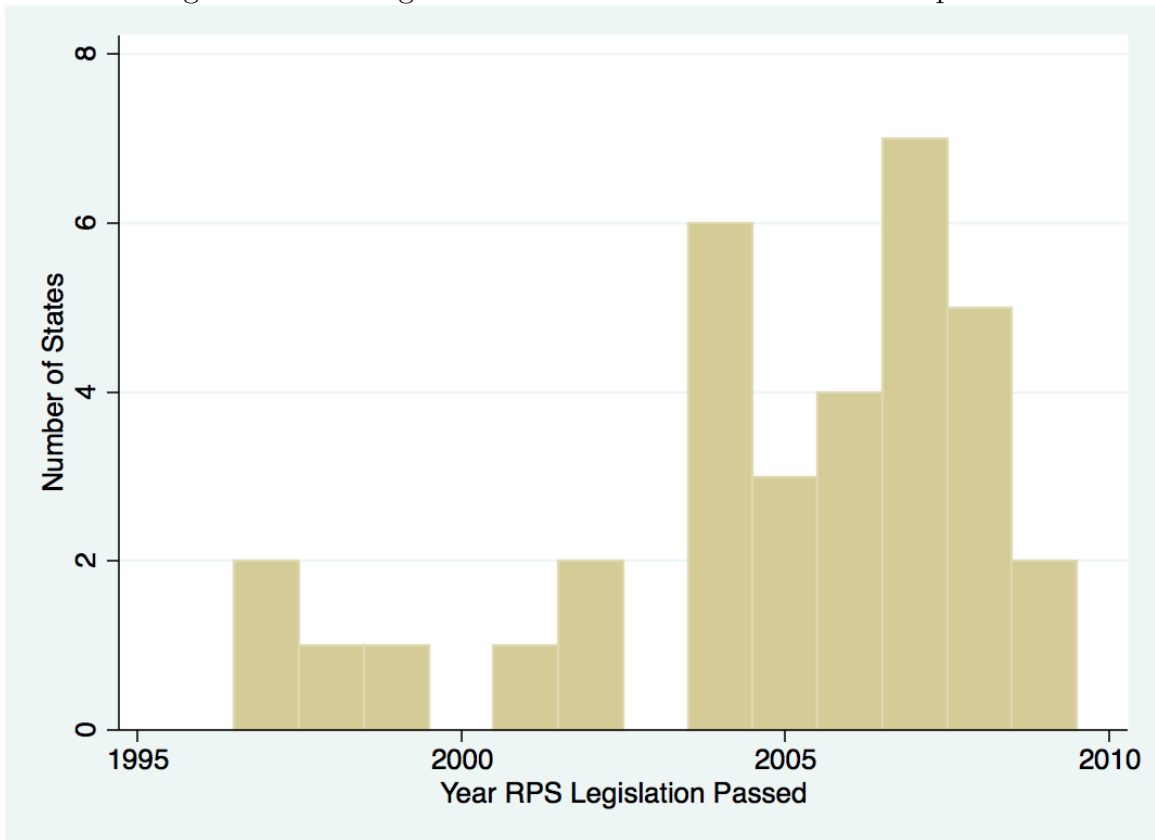


Figure 1.3: Timing of Renewable Portfolio Standard Adoption by State




Figure 1.4: Selected Wholesale Electricity Markets




Figure 1.5: Capacity for Renewable Generation by Fuel

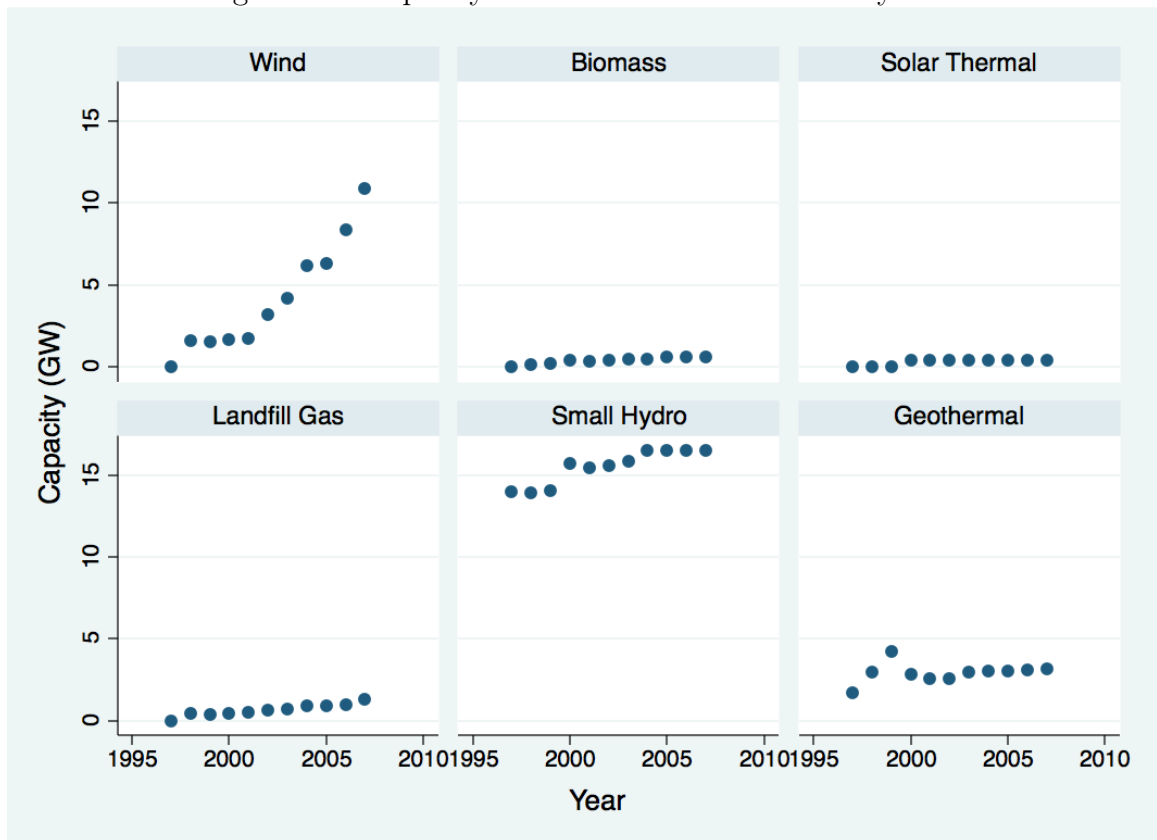


Figure 1.6: Electricity Capacity Market without RPS

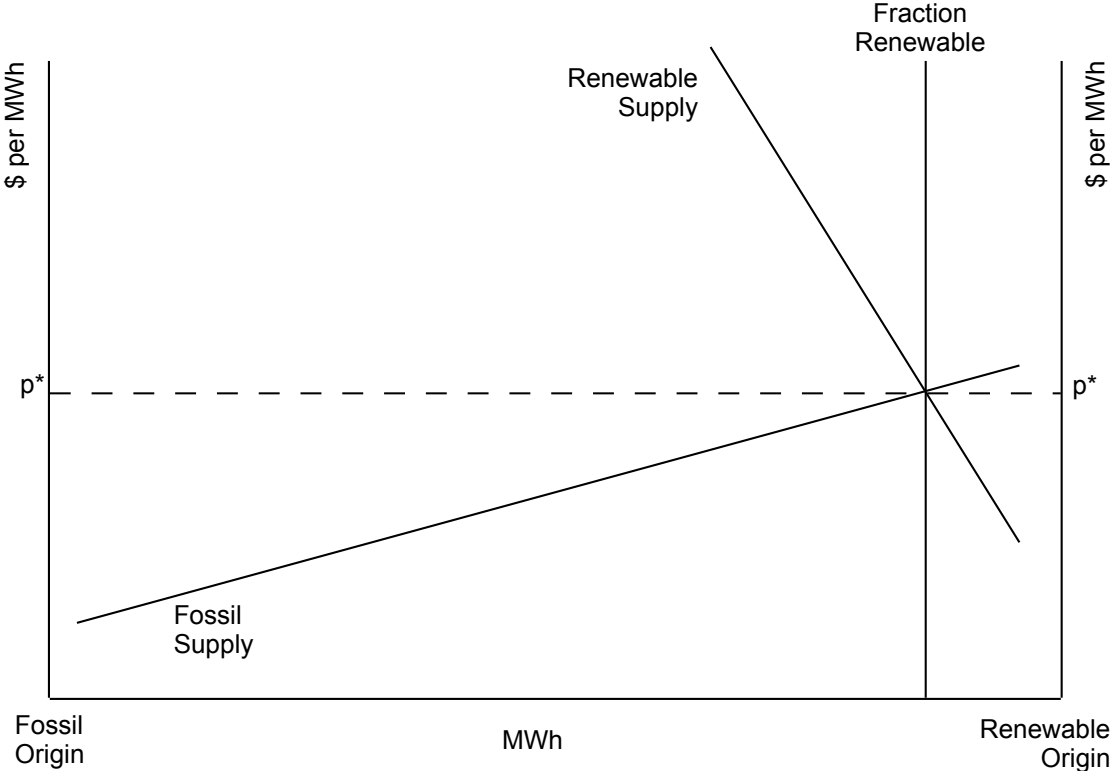


Figure 1.7: Electricity Capacity Market with RPS

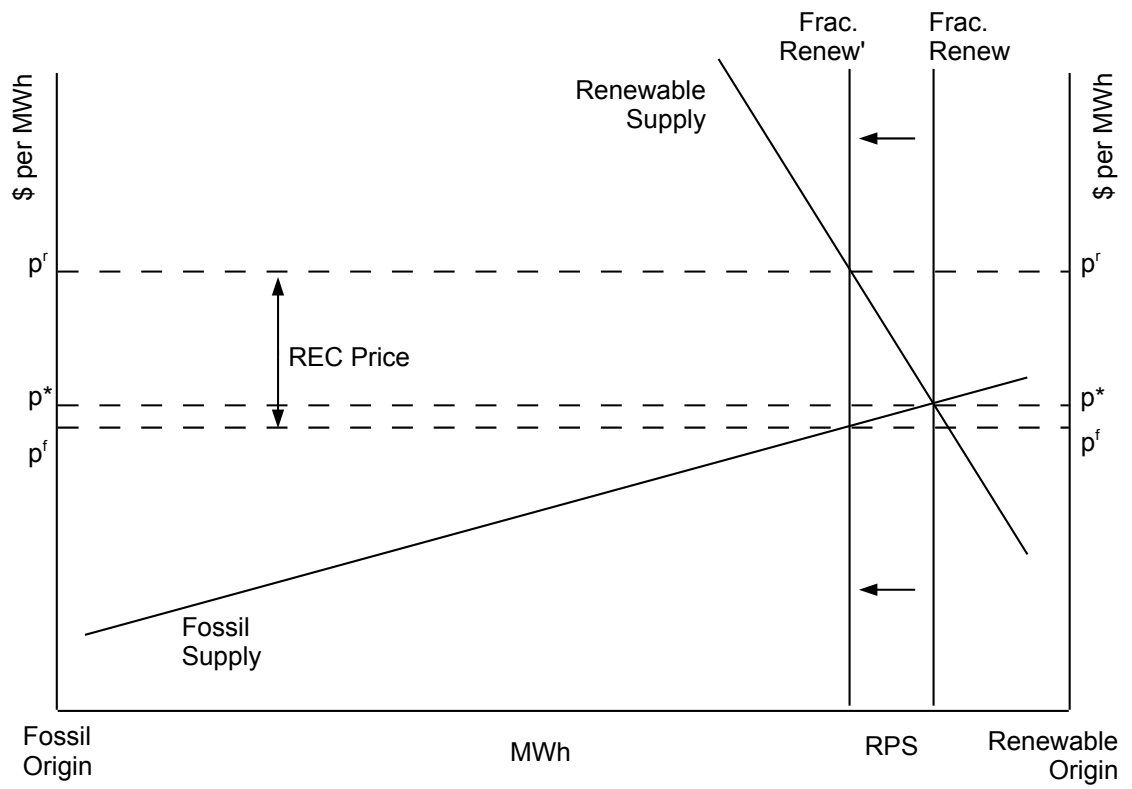



Figure 1.8: State REC Prices



CHAPTER II

Cost Effectiveness of Residential Radon Remediation with Household Mobility

2.1 Introduction

The Environmental Protection Agency (EPA) has termed exposure to residential radon "probably the biggest public health problem we have," causing from 7,000 to 30,000 lung cancer deaths per year [2]. Since 1986 the EPA has waged a series of publicity campaigns, urging all householders to test for the presence of radon and to reduce ambient levels of radon when airborne radiation from radon decay products exceeds 4 picoCuries/liter (pCi/l). [13]. If universally adopted, such a program would reduce exposure levels in about 5.7 percent of housing units, occupied by almost 5 percent of the population. The EPA has provided extensive technical analysis of and support for their recommended course of action, including cost-effectiveness analysis [2].

We re-examine the EPA's recommendations using a model that incorporates much of the important heterogeneity in the population. We find universal remediation of all houses with ambient exposure above 4pCi/l (and, possibly, lower levels) would pass a social benefit-cost test. However, we can imagine no conceivable set of circumstances under which there would be anything like general compliance with any action

level other than an extremely high one, assuming that people truly understand their individual risk and behave rationally. At 20 pCi/l, for example, an exposure level that is five times the recommended level for remediation, found in less than 0.01 percent of homes, there would not be universal compliance on the part of well-informed households.

The failure of the voluntary action-level approach arises due to the interaction of two phenomena, either one of which would cause difficulty. First, the remediation technology of choice has both a capital component (sealing, plugging, and installing fans) and an operating component (running and maintaining the fans). Households that do not expect to be in their houses very long generally will not be able to recoup the requisite capital investment unless it is capitalized into the price of the housing unit. (For reasons that we discuss below, we expect capitalization to be much less than 100 percent of the cost.) We estimate that with normal mobility, between 5 and 15% of the initial investment will be capitalized into the house price.

Second, even if operational fans were available in every house that had exposure above any plausible action level, many well-informed households would choose not to operate the fans, reflecting enormous variation in the benefits gained from doing so. The variation in benefits derives from variation in four important characteristics of households: the age, size, and smoking behavior of their members, and their subjective valuation of expected life-years saved. Large households with young smokers will derive relatively large benefits from turning on the fans; small households of elderly nonsmokers will derive essentially no benefit. Because the remediation technology requires that residents operate the system, only the most inconceivably draconian of regulations (someone checking to see that the fans are on) could assure widespread compliance with any action level.

The implication of our analysis is that policy (and the analysis of policy) must be designed to take into account residential mobility, the heterogeneity of the population,

and the consequences of radon-remediation technology for individual behavior.

Considering these factors, we conclude that the only way in which general installation of radon-reduction equipment will be undertaken is via government provision or mandatory regulation. We also conclude that, depending on the action level and on the distribution of willingness-to-pay per life-year-saved, a nontrivial fraction of the radon-remediation equipment will not be operated in a given year. Assuming that households are well informed about the risks, however, allowing them to choose whether or not to operate the equipment increases, rather than reduces, the efficiency of policies to deal with the health effects of residential radon.

This paper also adds to the economic literature on willingness to pay for pollution abatement and the capitalization of environmental amenities in housing values. A large literature use housing market transactions to examine how various pollution abatement policies such as the Clean Air Act [12] and Superfund clean-ups [15]. Generally these studies tend to find small changes in housing values after a relatively large change in environmental quality. In fact, Greenstone and Gallagher find that Superfund clean-ups “are associated with economically small and statistically indistinguishable from zero local changes in residential property values [and] property rental rates.” [15]. However, our model suggests that this result may not be particularly surprising if there is a high degree of residential mobility surrounding the clean-up site. Their results may simply reflect a belief on the nearby residents that the health risk from the polluted areas was sufficiently low if the exposure time was sufficiently short. Thus, after the clean-up we would expect to see a change in the rate of housing sales and likely a change in the distribution of ages surrounding a clean-up site.

In the next section we describe the public health and economic models we use to analyze the costs and effects of reducing residential radon exposure. In section 2.3 we describe the data we use to calibrate our models to estimate the behavior

of households. In section 2.4 we discuss the results and intuition of our simulations and in section 2.5 we discuss how varying our assumptions affect the outcomes of the model. Finally, section 2.6 concludes.

2.2 Model

To evaluate radon-reduction policies we construct a model of individual decisions about radon remediation and mobility. We use an agent based approach where each individual has full information about the health effects of the current level of radon in their house, their complete history of residential radon exposure levels, and their (exogenous) probability of moving out of their current house. Furthermore, we assume at the beginning of the simulation that each agent has done nothing in the past to remediate residential radon in their homes. This is likely similar to the radon exposure history for most of the US population since there has not been a large scale program to encourage households to test their houses for radon and to encourage radon remediation.^{1,2} In the rest of this section we will discuss the model we use to estimate the health effects of radon exposure, then discuss the options that are available to homeowners who want to reduce their radon exposure, and then the economic model of agents' behavior.

2.2.1 Effect of Radon on Health

Epidemiological studies of underground miners have documented that exposure to high levels of radioactive radon gas can cause lung cancer. The risk of lung cancer from radon can be calculated as a function of cumulative exposure over a person's lifetime. Being exposed to a higher concentration of radon in any year increases that

¹If some households have installed remediation equipment and then moved, the new occupants would also have to make the active decision to turn the installed fans on.

²If this assumption fails, it only affects each household's exposure *history* of exposure which is not the driving force in the model.

person’s chance of developing radon induced lung cancer, though the effect from a radon exposure in a particular year eventually fades over time.

There are strong interactions between radon induced lung cancer and smoking induced lung cancer. People who are current smokers not only are 22 times as likely to develop lung cancer from smoking than non-smokers, they are also approximately six times more likely than non-smokers to develop lung cancer from equivalent levels of radon exposure.

Exposure to residential radon is translated into lung-cancer mortality risk according to the model described by the Commission on the Biological Effects of Ionizing Radiation of the National Academy of Sciences [27]. This model is commonly referred to as BEIR VI.

The BEIR VI model estimates the relative risk that an individual faces for radon exposure based lung cancer as a function of their cumulative radon exposure, age, and if they have ever smoked. The exact model is specified as:

$$RR = 1 + (\Psi_{s=1}s + \Psi_{s=0}(1 - s))(\gamma_1w_{5-14} + \gamma_2w_{15-24} + \gamma_3w_{25+})$$

where RR is the relative risk of lung cancer from radon, s is an indicator equal to 1 for people who were ever smokers, $w_{x_1-x_2} = \sum_{t-x_2}^{t-x_1} w_t$ for radon levels w measured in pCi/l, Ψ is an age and smoking status specific constant and γ_1 , γ_2 and γ_3 are weights used to signify the decreased risk of radon induced lung cancer as the exposure date fades into a person’s radon exposure history. Table 2.1 shows the numeric values for the constants.³

³There are a few variants of this model described in the BEIR VI report. We use the same coefficient values as the EPA so as to make the analysis comparable.

2.2.2 Remediation Technology

The most prevalent method of remediating residential radon and therefore the only type of radon remediation we consider in this analysis is Active Slab Depressurization (ASD). ASD involves two steps. First, fans are installed in order to ventilate the radon trapped in the area below the foundation slab to the outdoors, while holes in the walls and floors are plugged and sealed. Second, the fans must be turned on, maintained, and occasionally repaired. Running the fans uses electricity and also increases the costs of heating a house [2]. Henschel [16] suggests that over 90 percent of all residential radon remediations use ASD. Initial installation of ASD systems, including testing, plugging and sealing, involves an average cost of \$1200 with a range of \$800-\$2,500.

Running the fans costs an average of about \$125 per year for electricity and increased heating costs.⁴ Annualized maintenance and testing costs, assuming testing every five years, come to another \$24. EPA assumes that the whole system lasts for 74 years, at which time it would have to be replaced. Over the 74-year life of the system, total costs come to \$5010 discounted at 3%. Note that for any plausible discount rate, the up front costs are small compared to the present value of the operating, testing, and maintenance costs.⁵

Provided that the fans are operating properly, ASD generally reduces radon exposure to an average of 2pCi/l or less, regardless of the initial level of exposure.⁶ (With

⁴The EPA estimates a range between \$75 and \$175. [2]

⁵The EPA's Technical Support Document [2] has somewhat higher costs, averaging \$1684 in upfront costs and annual costs of \$150.75. However, \$38.77 of their annual costs is annual radon testing, which we find highly implausible, and many of the more expensive technologies that are averaged into the initial costs are economically dominated by ASD. Moreover, other work [16] documents that ASD almost always provides high radon reductions, and is almost always the method of choice. Thus, we conduct our analysis under the simplifying assumption that all remediation is done via ASD.

⁶2pCi/l is the average assumed by the EPA in the Technical Support Document [2] and the amount used in this paper as the average exposure for remediated housing units. The "or less" reflects Henschel's remark that ASD "techniques provide high radon reductions, as high as 98 to 99+ percent." [16] Except for a very few houses, such reductions would imply post-remediation levels well below 2pCi/l. For our purposes, we accept EPA's estimate of remediation to 2 pCi/l.

only a few exceptions, this is true of all of the remediation strategies discussed by the EPA [2] and by Henschel [16]). To some extent, the post-remediation level of exposure depends on the characteristics of the house, but only slightly or not at all on the pre-remediation level of exposure.⁷ The implication for cost-effectiveness is straightforward: cost is approximately independent of the level of initial exposure, as is post-remediation exposure. This implies that effectiveness increases almost linearly with the initial level of exposure, so that cost-effectiveness will generally be greater the greater is the initial level of exposure. A given reduction in exposure has the same effect on health no matter what the initial level of exposure, but the higher the initial level, the more reduction can be obtained at a given cost. Note that this does not imply that only high-radon homes are worth remediating. It merely implies that the net payoff to remediation is highest in homes with the highest initial levels of radon.

In the following analyses, we adapt the EPA's conclusion that remediation will reduce exposure to an average of 2 pCi/l to the more analytically tractable assumption that all remediated housing units have exposure levels of exactly 2 pCi/l.

2.2.3 Behavioral Model

In this section we describe the economic model we use for a household's decision about whether to take action to remediate the ambient radon concentration in their house. Our model allows households to be forward looking about their remediation decisions and allows households to completely understand the health consequences of their decisions.

Since all residents of a particular household are necessarily exposed to the same radon concentrations, our model takes households as the relevant level of analysis.

⁷See Exhibit F-2 of the EPA's Technical Support Document [2]. The Exhibit shows that the cost of ASD generally varies with technical characteristics of the building foundation, but for all action levels below 20 pCi/l cost does not vary with action level. Above 20 pCi/l, only "Hard to Fix" houses with basements, accounting for 16.5 percent of houses above 20 pCi/l (about 0.6 percent of houses that would be remediated at an action level of 4 pCi/l) cost more than other houses that use ASD, and the difference is \$221.54 of initial cost.

Therefore, we assume that the household maximizes over all members and that the utility of the household is additive across periods. Thus, a household of two people incur twice the health costs of a household with one individual. More restrictively, we assume that the household has a particular smoking history instead of individuals within the household having a particular smoking history. This means that all adults in a household are current smokers, former smokers, or have never smoked. Moreover, we assume that all households have two adult members, both of whom are the same age. In addition to two adults, households contain between 0 and 5 children according to the frequency of households with children in the US population.

We model household behavior assuming that agents have full information about the health effects of radon exposure according to BEIR VI, the distribution of radon in the current housing stock, the radon level in their current house, and their complete radon history. Households are forward looking with respect to their probabilities of leaving their current house and the likelihood of dying before next period.

At the beginning of each period, a household observes whether or not its current house already has remediation equipment installed, the radon level in their current house, as well as their age, smoking status (current smoker, former smoker, or never smoker), the number of children living in the house, and everyone's radon exposure history (assumed to be the same for each member of the household).

After observing all of this, the household can choose whether or not to install radon remediation equipment in the house and pay the cost of the installation, k , or choose not to remediate the radon. If a household chooses to install remediation equipment or the house already has remediation equipment installed, the agent then chooses whether or not to use the remediation equipment. If the household uses the remediation equipment and pays the associated cost, c , its radon exposure for that period is assumed to be 2 pCi/l. If it chooses not to use the remediation equipment, its radon exposure for that period is the ambient radon concentration in the house.

See Figure 2.1 for a decision tree representing the choices.

The household will move in the next period with probability, δ . We assume that a household's decision to move is independent of its previous choices of purchasing radon remediation equipment and the concentration of radon in their house. However, if moving is endogenous to a household's decision to remediate the radon in the house, any investment that has already been made in remediation equipment is a sunk cost and should not influence its decision to move or not in the current period.

We allow our model to take on different values for the probability of moving, δ , and dying, λ , for different types of people. We define types of people, θ , by their age, radon exposure history, number of children, and smoking status.

This model leads to two Bellman equations of the form:

$$V_{\theta}^R(r, a, z) = \max\{\text{operate remediation fans, do not operate remediation fans}\} \quad (2.1)$$

and

$$V_{\theta}^{NR}(r, a, z) = \max\{\text{install remediation, do not install remediation}\} \quad (2.2)$$

where $V_{\theta}^R(\cdot)$ is the value function for people who already have radon remediation equipment installed in their house and $V_{\theta}^{NR}(\cdot)$ is the value function for people who do not already have remediation equipment installed.

In order to specify what $V_{\theta}^R(\cdot)$ and $V_{\theta}^{NR}(\cdot)$ look like, we will first need to introduce some more notation. Each period (assumed to be a year) households get flow utility, u , from their house and discount the future using a 3% discount rate. Both adult residents of each household are the same age, a , have between zero and five children who are all 10 years old, and have the same smoking histories, s , (current, former, never smokers). We assume that no one begins smoking before age 18 and that the distribution of current, former, and never smokers are the population average for that category at each age. Together age, number of children, and smoking status define a

type, θ . Each household type, θ , has an exogenous probability of moving each year, δ_θ , and a probability of dying during the year, λ_θ .

Each household knows the radon concentration of their current house, r , measured in picoCurries per liter (pCi/l), as well as the distribution of radon in the housing stock, $f(r)$.⁸ We then define each household's radon exposure history as a 25×1 column vector, $w = w_a, w_{a-1}, w_{a-t}, \dots, w_{a-25}$, where each entry is the cumulative sum of radon exposure up to time $a - t$. This vector of radon exposure history, w , along with current radon exposure, smoking status, and age determine the increased probability of lung cancer from radon exposure, $h(r, s, a, w)$, from the BEIR VI model described above. Once we multiply the probability of lung cancer by the value of a life-year, we can monetize the risk to the household. This monetized value allows us to calculate if that household finds it worthwhile to install and/or operate radon remediation equipment in its house.⁹

We assume that radon remediation equipment has a capital and installation cost, k , and a yearly cost of running the fans and extra heating costs, c . We then define the indicator variable, $f = 1$ if the household chooses to run the remediation fans. With positive probability, p , if a household changes houses, that house will already have remediation equipment installed in the house. In equilibrium, p will be higher for homes with a higher initial radon concentration, but this doesn't affect households' behavior since we assume households' choose a house independently of the radon concentration.

This is identical to assuming that households do not sort into houses based on the radon concentration. We think this is a plausible assumption for two reasons. First, to the extent that searching for a house is costly (in terms of time and/or money)

⁸The density of radon concentrations (before remediation) in the housing stock can be approximated with a log-normal density, with mean 1.25 and geometric standard deviation, 3.11; $r \sim \log N(-0.42, 1.13)$. [28]

⁹Our central case for the value of a life-year is \$300,000. We do a sensitivity analysis by using life-year values between \$100,000 and \$500,000.

households are unlikely to find a house that matches their preferences across all of the search parameters. Therefore households will need to choose which variables in a house are most important to them and weight each appropriately. Secondly, even in a scenario where search is costless, the housing market in any particular geographic region is unlikely to be thick enough in order for all households to be able to optimize their housing choices with respect to all of their choice variables, again necessitating a weighting of different housing attributes. We believe that it is unlikely that the ambient radon concentration is likely to receive much weight in a household's utility function, particularly given that remediation costs are relatively low compared to the purchase price of the house. This lack of perfect matching creates a situation in which radon remediation equipment will be partially capitalized into the house price.¹⁰ Figure 2.2 shows the probability, in equilibrium, that a house of a particular radon concentration will have remediation equipment installed.

Households that have remediation equipment installed will be able to recoup part of their capital expenditures on the remediation equipment through an elevated resale value of the house. One of the important features of our model is that we can estimate both the capitalization of remediation equipment into the housing value, π , and the probability that a house will have remediation installed, p in an internally consistent manner so that households make remediation decisions in part based on these parameters. We will describe how we solve for these parameters in Section

¹⁰If however, households perfectly sorted into houses based on radon concentrations, we would not expect to see any capitalization of radon remediation equipment, nor see any remediation equipment installed. The intuition for this is that only those households that have a low cost of radon exposure will move into high radon houses. Some households will have a low cost of radon either because the household is old and will not experience the full effect of the current radon exposure before they die, because it is a non-smoking household and therefore have a smaller health effect of radon exposure, or because it is a young household that is likely to move again soon. If these three groups of households with a low cost of radon exposure are not big enough to inhabit all of the high radon houses we would see some capitalization of remediation equipment. However, the size of the population that falls into one of these three groups far outweighs the number of houses with a high radon concentration. Because we are assuming household are choosing houses independently of radon concentrations, we will overestimate the capitalization of remediation equipment and the fraction of houses that are remediated.

2.2.3.2.

Given this notation, we can define $V_\theta^R(\cdot)$, the value function for households that already have remediation equipment installed as:

$$\begin{aligned}
V_\theta^R(r, a, w) = & \max_{f \in \{0,1\}} u - \underbrace{\left(h(r, s, a, w, E[r', w']) (1 - f) \right)}_{\text{health cost of not running fans}} - \\
& \underbrace{h(2, s, a, w, E[r', w']) f}_{\text{health cost of running fans}} - \\
& \underbrace{cf}_{\text{cost of running fans}} + \underbrace{\lambda_\theta (1 - \delta_\theta) \beta V_\theta^R(r, a', w')}_{\text{value of staying in the same house}} + \\
& \underbrace{\lambda_\theta \delta_\theta \beta p E_r[V_\theta^R(r', a', w')]}_{\text{value of moving to a house with remediation}} + \\
& \underbrace{\lambda_\theta \delta_\theta \beta (1 - p) E_r[V_\theta^{NR}(r', a', w')]}_{\text{value of moving to a house without remediation}} + \\
& \underbrace{(1 - \lambda_\theta)(0)}_{\text{value of dying}} \tag{2.3}
\end{aligned}$$

and $V_\theta^{NR}(\cdot)$, the value function for households that do not already have remediation equipment installed (either because the household has installed it in a previous period

or moved into a house that already had it installed) as:

$$\begin{aligned}
V_{\theta}^{NR}(r, a, w) = & \max \left\{ \left(u - h(r, s, a, w, E[r', w']) + \underbrace{\lambda_{\theta}(1 - \delta_{\theta})\beta V_{\theta}^{NR}(r, a', w')}_{\text{value of staying in the same house}} \right) + \right. \\
& \underbrace{\lambda_{\theta}\delta_{\theta}\beta p E_r[V_{\theta}^R(r', a', w')]}_{\text{value of moving to a house with remediation}} + \\
& \left. \underbrace{\lambda_{\theta}\delta_{\theta}\beta(1 - p) E_r[V_{\theta}^{NR}(r', a', w')]}_{\text{value of moving to a house without remediation}} + \underbrace{(1 - \lambda_{\theta})(0)}_{\text{value of dying}} \right), \\
& \left(\max_{f \in \{0,1\}} u_{\theta} - (h(r, s, a, w, E[r', w'])(1 - f) - \right. \\
& \left. h(2, s, a, w, E[r', w'])f) - cf - \pi k + \underbrace{\lambda_{\theta}(1 - \delta_{\theta})\beta V_{\theta}^R(r, a', w')}_{\text{value of staying in the same house}} + \right. \\
& \underbrace{\lambda_{\theta}\delta_{\theta}\beta p E_r[V_{\theta}^R(r', a', w')]}_{\text{value of moving to a house with remediation}} + \\
& \left. \left. \underbrace{\lambda_{\theta}\delta_{\theta}\beta(1 - p) E_r[V_{\theta}^{NR}(r', a', w')]}_{\text{value of moving to a house without remediation}} + \underbrace{(1 - \lambda_{\theta})(0)}_{\text{value of dying}} \right) \right\} \quad (2.4)
\end{aligned}$$

Note that for any variable x , x' is next period's realization of x . The age and smoking specific probabilities of moving, δ , and dying, λ , can be found in Table 2.2 and Figure 2.3 respectively. Using value function iteration we can solve the system of equations and then find each household's optimal remediation policy function.

2.2.3.1 Types of Agents

Our model has three dimensions that define each type of agent without considering the number of children in the household: age (between 20 and 110, inclusive), smoking status (current, former, and never smokers), and radon exposure history¹¹ (including current period unremediated ambient radon level for a total of 26 dimensions). This gives us a total of 28 state variables to integrate over to solve the model for each household size.

¹¹We use a 50 point grid of evenly spaced radon levels between 0 and 20 pCi/l. We have experimented with increasing the number of grid points to 100 which has very little effect on our results.

In order to solve our model we need to give each of our agents a history of radon exposure. We first assign an agent to each radon \times age \times smoking status cells. (This gives us a complete array of types in the $r \times a \times s$ dimensions.) We then impute a history, w_x , for each of these agents conditional on finding the agent in that cell, using age specific probabilities of having moved in the past and the level of radon in the agent's current house.

For instance, consider a 46 year old smoker who is currently living in a house with a radon level of 10 pCi/l. With probability 0.907 they were exposed to that level of radon when they were 45 (and with probability 0.093 had a randomly drawn other radon level). When the agent was 44 they were exposed to a radon level of 10 pCi/l with probability 0.773 ($=0.907 \times 0.852$) from being in the same house, with probability 0.079 ($=0.093 \times 0.852$) were exposed to the radon level from the house they may have moved into when they were 45, with probability 0.134 ($=0.907 \times 0.148$) received a new, random radon draw when they were 44 because they moved that period, and with probability 0.014 received a new, random radon draw when they were both 45 and 44 by moving two years in a row. The number of possible histories grows exponentially with the number of lagged state variables (dimension of w). This gives us a total possible number of types of agents of $(\dim(w))^{\text{(grid size)}} \times \dim(\text{age}) \times \dim(\text{smoking status}) \simeq 8 \times 10^{44}$ for each sized household. This is not a computable problem given current computing technology.

The typical solution to this problem is to reduce the number of state variables in the model. However, due to the fact that the BEIR VI model attenuates the effect of cumulative past radon exposure over 25 years and our assumption of persistence of radon concentrations in a household's history, we are not able to condense the number of state variables for the history of radon exposure concentrations, $\dim(w)$.

In order to get around this problem, we sample for the possible histories available in the population and integrate over our subsample of histories (instead of integrating

over all possible histories).¹² We have run the simulation with a number of historical types¹³ between 100 and 1,000¹⁴ and have found that the results do not seem to be terribly sensitive to the number of historical types. We report the results of the simulations with 1,000 historical types. Moreover, if people do not actually know their complete radon exposure history, our simulation is capturing much of the important variation by simply having a high likelihood of the previous few periods of radon exposure correct since most of the 8×10^{44} variations come from the exposure “tree” splitting further in the past.¹⁵

2.2.3.2 Calculation of the Capitalization Value of Remediation Equipment

In our economic model of households’ behavior, both the capitalization of remediation equipment and the probability that an agent moves into a house with remediation equipment already installed are determined endogenously.¹⁶ In order to make these parameters endogenous, after each step in the value function iteration, we calculate the proportion of houses that will have remediation and the probability that each agent will move into that house next period. We assume that agents do not choose houses based on either the ambient radon level or the existence of remediation equipment at a particular house. We do not believe this is a particularly restrictive assumption because the value of other amenities from a particular house are likely to be far greater than the costs of installing remediation fans. Therefore, in the next step

¹²We assume that none of our agents engaged in any radon remediation behavior until the date that the simulation started. Thus our distribution of radon concentrations across agents histories reflect the distribution of radon in an unremediated housing stock.

¹³A historical type is a particular moving history for each type (radon×age×smoking status×history×number of kids). Therefore, for each historical type we are actually solving the dynamic programming problem for $50 \times 91 \times 3 \times 6 = 81,900$ types.

¹⁴We stop increasing the number of historical types at 1,000 due to computational size since the matrices needed to be held in memory by the computer reach 1.1Gb.

¹⁵Table ?? lists all of the behavioral model parameters with a brief summary of each.

¹⁶In order to ensure that our value function iteration converged in under 1,000 iterations, we rounded the capitalization percent to the third decimal place.

of the value function iteration agents have the correct beliefs about these parameters and make their decisions accordingly until we find the fixed point.

2.3 Data

In order to make our model operational, we use data on age specific death rates, baseline lung cancer rates for current smokers, former smokers, and people who have never smoked.

We also use data from the U.S. Census Bureau on age specific probabilities of moving to inform our model [10]. From the age of 20 until 84, the probability of moving decreases for each age group, from a maximum of 35.5% of 20-24 year olds moving every year to a minimum of 4.3% of 65-84 year olds moving every year, with just a slight increase in the probability of moving for people 85 years old and over. All of the probabilities can be seen in Table 2.2.

We calibrate our model to match the age distribution in the US by using data from the 2006-2008 American Community Survey [11]. A density plot of this data is shown in Table 2.4.

2.4 Results

In this section we present the results in two separate sections. The first section describes the results of our simulations while varying the value of a life year between \$100,000 and \$500,000. Throughout section 2.4.1 we maintain the assumption that each household is comprised only of two adults (and no children). In section 2.4.2, we present the results of our simulations after relaxing the assumption that each household contains only 2 adults and no children. We present our results for simulations that consider households to have between zero and five children in proportion to the US population distribution of household size. However, throughout section 2.4.2 we

fix the value of a life year at \$300,000.

2.4.1 Simulations without Children

We present the results of our simulations in Tables 2.4 and 2.5. The first line of Table 2.4 shows the capitalization percentage for each of our simulations, varying both the value of a life-year and the number of household radon exposure histories we consider. As mentioned above, our preferred estimate is from the simulation with 1,000 household radon exposure histories. Within this row, the capitalization values vary between 5% and 15% depending on the assumed value of a life-year. The literature typically considers a life-year to be worth approximately \$300,000. Under this assumption the capitalization value of remediation equipment is 9.5% with 0.3% of houses having remediation equipment installed in them. To add some context, if houses with the highest 0.3% of radon concentrations had remediation equipment installed, houses with ambient radon concentrations of 13.4 pCi/l and above would have remediation equipment installed.¹⁷

However, as displayed in Figure 2.6 we can see there is significant variation in household's remediation decisions based on their age and smoking status. Figure 2.6 shows a contour plot of the policy functions for a set of agents with the same radon exposure history. Each point on the plot represents an agent with a particular age and current radon exposure, while holding the agent's smoking status and radon exposure history constant. The stalactite-shaped area protruding from the top of the each of the graphs is the boundary between where an agent chooses to install and/or run remediation equipment and fans in their house or not.¹⁸ The top row of contour plots show the policy functions for agents who live in houses with remediation equipment

¹⁷A full 7% of the housing stock has a radon concentration above the current EPA recommended action level of 4 pCi/l.

¹⁸We also allow for the possibility that an agent would want to install remediation equipment in their house but not pay the variable cost involved with running the fans. Our simulations suggest that this is never an optimal choice for an agent.

already installed, either because they have installed it in a previous period or because the house they live in had it installed when they moved in. The bottom row of contour plots show the policy functions for agents who live in houses that do not already have remediation equipment installed.

We first turn our attention to the bottom row of the contour plots. This row of contour plots show the policy functions for households who live in houses that do not have radon remediation equipment already installed. Thus, if the household wants to influence the concentration of radon in their house they will need to install the remediation equipment at a cost of \$1,200 and then also pay the annual cost of running the fans of approximately \$125.

The stalactite in the two right-most figures delineates the age \times radon concentration combinations where households find it worthwhile to install radon remediation equipment and operate it. The size of the stalactite increases as we move from the left picture in the figure to right picture of the figure. This shows that never smokers are least likely to install remediation equipment and/or run the fans and current smokers are the most likely to remediate the radon concentrations in their house. This fits the fact that the largest health benefits of remediation accrue to current and former smokers.¹⁹

We also see that the very young and very old people tolerate substantially higher radon levels before purchasing or operating radon remediation equipment. These results come from two different sources of variation. The oldest people in our simulation do not install or use radon remediation equipment because there is no health benefit that accrues to them to offset the cost of installation or use. BEIR VI models radon exposure as having no negative health implications until 5 years after the year of exposure. Thus, in the limiting case, 105-110 year olds will never have any adverse

¹⁹Though the BEIR VI model does not distinguish between current smokers and former smokers, these two groups have different baseline levels of lung cancer. These baseline levels drive the difference in actions between the two groups.

health effects of radon exposure during those years of their life since all of our agents are projected to die before they reach 111 years old. This effect attenuates rather rapidly as we examine the optimal policies of agents younger than 105. However, this is the effect that is driving the steep slope of the boundary between the two policies on that side of the contour plot.

On the left side of all of the graphs we see that young people tolerate substantially higher radon levels before purchasing or operating radon remediation equipment. This is due to the fact that young people move frequently and therefore will reap very few benefits from installing remediation equipment. A young person (age 20-24) living in a high radon house has a $\frac{1}{3}$ chance of moving next year whereby they will be living in a house that has an expected radon level of 1.25 pCi/l. Thus, in expectation, they are likely to have a one or two year spike in their radon exposure which will be less detrimental to their health than if they were exposed to that concentration of radon for the next 10 years. However, since the probability of moving falls by over 300% between the ages of 29 and 45, we also see that the incentive to install and/or operate remediation equipment increases dramatically with age.

Tables 2.4 and 2.5 display the trough point for each of contour plots (averaged across all radon history types). We see that the trough generally occurs near the age of 50 with an action level greater than 5 pCi/l depending on the smoking status of the household. These households are the most likely to directly benefit the most from their investment in remediation equipment since they only have a 1 in 10 or 1 in 15 chance of moving in an particular year. Thus, the capitalization value of the remediation equipment plays a substantially smaller role for these households than the younger households. Moreover, these households are likely to live long enough to capture the most of the health benefits of reducing their radon exposure.

We now turn our attention to the top row of contour plots in Figure 2.6. This row shows the contour plot of the policy functions for households who live in houses that

already have remediation equipment installed. This means that in order to decrease the radon concentration in their house, these households only need to pay the cost of operating the fans, approximately \$125. Unsurprisingly, younger households with houses that have lower concentrations of radon operate their fans than if they also needed to install them. In fact, we see that never smokers between the ages of 40 and 60 who have houses with extremely high concentrations of radon find it worthwhile to operate the remediation equipment in contrast to never smokers without equipment already installed.

These simulations all assume that there is no sorting on radon concentrations in the housing market. To the extent that the housing market is efficient at sorting people such that those with the lowest willingness-to-pay for radon reduction (in the extreme, mobile, nonsmoking elderly) tend to find the highest-exposure houses, there would be essentially no capitalization, even at high levels of exposure. Perfect sorting of this kind would also obviate radon exposure as a policy problem, because exposure would do relatively little damage. Even imperfect sorting, which we would expect to see in the housing market, would further reduce capitalization, thereby weakening the incentives for households to install remediation equipment.

2.4.2 Simulations with Children

We now turn our attention to our simulations where we consider the more realistic scenario where both children and adults in the same household. Table 2.6 shows the capitalization of the remediation equipment into housing prices and the fraction of houses that have remediation equipment installed. There are two complementary forces that make these simulations different from those without children in the household. First, since the cost of remediation does not depend on the number of people in a household, the net benefits of remediation increase for each adult age for a given radon concentration. Therefore, the threshold radon concentration for each age will

decrease. Because these higher net benefits create a larger demand for remediation equipment to be installed, the capitalization of the remediation equipment increases causing the threshold radon concentration to fall even more.

Comparing these numbers to the above results we see that accounting for the benefits from remediation that accrue to the children in a house increases both the capitalization percentage, from 9.5% to 10.9%, and the fraction of houses that have remediation equipment installed. This is an intuitive result since the cost of remediation does not change with the number of people in the household but the benefits increase at an approximately proportionally to the number of people in the house.

Table 2.7 shows the trough points for each of the contour plots for households with varying numbers of children. In general, these trough points occur at a lower radon concentration than the previous simulations. Again, this pattern is intuitive since there are more people in the household who are accruing the benefits of the remediation, making the decision to remediate beneficial at lower radon concentrations.

Examining the policy functions in the contour plots displayed in Figures 2.7 - 2.12 we see that households of all ages find it beneficial to remediate at lower radon concentrations the more children they have. Noticeably, even households who are 105 years or older find it beneficial to run the remediation equipment with as few as two children living in the house, though the adults receive no benefit from the remediation. We can see the threshold radon level decrease in each of these policy functions as the number of children in each household increases until there are 5 children in the household and the threshold radon concentration is approximately 10 pCi/l.

2.5 Discussion of Policy Implications

We have identified two reasons why, if people are well-informed and behave rationally, an EPA-style policy is unworkable. Only one of these – the positive externality

associated with installing remediation equipment – is realistically amenable to policy intervention. The second reason is that many households will choose not to operate a remediation system. If people are rational and well-informed as in our model, however, this is not really a policy problem at all, provided that households who would derive net benefits from operating radon remediation equipment live in residences that have remediation installed. The reason for this is straightforward: if well-informed households choose not to operate remediation systems, there is no market failure; we can infer that operating the systems would not be cost-beneficial.

There is no public (and virtually no private) health problem if a household consisting of elderly nonsmokers chooses not to operate the fans. Given any household's value of a life-year, if cost per life year saved in that household exceeds its value, economic efficiency is enhanced when the (well-informed) household chooses not to run the fan. The marginal savings exceed the marginal benefit. Thus the public health problem presented by residential radon involves inducing the installation (and repair) of ASD systems, not the operation of such systems.

Of course, if households are in some combination irrational or ill-informed, this straightforward theorem from revealed preference is not applicable. We speculate on these possibilities later in this section, but retain for now the assumption that *homo economicus* is deciding whether or not to turn on the fans.

The positive externality from installation of radon remediation equipment can be dealt with through a regulatory strategy, through public provision of radon remediation equipment, or by providing monetary incentives for private provision. Public provision of remediation equipment or providing monetary incentives essentially changes the capitalization value of the equipment in our model. Public provision of remediation equipment would correspond to a 100% subsidy with partial monetary incentives resulting in an increased capitalization percentage. At an action level of 4 pCi/l (the EPA's guideline), public provision of remediation equipment to 6,000,000

eligible housing units at a cost of \$1,200 each would cost approximately \$7.2 billion. It seems unlikely, in the current political climate, that such expenditure would be undertaken publicly. Private monetary incentives would also involve budget costs. Such incentives would also be very ineffective in the short run, given our estimate that over four-fifths of eligible households at any given time would obtain no surplus from having access to a remediation system - that is, they would not even turn on the fans.

Notice that even if the remediation equipment is completely paid for by the government, many households with radon concentrations above 4 pCi/l would not operate the equipment. This can be observed since a complete subsidy is identical (in our model) to all households already having remediation equipment installed (These policy functions are displayed in the top row of Figure 2.6). Never smokers rarely operate the fans and even many former smokers do not find it worthwhile to operate the fans. Our simulations suggests that the only subpopulation that is likely to operate the fans at the EPA's guideline of 4 pCi/l are middle aged smokers.

2.5.1 Alternative Scenarios

Up to this point we have only considered models where all households had complete information about the health effects of radon exposure, the capitalization of an investment in remediation equipment, and their probability of moving. We now consider two variants of the models.

The first variant of our model we consider is a situation where investment in remediation equipment is not capitalized into the housing value at all. If our model is incorrect to assume that every one has full information about the health effects of radon (as it likely is), we are also likely to misestimate the capitalization of remediation. If the population systemically underestimates their risk of radon induced lung cancer, our capitalization estimates will be too high. In the limiting case, where the

population either assumes they have no risk of radon induced radon lung cancer or have no knowledge of the risks (and therefore assumes there is none) there will not be any capitalization of remediation equipment. Figure 2.13 shows the policy functions of a household without any children in this scenario.

These policy functions look quite similar to those from the baseline scenario where there is full information about the radon health costs.²⁰ The stalactites are slightly smaller than the baseline scenario but there is little difference. This is a result of the fact that capitalization in the baseline scenario is approximately 10%, so even when fully informed, households are already bearing most of the capital cost of remediation.

The second variant of our model we consider is a situation where our households never move from the house they currently occupy. While this situation is clearly not something that would occur, most cost benefit analyses assume that all of the benefits of remediation equipment accrue to the current homeowner. In this situation, the externality associated with the installation of remediation equipment is internalized because the household will not move and receives all of the benefit of the remediation equipment. In this model we assume that households have moved at the average probabilities before the simulation begins, but then once made aware of the risk of radon, never move again.

The policy functions shown in Figure 2.14 come from this model, again for a household without any children.²¹ We can see that more households choose to install remediation equipment, particularly among younger cohorts. The substantial increase in installation among younger cohorts is due to the fact that they have a shorter radon exposure history and have the most to gain if they do not move in the future. We also see that, even at high levels of radon a substantial proportion of the population will

²⁰Clearly, only the decision of installing remediation equipment is affected by forcing the capitalization of the investment to be zero. If there are already remediation equipment in the house, the household's decision does not change.

²¹Again, this change to the model will not have any effect on household decisions to operate the fans once they are installed, so the top row of contour plots do not change.

not install remediation equipment, notably all households that have never smoked.

2.5.2 Information Requirements

To be effective, the "realistic" policy approach discussed in this section requires that householders be well-informed about the radon-related risks that they face and about the effectiveness of radon remediation systems in reducing that risk. A regulation requiring that remediation systems be installed will have no health benefits unless households turn on the fans. Our analysis assumes that households have sufficient information to make that decision.

Making such information generally available will surely require a costly public information campaign. Implicitly, the analysis that we have conducted here supposes that such a campaign would be inexpensive relative to the costs of remediation themselves (a reasonable assumption) and effective (an assumption that may or may not be reasonable). Thus, before pursuing any policy of public provision or subsidization of remediation equipment we would propose extensive test-marketing and experimentation of advertising protocols. The object would be to see what fraction of households, at what costs, would become well-informed about the costs and consequences of operating radon remediation systems. Note that if households systematically believe that such systems are substantially less effective than they are, a policy that promotes general availability of such systems will fail to meet a benefit-cost test. Thus the marketing research that we call for here is a fundamental prerequisite of a successful remediation policy.

Subsequent to such a study, it would be possible to reliably predict who would operate the systems were they available, and thus to reliably evaluate the policy of "remediate on transfer" that looks so promising based on the rational-behavior analysis we have reported here.

2.6 Conclusions

Because of mobility, remediation of residential radon exposure has all the earmarks of a classic public health problem: the benefit to the community of installing remediation equipment considerably exceeds the benefit derived directly by those members of the community undertaking the requisite action. Our results suggest that mobility could be a very significant barrier to reducing residential radon exposure because much of the initial investment in remediation equipment is not likely to be capitalized into housing prices. Our simulations suggest that in the case that everyone is well informed about the health risks of radon exposure, approximately 11% of the initial investment will be reflected in the resale price of the house, and only about 0.4% of houses would have remediation equipment installed.

Assuming that households are well-informed about the health risks associated with radon exposure, our simulations suggest that very few households will find it in their best interest to follow the EPA's guideline to install remediation equipment at a concentration of 4 pCi/l or above. However, if the remediation costs are heavily or completely subsidized, an action level between 4 and 6 pCi/l is justified for both current and former smokers and households with many children, though households who have never smoked and have few children are unlikely to find it worthwhile to operate remediation equipment even if capital costs are fully subsidized.

In the absence of a capital subsidy, very few households will find it worthwhile to install and operate remediation equipment. Our simulations suggest that among childless, former smokers only households between the ages of 40 and 60 in houses with a radon concentration above 12 pCi/l will install remediation equipment. The among childless, current smokers, the age range approximately doubles to 30-85 years old and the required radon concentration falls to approximately 8 pCi/l. Most strikingly, childless households who have never smoked will never find it worthwhile to install radon remediation equipment at any concentration of radon less than 20 pCi/l or age.

These ranges age ranges increase and threshold radon concentrations decrease as the number of children living in a household increase. Our simulations suggest that in households of former smokers with two children, households between the ages of 30 and 70 in houses with a radon concentration above 10 pCi/l will install remediation equipment. Households with two children of people who have never smoked between the ages of 45 and 55 find it worthwhile to install remediation equipment only at radon concentrations above 17 pCi/l.

Our simulations suggest that there is important dimensions of heterogeneity both across ages, smoking histories, and household size to consider when designing a policy toward residential radon. While previous studies have examined the aggregate costs and benefits of a homogeneous policy our results suggest that a targeting of policy at subpopulations, particularly smokers and large households, would increase the benefit cost ratio substantially.

Tables

Table 2.1: Constants used in the BEIR VI model of radon related lung cancer mortality

γ_1		1.00
γ_2		0.78
γ_3		0.51
$\Psi_{s=0}$	if age < 55	0.1536
$\Psi_{s=0}$	if 56 < age < 65	0.0876
$\Psi_{s=0}$	if 66 < age < 75	0.0446
$\Psi_{s=0}$	if 76 < age	0.0138
$\Psi_{s=1}$	if age < 55	0.06912
$\Psi_{s=1}$	if 56 < age < 65	0.03942
$\Psi_{s=1}$	if 66 < age < 75	0.02007
$\Psi_{s=1}$	if 76 < age	0.00621

Table 2.2: Percent of Population that Moved Residences by Age

Age	Percent Moved
20-24	35.2
25-29	32.4
30-34	22.0
35-44	14.8
45-54	9.3
55-64	7.0
65-84	4.3
85+	4.7

Source: US Census Bureau

Table 2.3: Parameters of Economic Model

u	Flow utility from housing
s	Smoking status (current, former, or never smoker)
a	Age (20-110)
r	Radon concentration, before remediation, in current house (0-20 pCi/l)
w	Vector of 25 years of cumulative radon exposure
f	Indicator variable = 1 if fans are operated in current period
c	Cost of operating the fans for 1 period; \$125
π	Capitalization of remediation equipment into housing value, endogenously determined
k	Cost of installing remediation equipment; \$1,200
θ	Denotes a "type"; unique combination of (s,a,r,w)
λ_θ	Probability of dying this period for type θ
δ_θ	Probability of moving next period for type θ
β	Discount factor; 0.971
p	Probability that the next house has remediation equipment installed, endogenously determined

Table 2.4: Capitalization and Prevalence of Remediation Equipment

	Value of a Life Year				
	\$100,000	\$200,000	\$300,000	\$400,000	\$500,000
Capitalization (Percent)	4.0	7.9	9.5	12.3	12.7
Percent of Housing Stock with Remediation	0.0	0.1	0.3	0.4	0.6

Table 2.5: Peak Age and Radon Levels for Remediation Decisions

	Value of a Life Year				
	\$100,000	\$200,000	\$300,000	\$400,000	\$500,000
Houses with Remediation Equipment Already Installed					
Non-Smokers - Age	—	—	50.1	49.7	49.5
Radon Level	—	—	18.3	14.2	11.8
Former Smokers - Age	49.7	48.8	48.2	47.1	48.9
Radon Level	14.7	8.3	6.2	5.2	4.5
Current Smokers - Age	48.8	45.9	43.8	42.1	43.7
Radon Level	7.5	4.8	4.0	3.6	3.2
Houses without Remediation Equipment Already Installed					
Non-Smokers - Age	—	—	—	—	—
Radon Level	—	—	—	—	—
Former Smokers - Age	—	45.8	47.4	45.6	45.3
Radon Level	—	18.0	12.4	9.6	8.0
Current Smokers - Age	41.8	45.6	46.3	45.3	45.2
Radon Level	14.6	9.0	6.6	5.3	4.8

Table 2.6: Capitalization and Prevalence of Remediation Equipment when some Households have Children

	Value of a Life Year \$300,000
Capitalization (Percent)	10.9
Percent of Housing Stock with Remediation	0.4

Table 2.7: Peak Age and Radon Levels for Remediation Decisions

	Number of Children in Household					
	No Children	1 Child	2 Children	3 Children	4 Children	5 Children
Houses with Remediation Equipment Already Installed						
Non-Smokers - Age	48.5	47.3	45.8	45.6	44.0	48.2
Radon Level	13.4	10.8	9.2	8.0	7.2	6.4
Former Smokers - Age	48.3	48.3	44.3	47.6	42.4	48.7
Radon Level	5.6	5.3	5.2	4.8	4.8	4.4
Current Smokers - Age	48.0	45.9	44.0	41.6	40.1	38.8
Radon Level	3.6	3.6	3.6	3.6	3.6	3.6
Houses without Remediation Equipment Already Installed						
Non-Smokers - Age	—	—	45.0	45.0	45.0	45.0
Radon Level	—	—	17.3	14.6	12.8	11.4
Former Smokers - Age	45.0	45.0	45.0	45.0	45.0	45.0
Radon Level	10.3	9.6	9.0	8.4	8.0	7.6
Current Smokers - Age	45.2	45.4	45.1	45.1	45.3	45.0
Radon Level	5.9	5.7	5.6	5.5	5.3	5.2

Figures

Figure 2.1: Decision Tree for Households

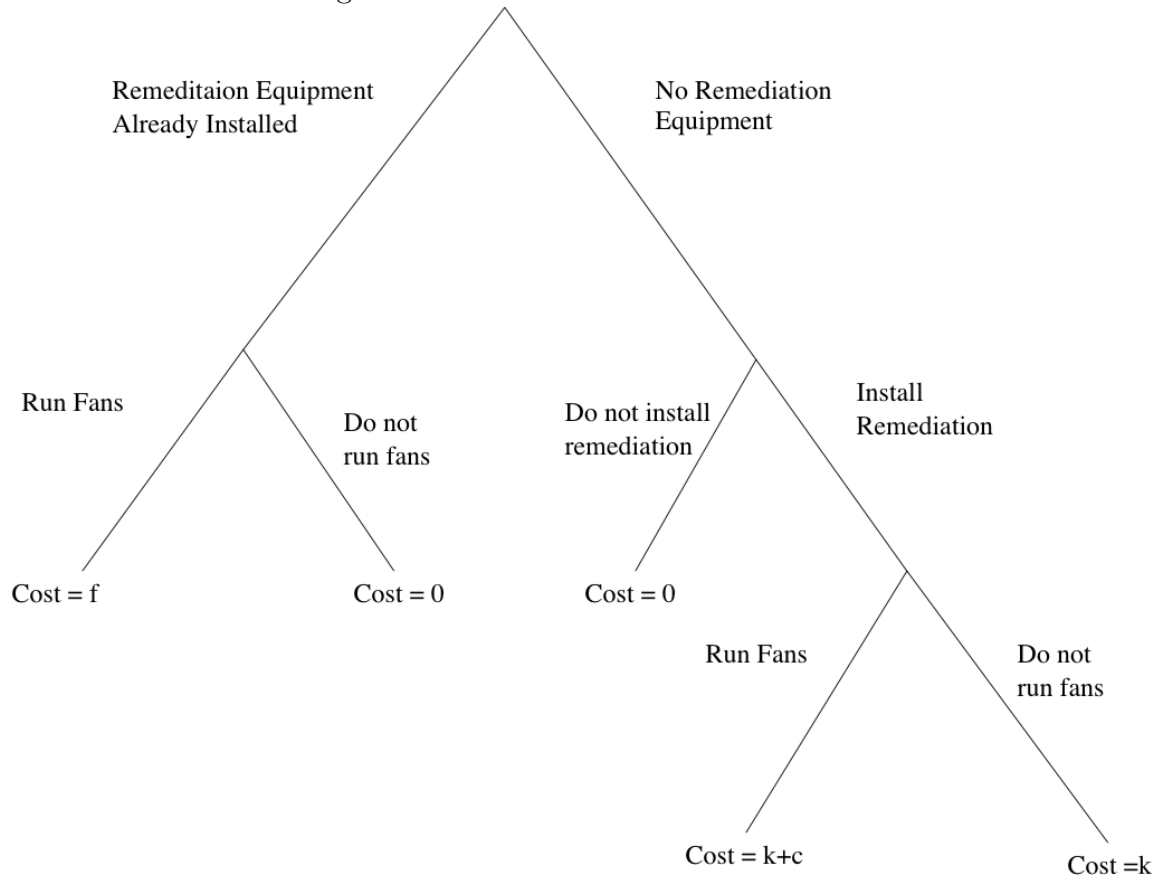


Figure 2.2: Probability of Remediation Equipment Being Installed in a House




Figure 2.3: Probability of Dying before reaching the next age




Figure 2.4: Density of Ages in the United States, 2006-2008




Figure 2.5: Density of Radon Concentrations in an Unremediated Housing Stock




Figure 2.6: Policy Function for a Typical Household's Remediation Decision with a Life-Year Valuation of \$300,000

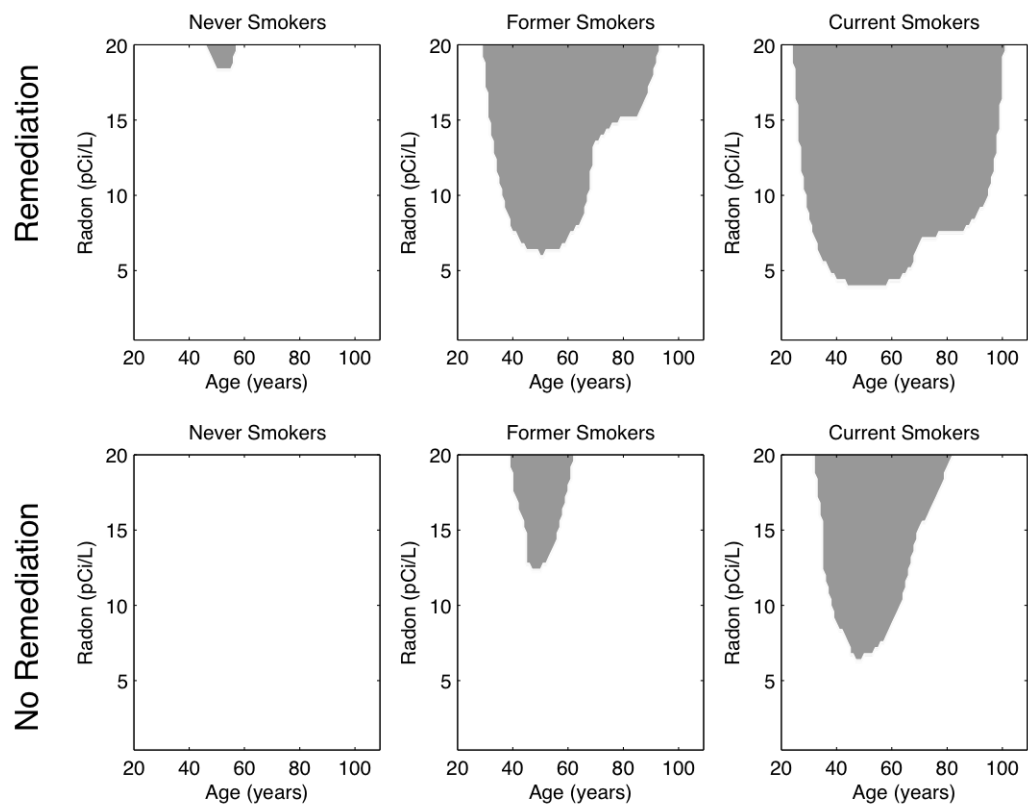


Figure 2.7: Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 with No Children

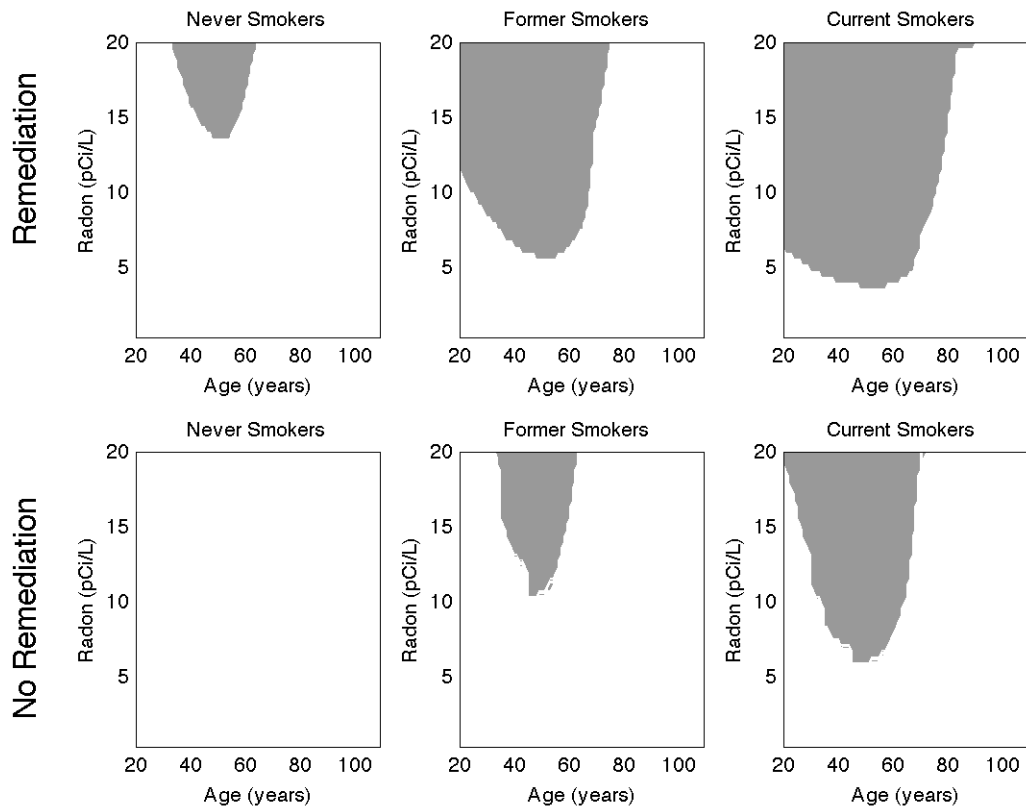


Figure 2.8: Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 with 1 Child

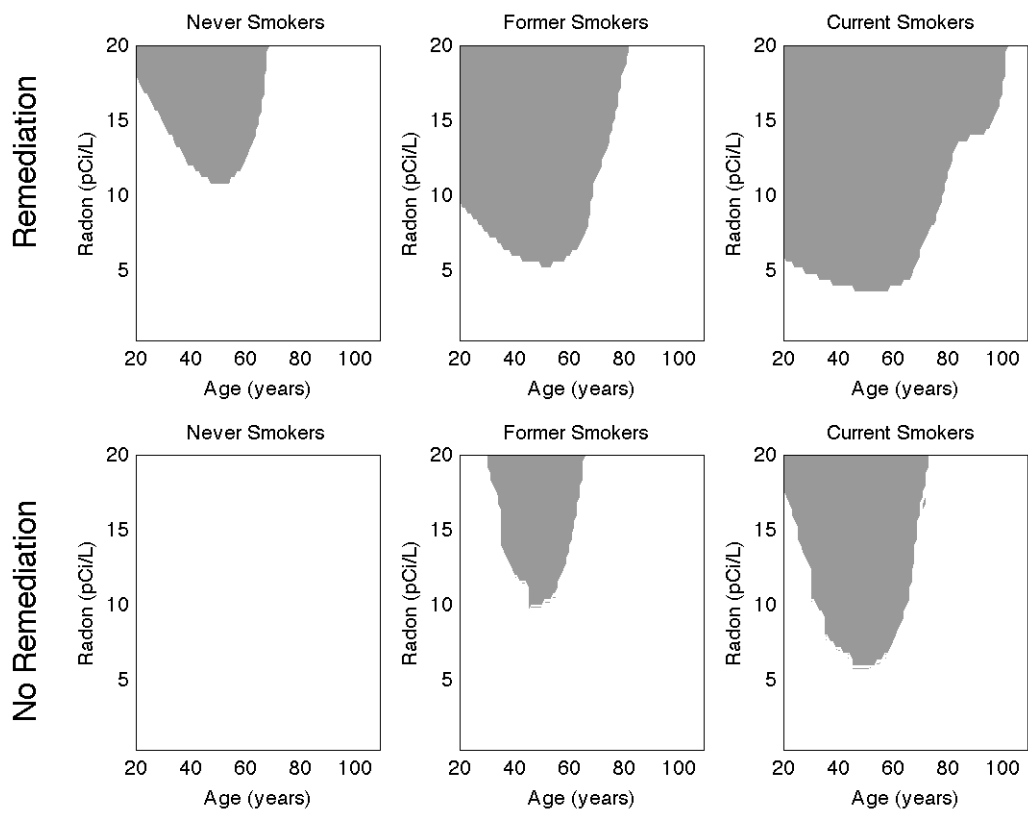


Figure 2.9: Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 with 2 Children

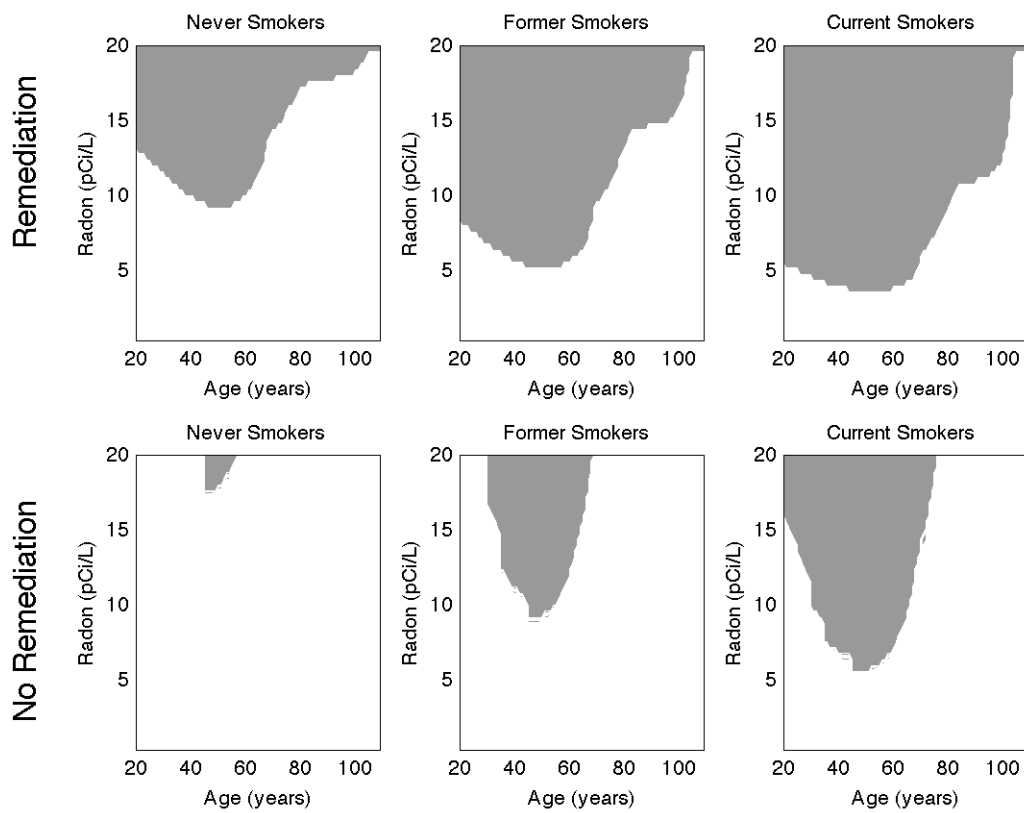


Figure 2.10: Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 with 3 Children

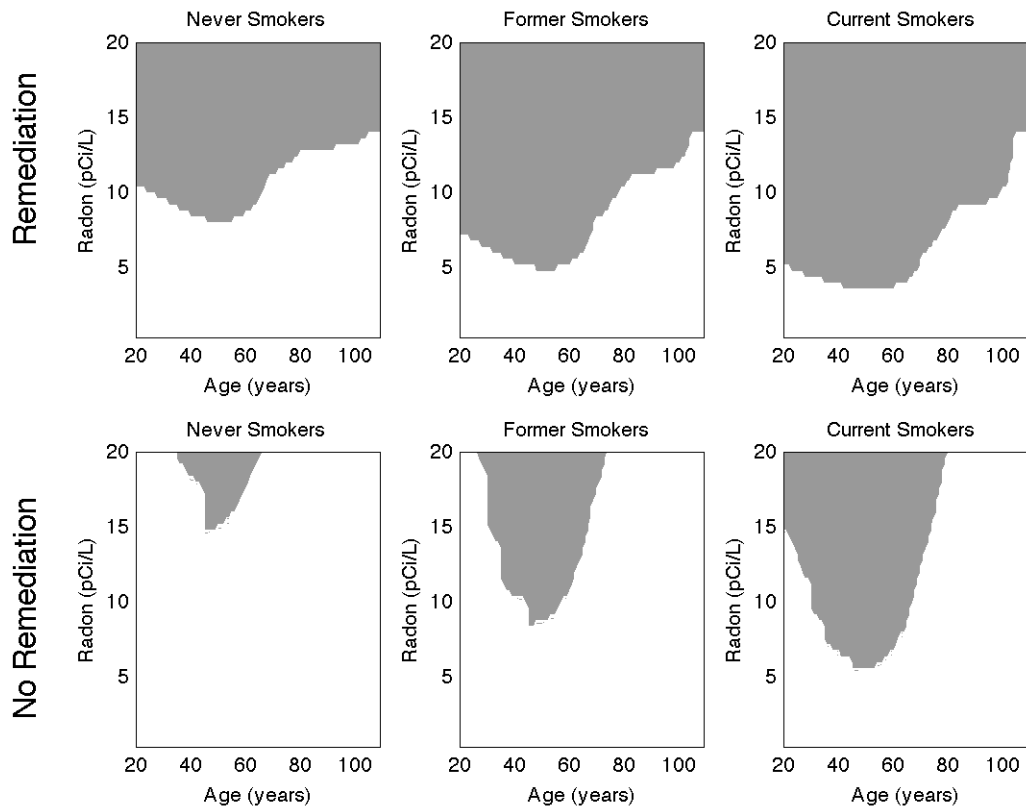


Figure 2.11: Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 with 4 Children

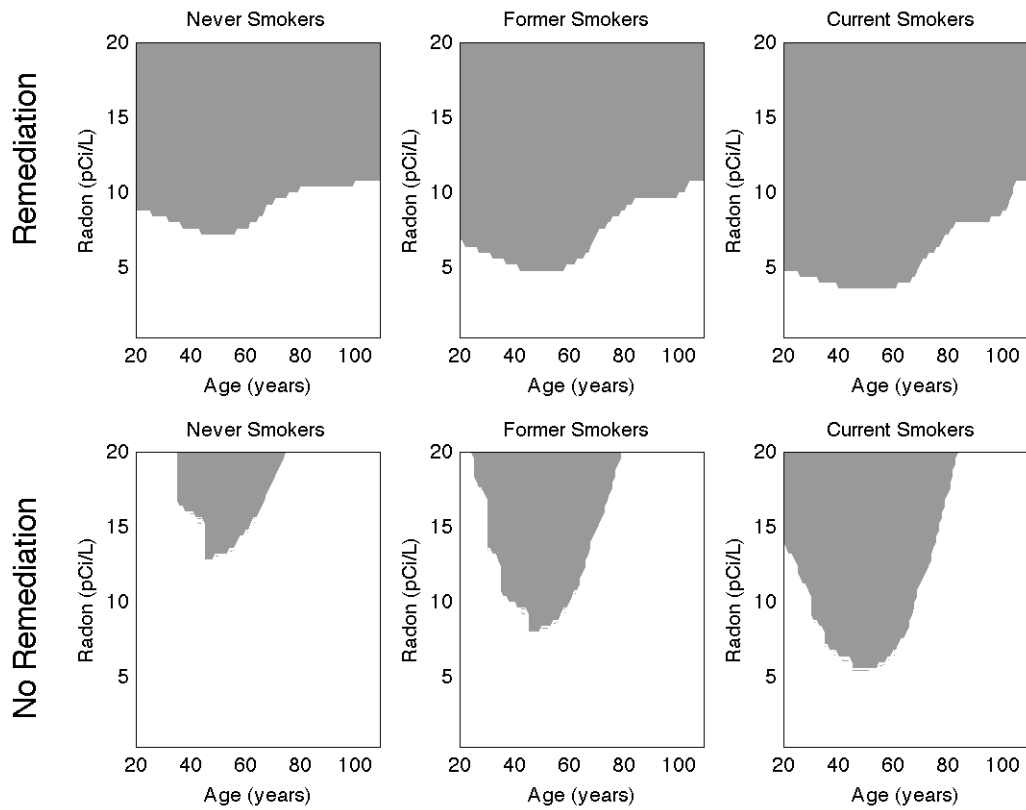


Figure 2.12: Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 with 5 Children

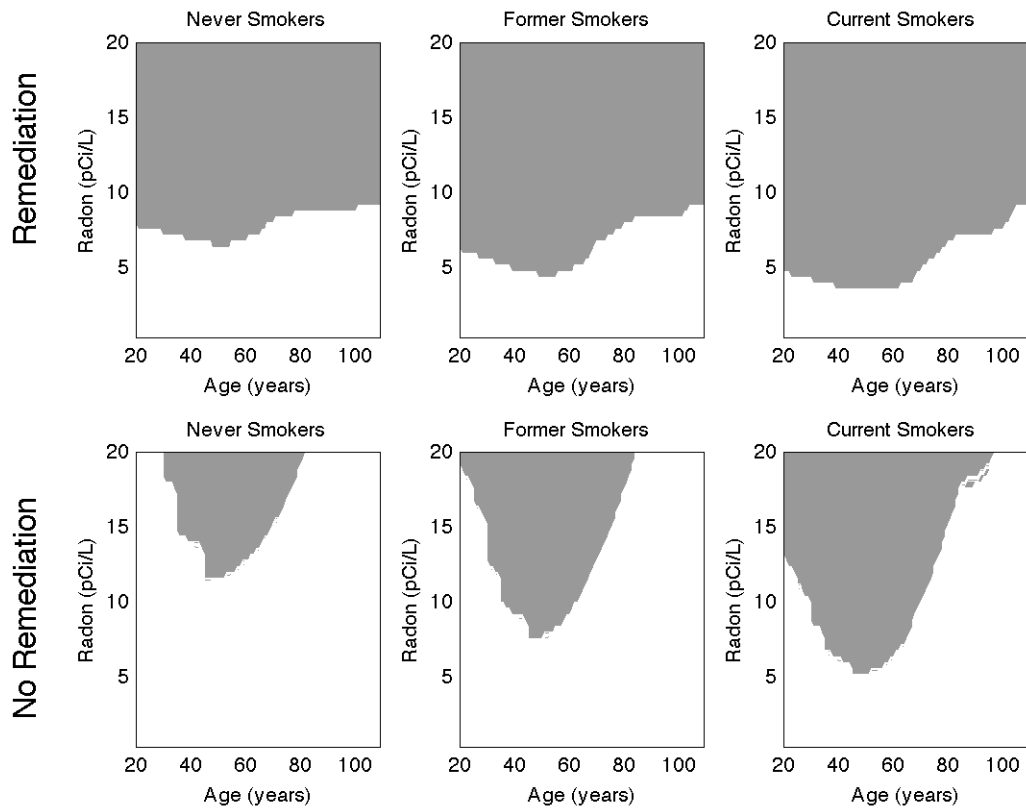


Figure 2.13: Typical Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 if There is No Capitalization of Remediation Equipment

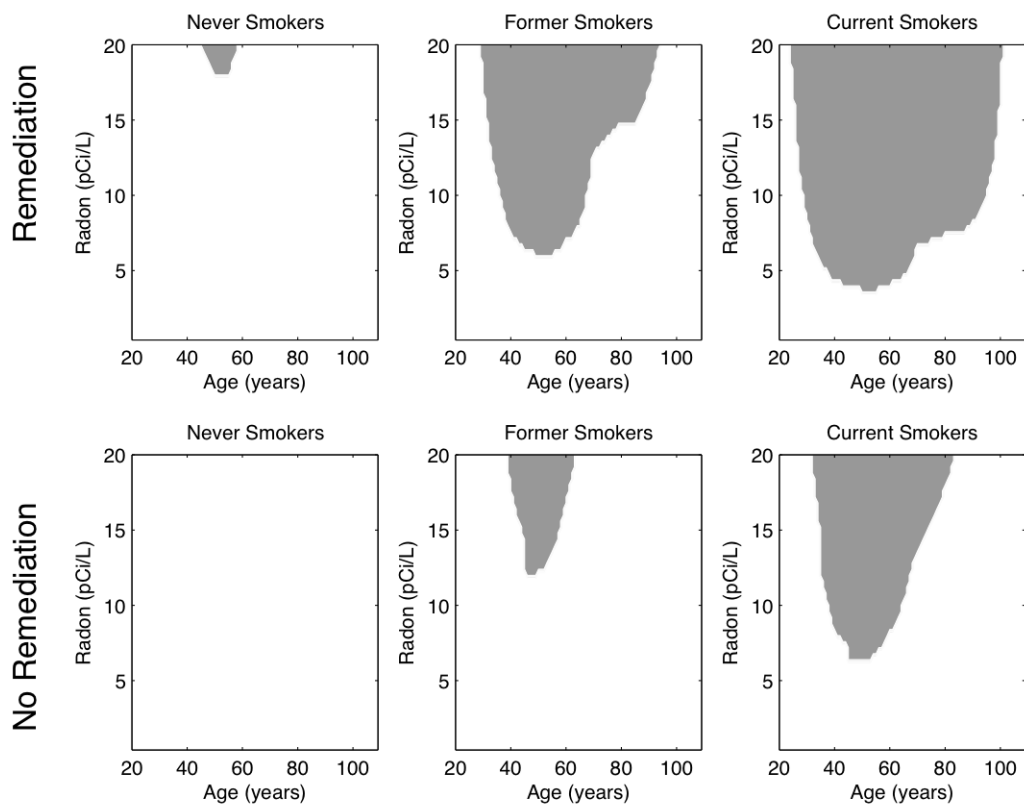
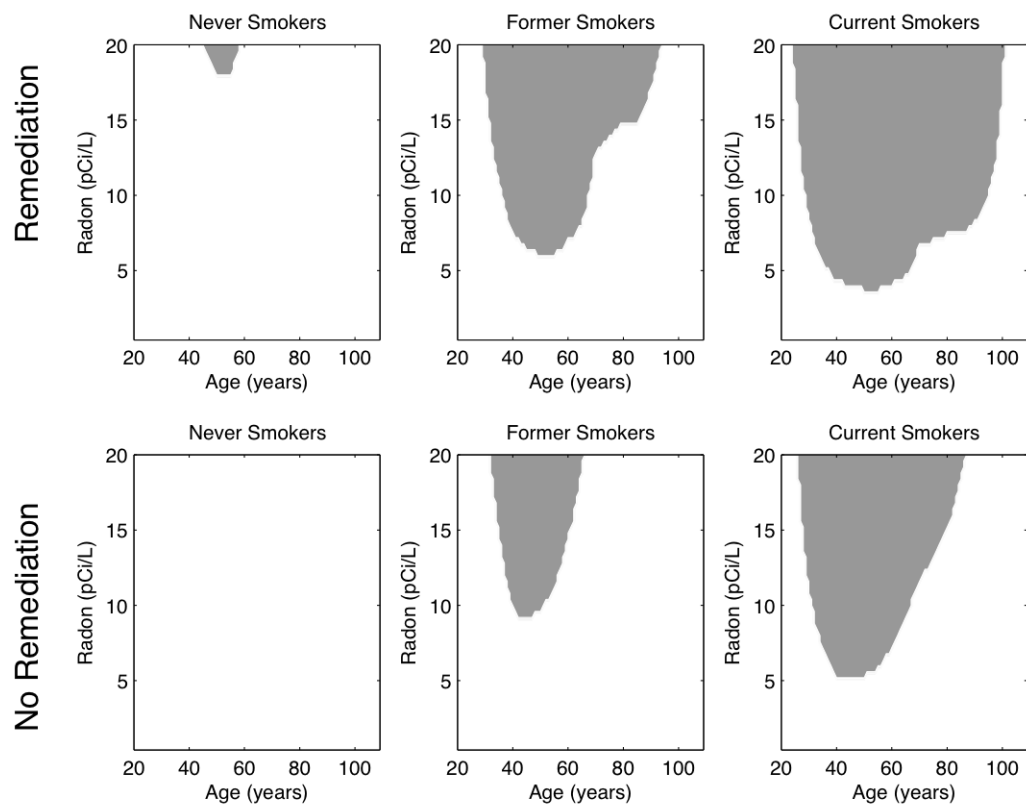


Figure 2.14: Policy Function for a Household's Remediation Decision with a Life-Year Valuation of \$300,000 if the Household Will Never Move



CHAPTER III

Monte Carlo Simulations of the Nested Fixed-Point Algorithm

3.1 Introduction

Structural estimation of dynamic programming problems has become a prominent tool in many economists' toolbox since the publication of John Rust's "The Optimal Replacement of GMC Bus Engines" [36], particularly in the context of dynamic discrete choice models¹. Papers that use structural estimation are generally characterized by a complete, explicit, usually dynamic, mathematical model of agents' behavior, then estimate the parameters of the model either through maximum likelihood or method of moments. However, these models often rely on assumptions about the distributions of unobservables and functional forms to make them tractable to estimate. Even with these assumptions, they usually result in highly nonlinear objective functions that present a challenge to estimate. There is a growing literature that examines the how numerical methods such as choice of optimization routines and starting points affect estimates. For example, Knittel and Metaxoglou [21] investigate how researchers' decisions about maximization algorithms and different starting points for each algorithm can lead to different answers. They find a very wide array

¹For a survey of the dynamic discrete choice literature see Rust [38], Pakes [31], Miller [26], and Mira [6].

of estimates can be obtained depending on the choices a researcher makes about the maximization algorithm and starting points. However, there are no papers to my knowledge that examine how distributional assumptions affect parameter estimates. This paper attempts to begin to fill this hole by examining a very simple dynamic structural model.

This paper uses the optimal stopping model from Rust [36] of GMC bus engine replacement as a starting point to evaluate how distributional assumptions affect the performance of the nested fixed point algorithm (NFXP). Rust develops a dynamic discrete choice model of bus engine replacement for the supervisor of the Madison Metro Bus Company, Harold Zurcher. In each period Mr. Zurcher observes the mileage that the bus has accumulated and has a discrete choice to make: replace the engine in the bus or use the current engine for another month. Rust then poses a functional form for Mr. Zurcher's utility function over bus engine replacements and assumes that Mr. Zurcher is a forward looking agent who dynamically maximizes this utility function. Rust then solves Mr. Zurcher's dynamic problem and finds parameter values that maximize the likelihood of the data. This involves solving the entire dynamic maximization problem for every set of parameter values. This is done through the nested fixed point algorithm. The nested fixed point algorithm is an inner loop that solves a dynamic programming problem for a given set of parameter values and an outer loop that uses a routine to maximize the likelihood function over the parameter space.

Since Rust developed this framework for solving dynamic discrete choice problems, there have been many algorithms proposed to solve similar problems. Hotz and Miller [18] showed that it is not necessary to solve the dynamic problem at every step like the nested fixed point algorithm requires. Instead, since there is a one to one mapping between conditional choice probabilities and normalized value functions, the conditional probabilities can be inverted into estimates of the value functions which in turn

allow the econometrician to update the conditional probabilities. Aguirregabiria and Mira [5] show that the nested fixed point algorithm and Hotz and Miller’s conditional choice probabilities estimator are two extreme cases of a general class of estimators.

There has generally been seen to be a tradeoff between efficiency (from the nested fixed point algorithm) and computation time (reduced by using the Hotz and Miller [18] routine). I have chosen to use the NFXP from Rust [36] as a starting point for this paper for two reasons. Firstly, since this was one of the first papers to employ a structural approach to a dynamic problem it has become one of the standards in evaluating new methods. This is partly because the algorithm is particularly easy to implement. For instance, Hotz et. al [19] perform Monte Carlo simulations to compare their conditional choice simulation estimator to the NFXP and examine the NFXP for sample sizes of 10,000 and more. Secondly, because the NFXP solves the dynamic programming problem at every step I expect that it would be more robust to specification error.

This paper contributes to the literature on structural estimation in two ways. First, it extends the range of sample sizes for which there is Monte Carlo evidence for the validity of the nested fixed point algorithm and similar algorithms. In this paper I simulate the NFXP for datasets with as few as 500 observations and as many as 11,800. (Previously, the literature had only examined sample sizes as small as 10,000 [19].) Given that many papers that use structural estimation of a dynamic programming problem rely on sample sizes less than 10,000 I feel this is the relevant range of observations. Second, I examine how distributional assumptions on the unobserved state variable effect the estimates of the structural parameters. While this is obviously a context specific effect it is still important to have a sense of how important these assumptions may be in parameter estimates.

The remainder of this paper is organized as follows. Section 2 describes in detail Rust’s model of bus engine replacement. Section 3 describes the data that is used in

Rust [36]. Section 4 discusses Rust's results and my replication of his results. Section 5 discusses asymptotic results for the estimation procedure and how I simulate the data. Section 6 discusses the simulation results and Section 7 concludes.

3.2 Rust's Model

Rust provides two versions of his model. The first model, which I name the simple model, imposes strict functional form assumptions on the transition probabilities and assumes that there are no unobserved state variables. The second model, which I call the relaxed model, relaxes the functional form assumption on the transition probabilities and introduces an unobserved (to the econometrician) state variable, ϵ_t , that Rust assumes has very specific properties.

3.2.1 The Simple Model

John Rust [36] models the behavior of the superintendent of the Madison Wisconsin Metropolitan Bus Company, Harold Zurcher, when deciding whether or not to replace the engine in one of the company's buses. The model takes the form of a regenerative optimal stopping problem. Each month, Mr. Zurcher must choose either to (i) leave the bus in service for another month, while doing "normal maintenance" and incur operating costs $c(x_t, \theta_1)$ or (ii) take the bus out of service for the month and completely replace the engine for a cost of \bar{P} and sell the old engine for scrap for a price of \underline{P} . (Let the replacement cost of the engine, $RC = \bar{P} - \underline{P}$.) Mr. Zurcher is assumed to be a rational actor who minimizes the expected discounted costs of maintaining the fleet of buses. It is assumed that a bus with a newly replaced engine is just as good as a new bus in terms of the future decisions of whether or not to replace the engine. Therefore, the optimal stopping problem takes the form:

$$V_\theta(x_t) = \sup_{\Pi} E \left\{ \sum_{j=t}^{\infty} \beta^{j-t} u(x_j, f_j, \theta_1) \middle| x_t \right\} \quad (3.1)$$

where

$$u(x_t, i_t, \theta_1) = \begin{cases} -c(x_t, \theta_1) & \text{if } i_t = 0 \\ -(RC + c(0, \theta_1)) & \text{if } i_t = 1 \end{cases} \quad (3.2)$$

where Π is an infinite sequence of decision rules $\Pi = f_t, f_{t+1}, \dots$ where each f_t specifies Mr. Zurcher replacement decision at time t as a function of the entire history of the process, $i_t = f(x_t, i_{t-1}, x_{t-1}, i_{t-2}, \dots)$ and the expectation is taken with respect to the controlled stochastic process, $\{x_t\}$ whose probability distribution is defined from Π and the transition probability $p(x_{t+1}|x_t, i_t, \theta_2)$. If an exponential distribution is assumed for $p(x_{t+1}|x_t, i_t, \theta_2)$ then the transition probabilities take the form,

$$p(x_{t+1}|x_t, i_t, \theta_2) = \begin{cases} \theta_2 \exp[-\theta_2(x_{t+1} - x_t)] & \text{if } i_t = 0 \text{ and } x_{t+1} \geq x_t \\ \theta_2 \exp[-\theta_2(x_{t+1})] & \text{if } i_t = 1 \text{ and } x_{t+1} \geq 0 \end{cases} \quad (3.3)$$

Therefore, if the current engine is kept ($i_t = 0$) the next period's mileage is given by a draw from the exponential distribution $1 - \exp[-\theta_2(x_{t+1} - x_t)]$, but if the engine is replaced ($i_t = 1$) then x_t regenerates to 0 and the next period's mileage is drawn from the exponential distribution $1 - \exp[-\theta_2(x_{t+1} - 0)]$.

I can write the Bellman's equation to this system as:

$$V_\theta(x_t) = \max_{i_t \in \{0,1\}} [u(x_t, i_t, \theta_1) + \beta EV_\theta(x_{t+1}, i_{t+1})] \quad (3.4)$$

This should imply a deterministic cut-off rule such that

$$i_t = f(x_t, \theta) = \begin{cases} 1 & \text{if } x_t > \gamma(\theta_1, \theta_2) \\ 0 & \text{if } x_t \leq \gamma(\theta_1, \theta_2) \end{cases} \quad (3.5)$$

for some function $\gamma(\cdot)$.

However since in the data, we do not observe this type of deterministic cut-off rule, we assume that there is an unobserved state variable, ϵ_t , that Mr. Zurcher observes but the econometrician does not observe.

3.2.2 The Relaxed Model

Rust now adds two parts to the model. We add the unobserved state variable, ϵ_t , which is assumed to be additively separable from the rest of the utility function. Also, he relaxes the assumption the the mileage is drawn from an exponential distribution with parameter θ_2 and allow the mileage process to have an arbitrary density and define the difference between this month's mileage and last month's mileage to have arbitrary density $g(\cdot)$. These new assumptions lead to the Bellman's equation:

$$V_\theta(x_t, \epsilon_t) = \max_{i_t \in \{0,1\}} [u(x_t, i_t, \theta_1) + \epsilon_t(i) + \beta EV_\theta(x_{t+1}, \epsilon_{t+1})] \quad (3.6)$$

which has the solution

$$f(x_t, \epsilon_t, \theta) = \arg \max_{i_t \in \{0,1\}} [u(x_t, i, \theta_1) + \epsilon_t(i) + \beta EV_\theta(x_{t+1}, \epsilon_{t+1})] \quad (3.7)$$

Because the unobserved state variable, ϵ_t , enters non-linearly into the unknown function, EV_θ Rust makes a "Conditional Independence" assumption to circumvent this problem. The conditional independence assumption can be stated as:

Assumption III.1. *Conditional Independence: The transition density of the controlled process $\{x_t, \epsilon_t\}$ factors as*

$$p(x_{t+1}, \epsilon_{t+1} | x_t, \epsilon_t, i, \theta_2, \theta_3) = q(\epsilon_{t+1} | x_{t+1}, \theta_2) p(x_{t+1} | x_t, i, \theta_3)$$

This assumption introduces two restrictions. First it requires that x_{t+1} is a sufficient statistic for ϵ_{t+1} , which means that any dependence between ϵ_t and ϵ_{t+1} is transmitted through x_{t+1} . Secondly, it requires that the probability density of x_{t+1} depends only

on x_t and not ϵ_t .²

If we further impose that $q(\epsilon|y, \theta_2)$ is given by a type 1 extreme value distribution then we can state the formula for the choice probability, $P(i|x, \theta)$ as follows:

$$P(i|x, \theta) = \frac{\exp[u(x, i, \theta_1) + \beta EV_\theta(x, i)]}{\sum_{j \in \{0,1\}} \exp[u(x, j, \theta_1) + \beta EV_\theta(x, j)]} \quad (3.8)$$

which is the familiar multinomial logit formula.

This allows us to estimate the structural parameters, $\theta \equiv \{RC, \theta_1, \theta_3\}$, of the controlled process $\{i_t, x_t\}$ through maximum likelihood as shown in Rust [37]. The likelihood function ℓ^f take the form

$$\ell^f(x_1, \dots, x_T, i_1, \dots, i_T | x_0, i_0, \theta) = \prod_{t=1}^T P(i_t | x_t, \theta) p(x_t | x_{t-1}, i_{t-1}, \theta_3) \quad (3.9)$$

This likelihood function can be estimated in three stages. The first stage is to estimate

$$\ell^1(x_1, \dots, x_T, i_1, \dots, i_T | x_0, i_0, \theta) = \prod_{t=1}^T p(x_t | x_{t-1}, i_{t-1}, \theta_3) \quad (3.10)$$

which is the transition probabilities between mileage bins. The second stage is to estimate

$$\ell^2(x_1, \dots, x_T, i_1, \dots, i_T | x_0, i_0, \theta) = \prod_{t=1}^T P(i_t | x_t, \theta) \quad (3.11)$$

which requires the computation of the fixed point to get estimates of θ_1 and RC , the variable cost parameter and the replacement cost of the engine respectively. Since estimating both ℓ^1 and ℓ^2 give consistent estimates of the parameters, I can then use these consistent estimates to estimate ℓ^f and get efficient estimates of all of the structural parameters.

²For proofs of these results see Rust [37].

3.3 Data

I have obtained the relevant parts of Rust's original data from the Madison Metropolitan Bus Company that contains monthly maintenance records for every bus in the Madison bus fleet from December, 1974 to May, 1985.³ The observations consist of odometer readings on each bus and an indicator specifying if the engine was replaced that month. In addition to the original data that I have obtained, Rust's original data also consisted of a maintenance diary that records all repairs that were made on a bus such as replacing brakes, oil changes, etc. Rust considers all events that are not a complete engine replacement "normal maintenance" and disregards that information for the sake of his exercise. I proceed likewise.

The data that I have from Rust contains the mileage for each bus at the end of every month, an indicator if the bus' engine was replaced in that month, and the model of the bus. There are eight types of buses in the Madison Metro fleet over the covered time period. See Tables 1a and 1b for summary statistics of the data.

In order to compute the value function in the dynamic programming problem, I will need to do a grid search. This means that I will need to discretize our mileage data into bins. I discretize the continuous mileage variable into 90 bins of 5,000 miles each.⁴ This gives bins up to 450,000 miles to allow for that value function to be estimated for mileages above what I observe in the data (the maximum mileage I observe is 387,300) and allows for the possibility that it may be optimal to replace the engine at a mileage level large than I observe.

Now that I have discretized the mileage process, I can rewrite the transition density as the difference between last month's bin and this month's bin, giving density

³The data used in the original paper is available at <http://gemini.econ.umd.edu/jrust/nfxp.html>

⁴Rust [36] does some sensitivity analysis by increasing the number of mileage bins and finds results similar to those with 90 bins.

$$p(x_{t+1}|x_t, i_t, \theta_3) = \begin{cases} g(x_{t+1} - x_t, \theta_3) & \text{if } i_t = 0 \\ g(x_{t+1} - 0, \theta_3) & \text{if } i_t = 1 \end{cases} \quad (3.12)$$

In the data I only have buses where $(x_{t+1} - x_t) \in \{0, 1, 2\}$, thus I define θ_{30} as the probability that you stay in the same mileage bin as you were last month, θ_{31} , as the probability that you move to the next mileage bin, and θ_{32} as the probability that you move up two mileage bins. This reduces to a multinomial distribution with parameters θ_{30} , θ_{31} (and $\theta_{32} = 1 - \theta_{30} - \theta_{31}$).

In order to use the nested fixed point algorithm I need to assume that there is no heterogeneity in our data between the different types of buses. Rust tests the hypothesis that the mileage process is different for various groupings of bus types and cannot reject the null that bus types 1-4 have the same mileage process, while you can reject the null that bus types 1-4 have a different mileage process from types 5-8.⁵ Therefore, I proceed with the exercise only using bus types 1-4.

Next, I need to specify a functional form for the cost function. Rust did not find one particular functional form to fit the data statistically better than any other that he tried and therefore used a linear cost function⁶ with one unknown parameter defined as $c(x, \theta_1) = .001\theta_{11}x$.

Following Rust, I choose to fix β instead of estimating it since it is highly collinear with the fixed cost of replacement, RC . This collinearity becomes obvious by examining the value function since a lower discount factor will weight the present higher, which has the same effect as raising RC . Following Rust, for the rest of the paper I fix $\beta = 0.9999$.

⁵See Rust [36] for a detailed analysis and results.

⁶A square root cost function, $c(x, \theta_1) = \theta_{11}\sqrt{x}$ was also used, but results not reported.

3.4 Rust's Results and Replication

Using the exact data that Rust [36] uses, I proceed with the replication using Rust's methods, detailed in Rust [39]⁷. Rust [36] uses the nested fixed point algorithm to solve Mr. Zurcher's dynamic discrete control problem. This algorithm consists of two loops. The inside loop uses a combination of two methods to compute the fixed point. The first method used is the commonly used value function contraction iterations. Value function iteration defines a fixed point as $EV_\theta = T_\theta(EV_\theta)$ where $T_\theta(W)$ is defined as

$$T_\theta(W)(x) = \int_0^\infty \log[\exp\{-c(x+y), \theta\} + \beta W(x+y)] + \exp\{-RC - c(0, \theta) + \beta W(0)\}]g(dy|\theta)$$

This method begins with an arbitrary guess for EV_θ (usually equal to zero) and evaluates the value function given parameters θ and iterates the process k times. The k^{th} iteration can be written as $EV_k = T_\theta^k(EV_0)$ and as $k \rightarrow \infty$ it can be shown that $EV_k \rightarrow EV_\theta$.

Value function iteration converges at a linear rate to EV_θ . An alternative method, known as Newton-Kantorovich iteration uses an alternate method of iteration that converges at a quadratic rate when in the neighborhood of EV_θ . Thus, I use Werner's method [40] which uses value function iteration for the first few contraction steps and then switches to the Newton-Kantorovich method. Werner [40] showed that this produces a faster rate of convergence than either method alone.

Once the value function has converged, I evaluate the log-likelihood function using the assumed parameters $\theta_{11}, \theta_{30}, \theta_{31}, RC$. To get a new guess for the structural parameter, I use the outer hill climbing algorithm to find the parameters that maximize

⁷The replication was done using similar GAUSS code to that available through John Rust's website (<http://gemini.econ.umd.edu/jrust/nfxp.html>)

the likelihood function. Following Rust, I use the BHHH algorithm, which is similar to the Gauss-Newton and Newton-Raphson algorithms.

The results are presented in Table 2. As can be seen, my estimates of the transition probabilities are nearly identical, though the estimates of the cost function parameter, θ_{11} , and the replacement cost, RC , differ somewhat. It seems likely that I have found a slightly different local maximum than the original paper. Ideally, I would start the NFXP routine at many starting values and compare the value of the likelihood function at all maximum that the algorithm converges to in order to choose the global maximum. However, since for this paper it is only important that the routine always find "the same" maximum I will always initialize the algorithm to the same starting values so that it will likely head to the same local maximum.

3.5 Asymptotic Results and Simulation Procedure

Rust [37] shows that parameters estimated using the nested fixed-point maximum likelihood(NFXP) algorithm, $\hat{\theta}$, converges to the true value, θ^* with probability 1 as either N , the number of observations, or T , the number of periods, approaches infinity. He also shows that $\sqrt{N}(\hat{\theta} - \theta^*) \xrightarrow{d} N(0, -H(\theta^*)^{-1})$ where $-H(\theta^*)^{-1}$ is the negative inverse of the Hessian for θ^* . The main assumptions needed to make this result hold is that the model is correctly specified, the Conditional Independence assumption:

$$p(x_{t+1}, \epsilon_{t+1}|x_t, \epsilon_t, i, \theta_2, \theta_3) = q(\epsilon_{t+1}|x_{t+1}, \theta_2)p(x_{t+1}|x_t, i, \theta_3),$$

and some regularity conditions⁸.

Having validated the results reported in Rust [36], I examine the finite sample properties of the maximum likelihood estimator using the nested fixed point algo-

⁸For a complete proof and set of assumptions see Theorem 4.3 of Rust [37].

rithm. To do this, I simulate 1,000 datasets that are generated assuming that the model posed by Rust is correct and use the values that I estimated in the replication section as the "true" values of the model.

In order to simulate bus replacement data, there are two levels of randomness that need to be incorporated. First, the mileage bin that a bus falls into in a given month is a random variable that I model with a multinomial distribution, the parameters of the distribution, $\{\theta_{30}, \theta_{31}\}$ are random variables. Secondly, the model assumes an unobserved state variable, ϵ , that enters additively into the utility function and is independent across time and choices that is drawn from a Type I extreme value distribution.

In order to perform the actual simulation I need to proceed in a chronological order for each bus. Each bus is assumed to have an odometer reading of zero at the beginning of the simulation. In the first period the each bus receives a draw from the multinomial distribution for which mileage bin it will end that period in.

Once I know which mileage bin each bus ended the period in, I then evaluate the solution to the Bellman's equation, given parameters, θ .

$$f(x_t, \epsilon_t, \theta) = \arg \max_{i_t \in \{0,1\}} [u(x_t, i, \theta_1) + \epsilon_t(i) + \beta EV_\theta(x_{t+1}, \epsilon_{t+1})] \quad (3.13)$$

where

$$u(x_t, i, \theta_1) = \begin{cases} -[0.001\theta_{11}x_t + \epsilon_t(0)] & \text{if } i_t = 0 \\ -[RC + 0.001\theta_{11}x_0 + \epsilon_t(1)] & \text{if } i_t = 1 \end{cases} \quad (3.14)$$

and $\epsilon_t(\cdot)$ is drawn from a Type I extreme value distribution with mean 0 and variance $\frac{\pi^2}{6}$. If the value of replacement is larger than the value of not replacing the engine, then $i_t = 1$ and the bus starts over at mileage bin zero the next period. Once I have done this simulation for one month, I repeat the process for each of the 110 buses in each dataset for 118 months⁹. This leaves each dataset with approximately 13,000

⁹I chose 118 months since this is the maximum duration of data that is used for estimation in

bus-month observations.

3.6 Simulation Results

I first report the results from the simulations using datasets with relatively small sample sizes and then will discuss the results from simulations where the data generating process (DGP) is not the assumed DGP in the model. I find two largely consistent themes across all of the simulations. First, the estimator is biased in all samples examined, for all 4 parameters, though the bias decreases as I get closer to the assumptions of the model (unobservables become closer in distribution to the assumed EV1 unobservables). Second, the asymptotic variance is substantially smaller than the observed variance of the distribution of parameters.

3.6.1 "Small" Sample Results

Using the procedure described above, I produced datasets with 1,000, 2,000, 4,000, and 8,000 bus-month observations. Knittel and Metaxoglou [21] have shown that the choice of starting values can create very different results in highly non-linear environments. Therefore, in all of the simulations I use the same starting value, which is within 0.1 of the true value.

As can be seen in Figures 3.1-3.2 two of the four parameters that I estimate appear to have distributions close to their theoretical asymptotic distributions (shown in the red dotted line). However, in all of these simulations we get a biased estimated and in general a slightly larger variance. Examining the two multinomial transition probability parameters, θ_{30} and θ_{31} , we see that the both have a relatively large bias and a substantially larger variance than they should. This pattern also holds when the sample size increases to 13,000.

The mean and standard deviation of these distributions are shown in Table 3.4.

Rust [36].

We can see that the mean squared error decreases proportionally to the increase in sample size for the smallest sample sizes, with only a marginal decrease in the mean squared error between the 8,000 observation sample and 13,000 observation sample.

3.6.2 Distributional Results

Using the procedure described above, I produced simulated datasets, each containing 13,000 bus-month observations and used the nested fixed-point algorithm to estimate $\{RC, \theta_{11}, \theta_{30}, \theta_{31}\}$. Knittel and Metaxoglou [21] have shown that the choice of starting values can create very different results in highly non-linear environments. Therefore, in all of the simulations I use the same starting value, which is within 0.1 of the true value.

In order to explore the sensitivity of the nested fixed point algorithm to the assumption that the errors are distributed Type I extreme value, I will generate datasets with errors from three different distributions: Type I extreme value, Gaussian, and a Student's T with 3 degrees of freedom. I choose these distributions since they all have unbounded support. The Gaussian distribution is useful since it is similar in shape to the Type I extreme value. Meanwhile the Student's T_3 simulation will allow me to examine the behavior of the NFXP when I have many "large" errors.

3.6.2.1 Extreme Value Unobservables

Of the 3,000 datasets that were created, 1,148 datasets produced results that converged using our criterion that when the gradient times the direction is less than 1×10^{-8} . The remaining datasets produced parameter estimates where the gradient was ∞ or $-\infty$. While there are many different ways that the algorithm may be modified to get these datasets to converge, I will simply throw out these datasets from the analysis to focus on this particular procedure.

The means and standard deviations of these parameters can be seen in Table 3.5.

Table 3.5 suggests that the nested fixed-point algorithm does not provide unbiased estimates of the true parameters in our sample. Two possible explanations for this naturally present themselves. First, by only using 100 buses each for 118 months, I may not have gotten close enough to ∞ to have a consistent estimate of the parameters. Secondly, my results may be biased because of the datasets that did not converge. It seems likely that these datasets may be different in some systematic way.

I reject the null hypothesis that the mean of the parameters from the simulations is equal to true mean.¹⁰ Figures 3.5-3.8 display the full distribution of parameters from the simulations. All of the parameters appear to be distributed approximately Gaussian as the theory suggests, however formal tests, such as the Shapiro Wilk tests reject the null hypothesis that the data are Gaussian.

3.6.2.2 Gaussian Unobservables

Qualitatively, the simulation results from datasets that have Gaussian disturbances are similar to those with Extreme Value disturbances. There were 1,163 datasets that converged using the same convergence criterion, while the rest produced parameter estimates where the gradient was ∞ or $-\infty$.

The second panel of Table 3.5 shows descriptive statistics of these parameter. Note that the mean squared error from these results is approximately 1.5 times larger than that from the simulations with extreme value disturbances. The mean squared error for the parameters of the multinomial distribution does not change much in relative terms across any of the simulations, suggesting that the likelihood function is relatively well behaved in these dimensions.

The full distribution of parameters is displayed in Figures 3.5-3.8 with an overlaid Gaussian distribution. Again, using formal tests of normality I reject the null that

¹⁰I have not taken into account that the observations are estimated and therefore the standard errors should be larger. However, in my opinion, it is unlikely that this would change the results substantively.

the estimates are distributed Gaussian.

3.6.2.3 Student's T_3 Unobservables

The simulation results from datasets that have Student's T_3 disturbances are quite different from the other two simulations. The center of the parameter estimates is biased substantially downward, with a mean square error of 10,000 times that from the extreme value disturbances for one parameter and 200 times larger for another parameter. These results are shown in the bottom panel of Table 3.5 with the full distribution of parameters displayed in Figures 3.5-3.8. Since the Student's T_3 distribution has a higher probability of getting extreme values for the disturbance term, particularly extreme negative values, it makes sense that I end up with results that are biased substantially downward.

3.7 Conclusion

Empirical applications of highly nonlinear estimators has grown extensively recently. Naturally, these studies rely on asymptotic properties derived in the literature. However, there has been little examination of how these estimators perform in finite samples.

This paper adds to the growing literature that explores the numerical and finite sample behavior of nonlinear structural estimators. This study asks the question of how much data is "enough" to use asymptotic results for inference about the estimated structural parameters. By simulating datasets produced knowing that the model is correctly specified, I examine the marginal distributions of parameter estimates and find them to be non-Gaussian.

The results suggest that the NFXP performs relatively poorly for samples sizes smaller than 13,000. However, there appears to be substantial gains in terms of mean squared error up to at least 8,000 observations, at which point we see the mean

square error decreasing less rapidly than moving between smaller sample sizes. I can reject the null hypothesis that the empirical distribution is Gaussian for all of the parameter \times sample sizes in this paper, with the distributions of the transition probabilities performing the most poorly.

One reason that the estimates may appear non-Gaussian is that the simulated datasets do not have enough observations (results are proved as either $T \rightarrow \infty$ or $N \rightarrow \infty$). However, each of these datasets contain at least 13,000 observations, which is more than many structural models have at their disposal.¹¹ Therefore, we should be cautious about inference that we draw from finite samples smaller than our simulation sample size.

This paper has also explored to what extent one particular estimator, the nested fixed point algorithm, depend upon distributional assumptions. Though the NFXP is rarely used due to the computational burden of computing a fixed point at every iteration, it is part of a larger class of nested-pseudo likelihood estimators that depend on distributional assumptions. We have found that when the distributional assumptions are met, the estimator performs similarly to the theory. However, as we move away from the assumed distribution, we get worse parameter estimates with a mean square error of up to 10,000 times larger than the mean square error when the assumptions are met.

¹¹Rust [36] estimates his model on 8,156 observations and Berry, Levinsohn, and Pakes [7] use 2,271 model/year observations in their seminal paper. Berry et. al do not use the nested fixed-point algorithm though their objective function is also highly nonlinear.

Tables

Table 3.1: Summary of Replacement Data (Buses where at least 1 replacement occurred)

Bus Group	Mileage at Replacement				Elapsed Time (Months)				
	Max	Min	Mean	Standard Deviation	Max	Min	Mean	Standard Deviation	Number of Observations
1	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0
3	273,400	124,800	199,733	37,459	74	38	59.1	10.9	27
4	387,300	121,300	257,336	65,477	116	28	73.7	23.3	33
5	322,500	118,000	245,291	60,258	127	31	85.4	29.7	11
6	237,200	82,400	150,786	61,007	127	49	74.7	35.2	7
7	331,800	121,000	208,963	48,981	104	41	68.3	16.9	27
8	297,500	132,000	186,700	43,956	104	36	58.4	22.2	19
Full Sample	387,400	83,400	216,354	60,475	127	28	68.1	22.4	124

Source: Rust [36].

Table 3.2: Censored Data (Subsample of buses for which no replacements occurred)

Bus Group	Mileage at Replacement					Elapsed Time (Months)					Number of Observations
	Max	Min	Mean	Standard Deviation	Max	Min	Mean	Standard Deviation	Max		
1	120,151	65,643	100,117	12,929	25	25	25	0	25	0	15
2	161,748	142,009	151,183	8,530	49	49	49	0	49	0	4
3	280,802	199,626	250,766	21,325	75	75	75	0	75	0	21
4	352,450	310,910	337,222	17,802	118	117	117.8	0.45	117.8	0.45	5
5	326,843	326,843	326,843	0	130	130	130	0	130	0	1
6	299,040	232,395	265,264	33,332	130	128	129.3	1.15	129.3	1.15	3
7	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0
Full Sample	352,450	65,643	207,782	85,208	130	25	66.4	34.6	66.4	34.6	49

Source: Rust [36].

Table 3.3: Structural Estimates from Rust 1987 Fixed Point Dimension = 90, $\beta = 0.9999$

Parameter	Rust Estimate	Replication Estimate
RC	9.7558	8.5075
θ_{11}	2.6275	0.7571
θ_{30}	0.3489	0.3491
θ_{31}	0.6394	0.6396
Log-Likelihood	-6055.25	-6053.55
Observations	8,156	8,156

Source: Rust [36] and author's calculations.

Table 3.4: Summary Statistics of Parameter Estimates Model Assumptions Satisfied

Observations	Parameter	True		Standard Deviation	Mean Squared Error ($\times 10^3$)
		Mean	Mean		
1,000	RC	8.50	8.940	4.772	22962.9
	θ_1	0.76	0.862	0.561	326.2
	θ_{30}	0.35	0.334	0.023	0.7
	θ_{31}	0.64	0.614	0.038	2.1
2,000	RC	8.50	8.999	1.706	3149.7
	θ_1	0.76	0.880	0.360	144.8
	θ_{30}	0.35	0.335	0.020	0.6
	θ_{31}	0.64	0.617	0.033	1.6
4,000	RC	8.50	8.906	1.176	1541.9
	θ_1	0.76	0.846	0.263	76.9
	θ_{30}	0.35	0.336	0.018	0.5
	θ_{31}	0.64	0.617	0.032	1.5
8,000	RC	8.50	8.962	0.906	1025.7
	θ_1	0.76	0.847	0.202	49.1
	θ_{30}	0.35	0.336	0.017	0.5
	θ_{31}	0.64	0.618	0.031	1.4

Table 3.5: Summary Statistics of Parameter Estimates 13,000 Observations

		True		Standard	Mean Squared
	Parameter	Mean	Mean	Deviation	Error ($\times 10^3$)
EV1	RC	8.50	8.903	0.725	682.0
	θ_1	0.76	0.835	0.159	31.4
	θ_{30}	0.35	0.336	0.016	0.4
	θ_{31}	0.64	0.617	0.029	1.3
Gaussian	RC	8.50	9.008	0.812	908.3
	θ_1	0.76	0.882	0.170	44.4
	θ_{30}	0.35	0.335	0.016	0.5
	θ_{31}	0.64	0.616	0.030	1.5
T_3	RC	8.50	5.875	0.286	7016.1
	θ_1	0.76	0.510	0.095	70.0
	θ_{30}	0.35	0.341	0.021	0.5
	θ_{31}	0.64	0.627	0.038	1.6

Figures

Figure 3.1: Simulations of "Small" Data Sets: Replacement Cost Parameter

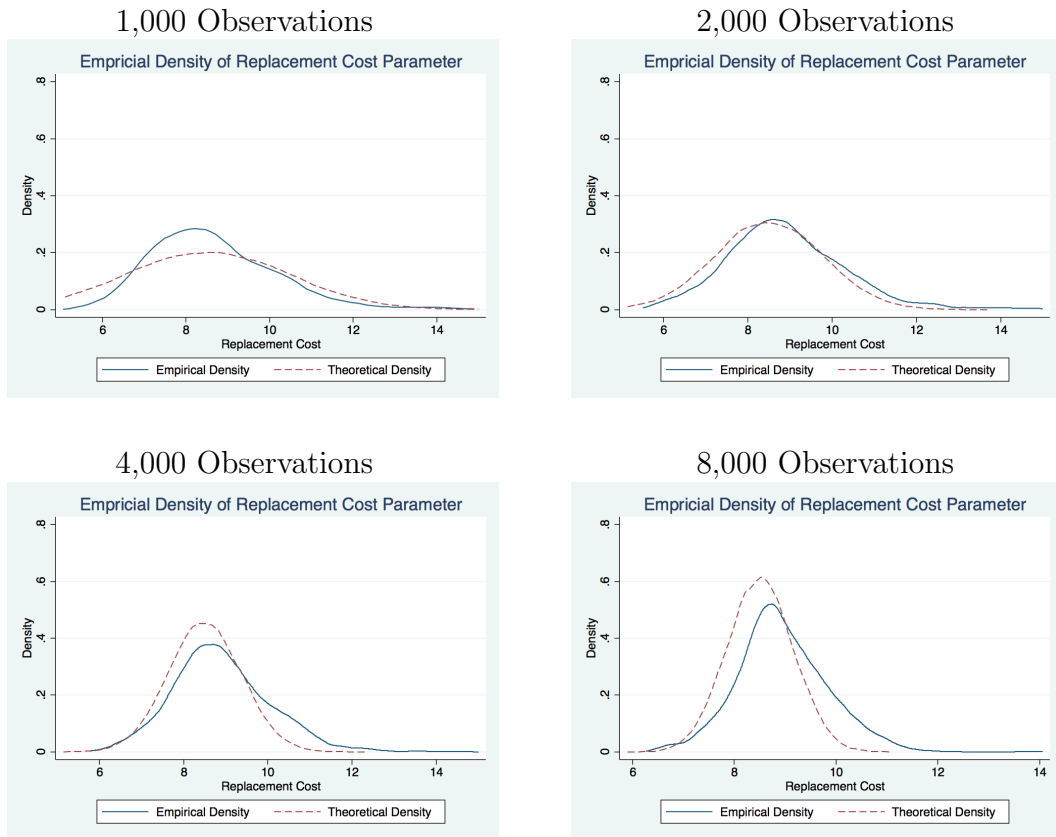


Figure 3.2: Simulations of "Small" Data Sets: Cost Function Parameter

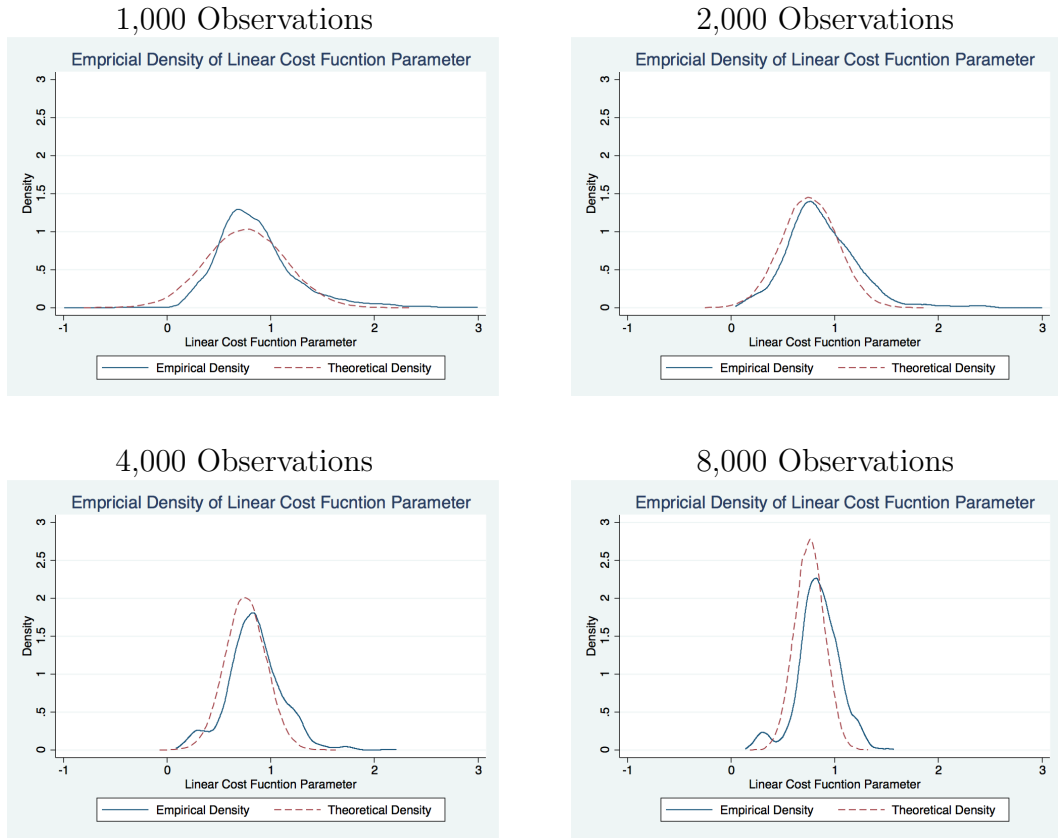


Figure 3.3: Simulations of "Small" Data Sets: $P(x_{t+1} - x_t = 0)$

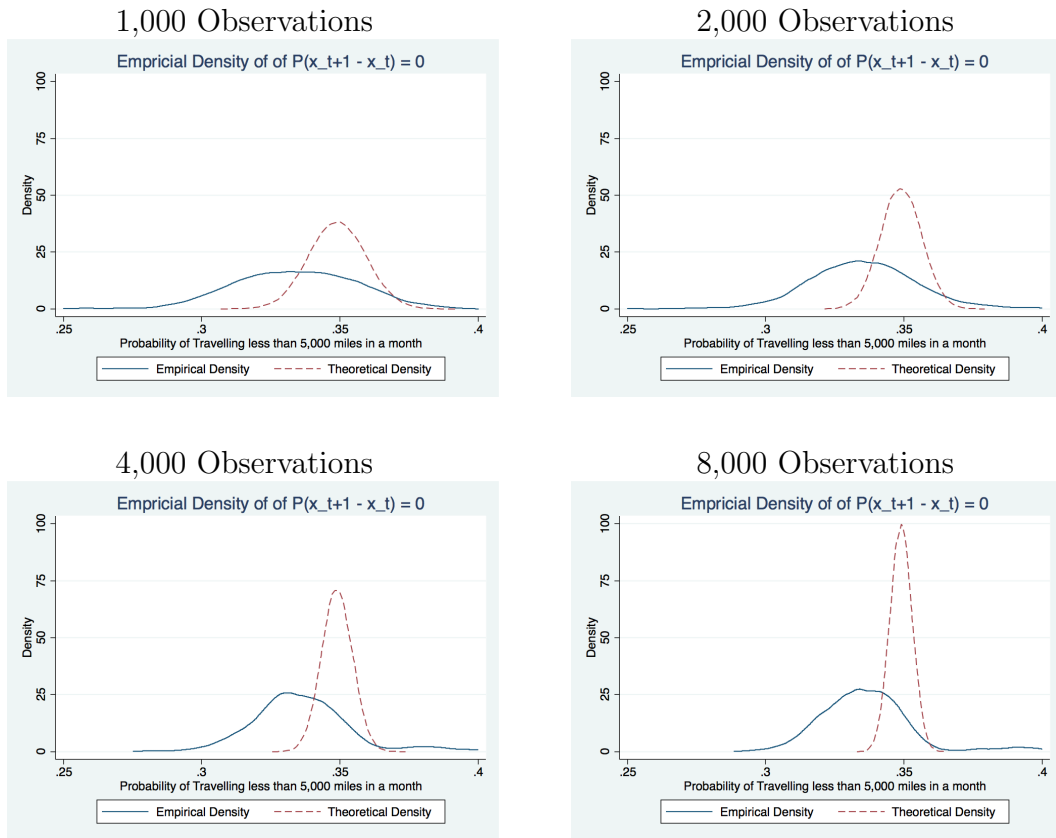


Figure 3.4: Simulations from "Small" Data Sets: $P(x_{t+1} - x_t = 1)$

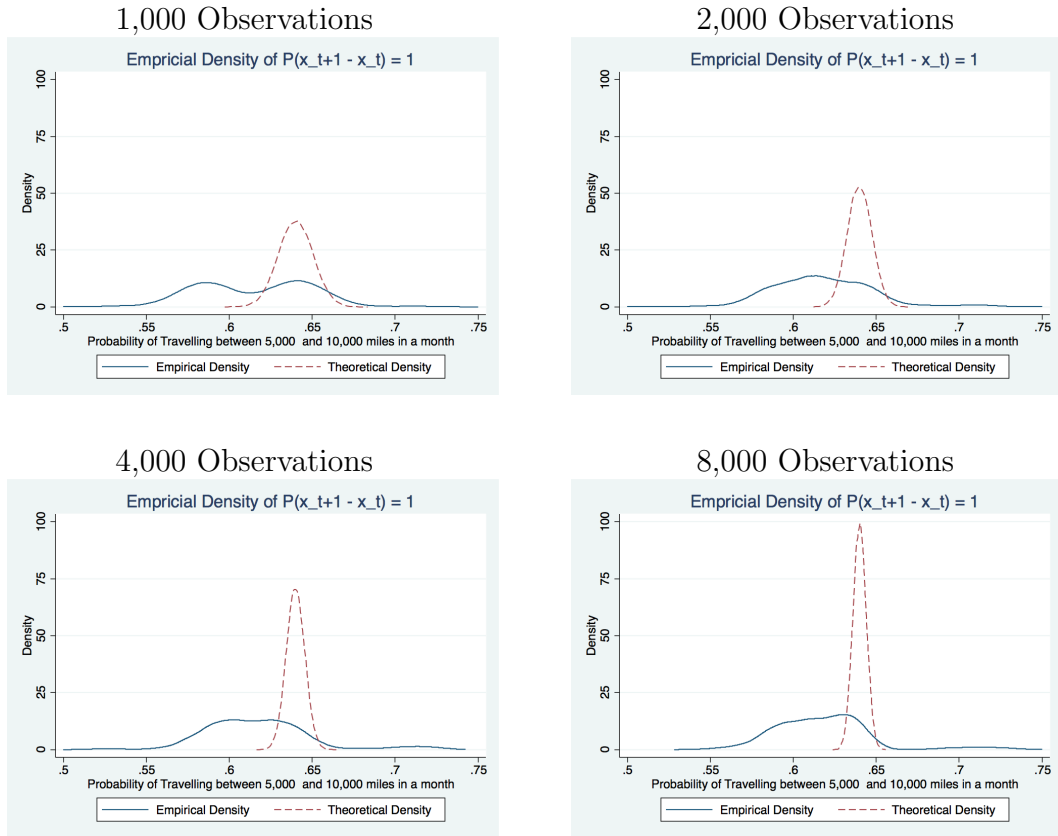


Figure 3.5: Simulations from Different Unobservable Distributions: Replacement Cost Parameter

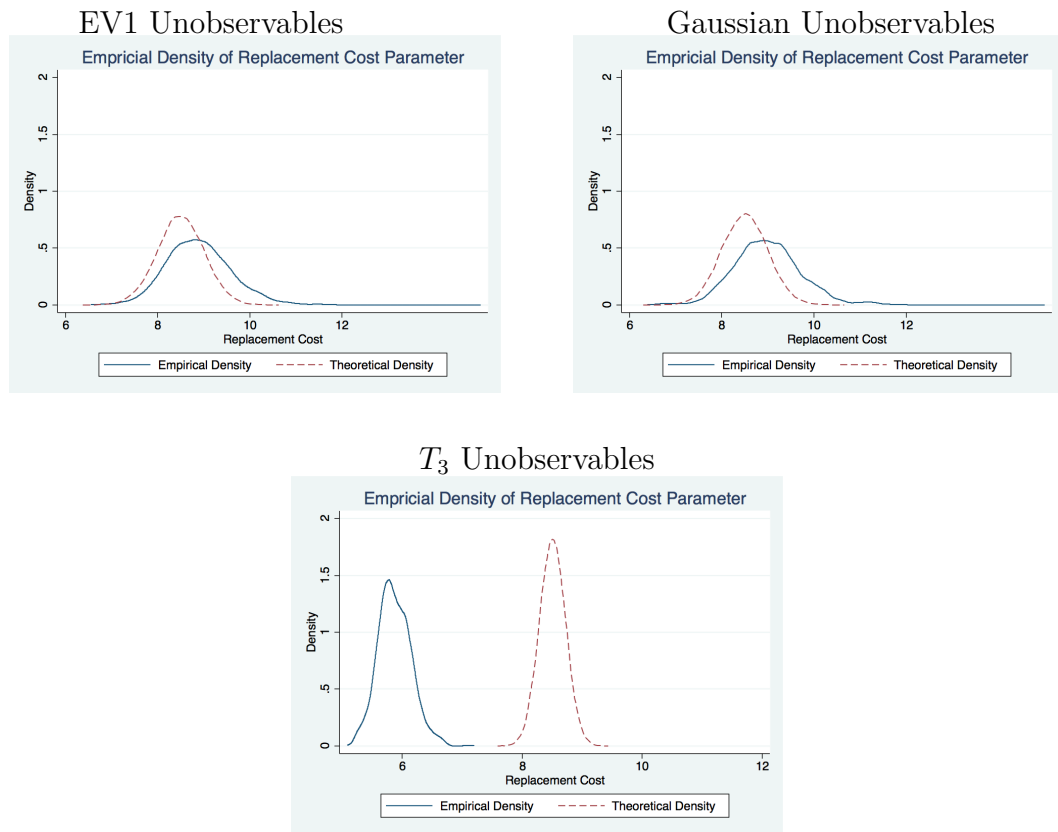


Figure 3.6: Simulations from Different Unobservable Distributions: Cost Function Parameter

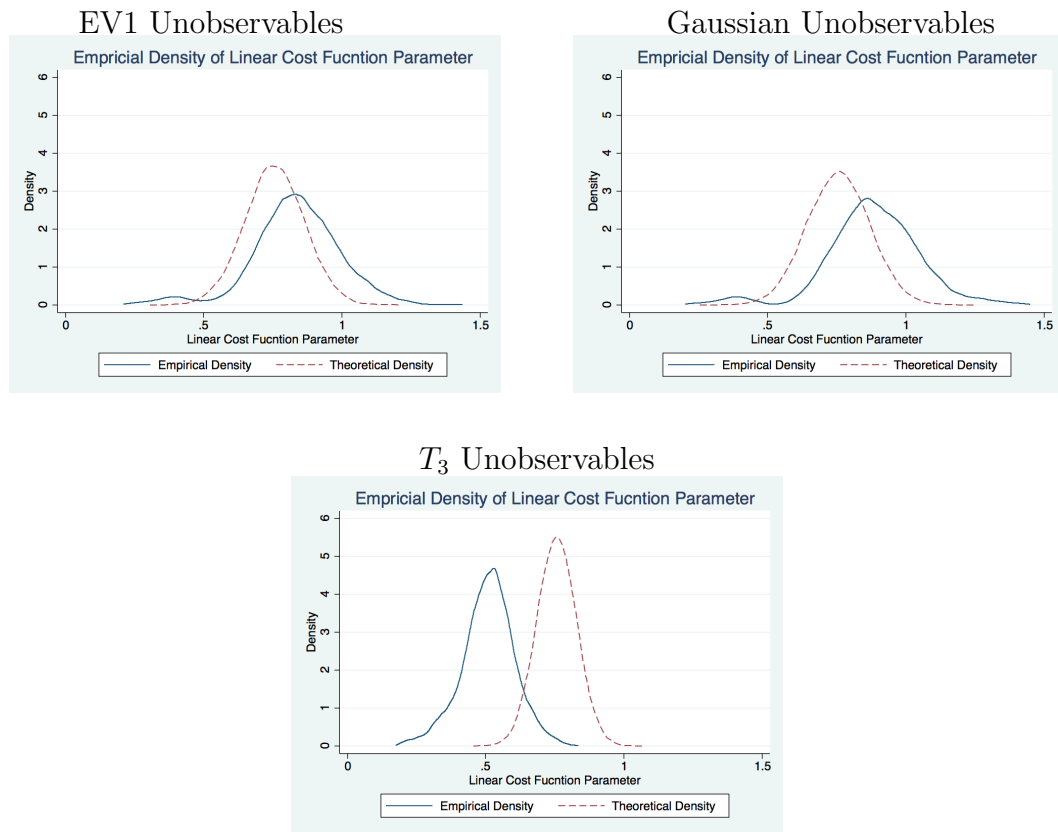


Figure 3.7: Simulations from Different Unobservable Distributions: $P(x_{t+1} - x_t = 0)$

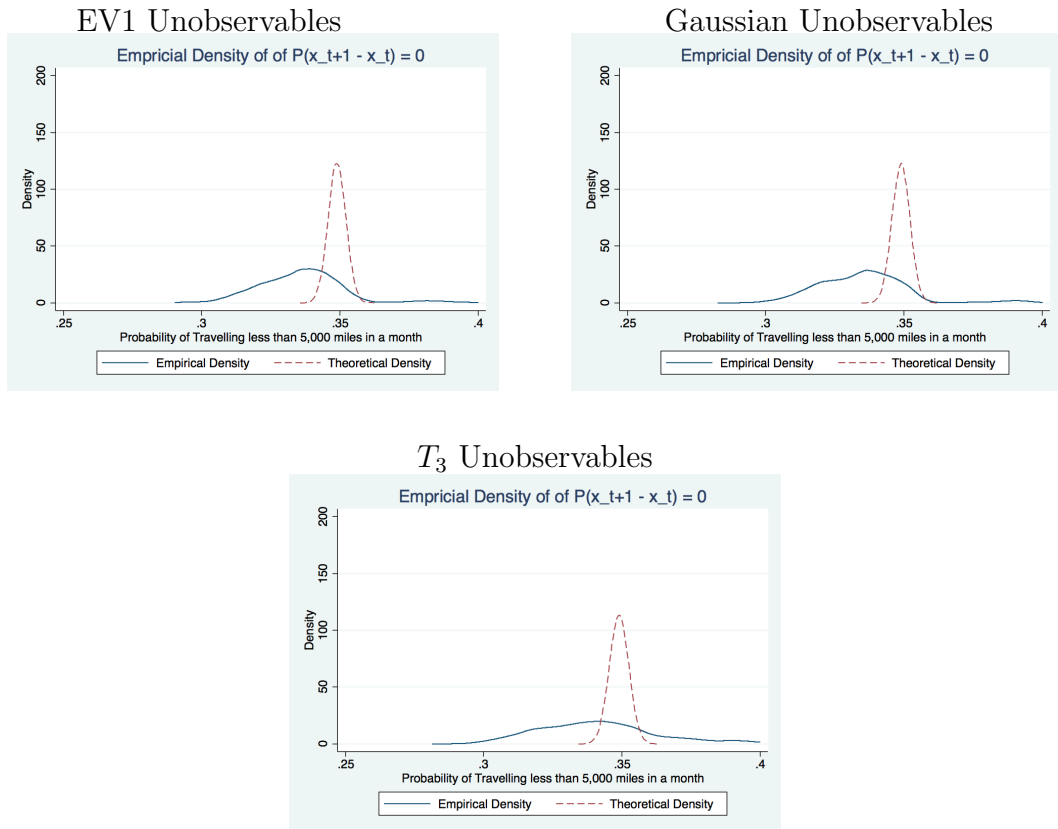
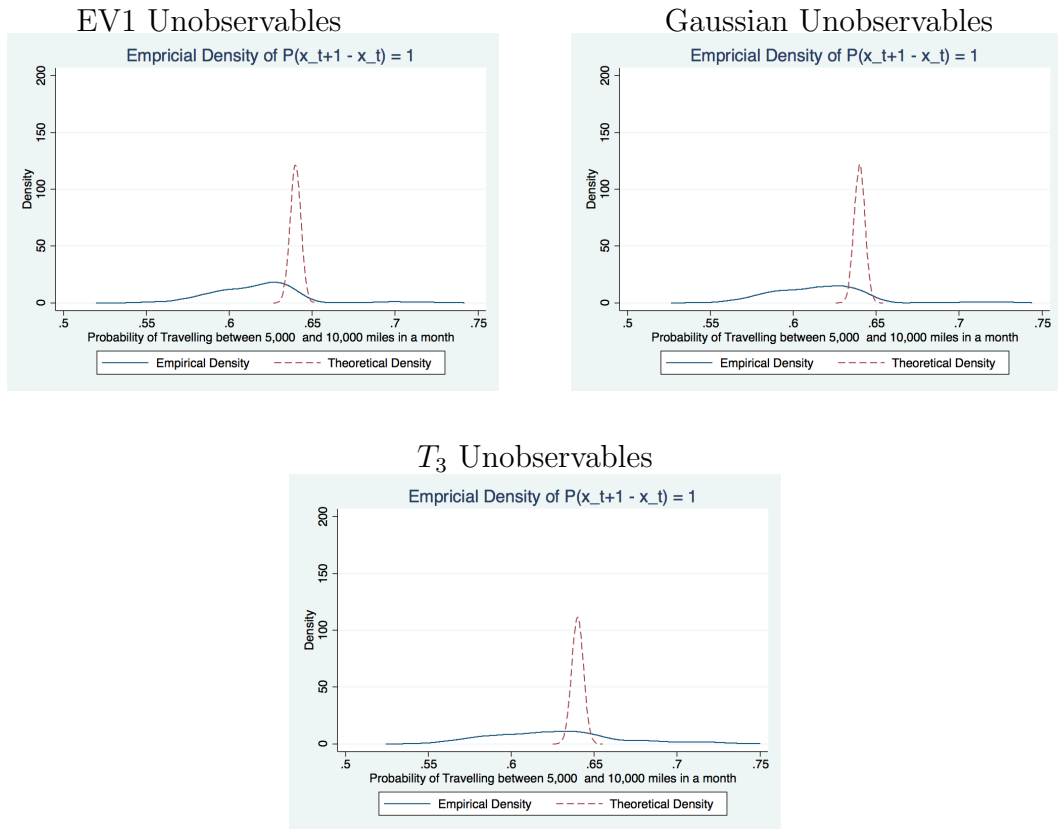


Figure 3.8: Simulations from Different Unobservable Distributions: $P(x_{t+1} - x_t = 1)$



BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] Fact sheet: The regional greenhouse gas initiative (rggi). www.rggi.org.
- [2] *Technical Support Document for the 1992 Citizen's Guide to Radon*. U.S. Environmental Protection Agency, Washington, D.C., 1992.
- [3] Tax expenditures for energy production and conservation. *Joint Committee on Taxation*, May 2009.
- [4] Soji Adelaja and Yohannes Hailu. Projected impacts of renewable portfolio standards on wind industry development in michigan. *Land Policy Institute, Michigan State*, Jan 2008.
- [5] Victor Aguirregabiria and Pedro Mira. Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models. *Econometrica*, 70(4):1519–1543, July 2002.
- [6] Victor Aguirregabiria and Pedro Mira. Dynamic discrete choice structural models: A survey. *University of Toronto Working Paper 297*, 2007.
- [7] Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, July 1995.
- [8] Severin Borenstein. Valuing the time-varying electricity production of solar photovoltaic cells. *Center for the Study of Energy Markets Working Paper 142*, Apr 2005.
- [9] James Bushnell, Carla Peterman, and Catherine Wolfram. Local solutions to global problems: Policy choice and regulatory jurisdiction. *NBER Working Paper 13472*, Oct 2007.
- [10] US Census. Current population reports. *US Census*, 2001.
- [11] US Census. American community survey. *US Census*, 2008.
- [12] Kenneth Chay and Michael Greenstone. Does air quality matter? evidence from the housing market. *Journal of Political Economy*, 113(2):376–424, Mar 2005.
- [13] L.A. Cole. *Element of Risk: The Politics of Radon*. AAAS Press, Washington, D.C., 1992.

- [14] Meredith Fowlie. Incomplete environmental regulation, imperfect competition and emissions leakage. *Working Paper, University of California, Berkeley*, Jan 2008.
- [15] Michael Greenstone and Justin Gallagher. Does hazardous waste matter? evidence from the housing market and the superfund program. *MIT WP*, 05-27:1–59, Jan 2008.
- [16] D.B. Henschel. Analysis of radon mitigation techniques used in existing us houses. *Radiation Projection Dosimetry*, 56:21–27, 1994.
- [17] Stephen Holland, Jonathan Hughes, and Christopher Knittel. Greenhouse gas reductions under low carbon fuel standards? *American Economic Journal: Economic Policy*, 1(1):106–146, Apr 2009.
- [18] V. Joseph Hotz and Robert A. Miller. Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60:497–529, 1993.
- [19] V. Joseph Hotz, Robert A. Miller, Seth Sanders, and Jeffrey Smith. A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies*, 61(2):265–289, April 1994.
- [20] Joshua Kneifel. Effects of state government policies on electricity capacity from non-hydropower renewable sources. *Unpublished Manuscript*, Jun 2008.
- [21] Christopher R Knittel and Konstantinos Metaxoglou. Estimation of random coefficient demand models: Challenges, difficulties and warnings. *Unpublished Manuscript*, 2008.
- [22] Ole Langness and Ryan Wiser. The renewables portfolio standard in texas: An early assessment. *Energy Policy*, 31:527–535, Jan 2003.
- [23] Thomas Lyon and Haitao Yin. Why do states adopt renewable portfolio standards?: An empirical investigation. *The Energy Journal*, 31(3):133–158, May 2010.
- [24] Fredric Menz and Stephan Vachon. The effectiveness of different policy regimes for promoting wind power: Experiences from the states. *Energy Policy*, 34:1786–1796, Apr 2006.
- [25] Gilbert Metcalf. Low-carbon tax policies. *Working Paper*, Jun 2009.
- [26] Robert A. Miller. Estimating models of dynamic optimization with microeconomic data. In H. Pesaran and P. Smidh, editors, *Handbook of Applied Econometrics: Microeconomics: Volume 2*. Blackwell, 1997.
- [27] (BEIR VI) National Research Council, Committee on the Health Risks of Exposure to Radon. *The Health Effects of Exposure to Indoor Radon*. National Academy Press, 1999.

- [28] A.V Nero, M.B Schwehr, W.W Nazaroff, and K.L Revzan. Distribution of airborne radon-222 concentrations in u.s. homes. *Science*, 234(4779):992–997, Apr 1986.
- [29] Whitney Newey and Kenneth West. Automatic lag selection in covariance matrix estimation. *Review of Economic Studies*, 61:631–653, Jan 1994.
- [30] Intergovernmental Panel on Climate Change. Greenhouse gas inventory reference manual. Jul 1997.
- [31] Ariel Pakes. Dynamic structural models, problems and prospects. In C. Sims, editor, *Advances in Econometrics*. Cambridge University Press, 1994.
- [32] Karen Palmer and Dallas Burtraw. Cost-effectiveness of renewable electricity policies. *Energy Economics*, 27:873–894, Nov 2005.
- [33] Nicholas Powers and Haitao Yin. Do state renewable portfolio standards promote in-state renewable generation? *Energy Policy*, 38(2):1140–1149, Nov 2010.
- [34] Barry Rabe. Race to the top. *Pew Center on Global Climate Change*, Jun 2006.
- [35] Barry Rabe. States on steroids: The intergovernmental odyssey of american climate policy. *The Review of Policy Research*, 25(2):105–128, Sep 2008.
- [36] John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, September 1987.
- [37] John Rust. Maximum likelihood estimation of discrete control processes. *SIAM Journal of Control and Optimization*, 26(5):1006–1024, September 1988.
- [38] John Rust. Estimation of markov decision processes. In R.E. Engle and D. McFadden, editors, *Handbook of Econometrics Volume 4*. 1994.
- [39] John Rust. Nested fixed point algorithm documentation manual. *Manuscript*, 2000.
- [40] Wendelin Werner. Newton-like methods for the computation of fixed points. *Computational Mathematics with Applications*, 10(1):77–86, 1984.
- [41] Ryan Wiser, Mark Bolinger, and Galen Barbose. Using the federal production tax credit to build a durable market for wind power in the united states. *Electricity Journal*, 20(9):77–88, Nov 2007.
- [42] Ryan Wiser, Christopher Namovicz, Mark Gielecki, and Robert Smith. Renewables portfolio standards: A factual introduction to experience from the united states. *Lawrence Berkeley National Laboratory Working Paper 62569*, Apr 2007.
- [43] Ryan Wiser, Kevin Porter, and Robert Smith Grace. Evaluating experience with renewables portfolio standards in the united states. *Mitigation and Adaption Strategies for Global Change*, 10:237–263, Apr 2005.

- [44] Ryand Wiser and Galen Barbose. Renewables portfolio standards in the united states - a status report with data through 2007. *Lawrence Berkeley National Laboratory Working Paper 154-E*, Apr 2008.
- [45] Sally Wright. Wind power: Capacity factor & intermittency. *Renewable Energy Research Laboratory, University of Massachusetts at Amherst*, Jun 2004.