The Drivers of Coal Generator Retirements

Rebecca J. Davis J. Scott Holladay Charles Sims*

January 2022

Abstract

We investigate the drivers of coal generator retirement using data on coal generator turnover, delivered coal prices, and wholesale electricity prices. To do this we identify the impact of sunk retirement costs on the probability of coal generator retirement. By pairing an optimal stopping model of firms' generator retirement decisions with the retirement timing of almost 200 coal-fired generators across the U.S., we estimate retirement costs that are not typically disclosed by firms and are not publicly available. Because our real options model cannot impute the retirement costs for coal generators that have not retired, we use a machine learning algorithm to impute retirement cost amounts for active coal generators. With this data, we used a parametric approach to estimate the impact of retirement costs on the probability of retirement and found that a one standard deviation increase in retirement costs resulted in a 0.2 percent reduction in the probability of retirement. We use the real-options model to predict the retirement date for every active coal generator in the U.S. We conduct two counterfactual analyses, modelling a carbon tax of $40 and a fuel cost subsidy. We find that a carbon tax brings forward the average retirement age by more than six years and that a fuel cost subsidy would have to cover more than half the cost of delivered fuel to extend the average life of a coal generator by 5 years.

Keywords: energy; market structure; barriers to exit; optimal stopping; uncertainty; irreversibility

JEL Codes: Q4, L1, L5, H4

1 Introduction

Coal's share of electricity generation in the United States has fallen from around half to a quarter over the last decade. This shift has been driven primarily by the retirement of existing coal-fired generators. These retirements have consequences for the economy (Hoag and Wheeler, 1996; Black, McKinnish, and Sanders, 2005) and environment (Johnsen, LaRiviere, and Wolff, 2019; Holladay and LaRiviere, 2017; Fell and Kaffine, 2018; Linn and Muehlenbachs, 2018; Schiavo and Mendelsohn, 2019). Regions of the country that produce coal are likely to struggle economically as consumption falls. These regional economic impacts have motivated several recent proposals to subsidize coal-fired generation and other technologies with on-site fuel sources (CITE). Coal generators are also one of the largest sources of carbon dioxide in the country and their exit could lead to significant reductions in U.S. carbon emissions. Carbon pricing policies are expected

*Davis: Assistant Professor, Stephen F. Austin State University Department of Economics, davisrj4@sfasu.edu. Holladay: Associate Professor, University of Tennessee Department of Economics and Howard H. Baker Jr. Center for Public Policy Fellow, jhollad3@utk.edu. Sims: Associate Professor, University of Tennessee Department of Economics and Faculty Fellow at the Howard H. Baker Jr. Center for Public Policy, cbsims@utk.edu. This research project was funded in part by the Alfred P. Sloan Foundation grant no. G-2015-14101, "Pre-Doctoral Fellowship Program on Energy Economics," awarded to the National Bureau of Economic Research.

to perpetuate this shift away from coal generators (Cullen and Mansur, 2017a). Whether such policies will alter the trajectory of coal retirements remains an unanswered question.

In this paper, we investigate coal-fired power plant retirements in the U.S., to understand how market forces and environmental regulation have affected the composition of the generating fleet. We predict the expected retirement time for every coal-fired generator in the country and then evaluate how environmental and industrial policy would affect those retirement dates. In our no-policy baseline we find the average remaining life expectancy at coal fired generators is around X years. We find that a $40 carbon tax reduces the average retirement age by six years, compared to a current mean predicted retirement date of XXXX. A coal fuel subsidy would have to be around half of fuel costs to produce a five year extension of average retirement times.

A growing literature attempts to document the drivers of the decline in coal-fired generation. Several studies consider the marginal effects of natural gas prices and/or environmental policies on coal-fired generation generally (Cullen and Mansur, 2017a; Holladay and LaRiviere, 2017; Fell and Kaffine, 2018; Linn and Muehlenbachs, 2018). This marginal approach provides little insight on the long-run response of coal-fired generation to the retirements seen over the past decade. For example, following a relaxation of emission control policies or an increase in natural gas prices, reductions in coal-fired generation will only persist if they were caused by retirements.

A second smaller group of studies focuses on retirements by accounting for the effect of both marginal and fixed costs. Schiavo and Mendelsohn (2019) examine the statistical relationship among coal-fired retirement rates, natural gas prices, and regional air pollution regulations (SO2 trading programs, Clean Air Interstate Rule, and Cross State Air Pollution Rule). They find that air pollution regulations were a more influential factor in coal plant retirements than low natural gas prices. Linn and McCormack (2019) use a structural or computational model of the electricity system to consider the effect of a variety of market forces and policies on the long-run composition of coal in the generation portfolio. In contrast to the reduced form approach taken in Schiavo and Mendelsohn (2019), they find that market forces (i.e, natural gas prices, renewable generation, and electricity consumption) have a much larger effect on coal retirements than regional pollution policies (NOx caps and MATS).

Instead of a post-mortem analysis of the drivers of retirement, we consider the effect of policies that may exacerbate or alleviate future coal retirements. Predicting the effect of policies on the timing of retirement remains a challenge due to data limitations. In particular, retirement costs (decommissioning costs, salvage values, and environmental remediation) are typically not disclosed by utilities but are a key unobservable variable in modelling retirement decisions because they act as a barrier to exit (Caves and Porter, 1976;

Siegfried and Evans, 1994).[1] Previous work on market exit attempts to proxy for the sunk costs associated with market exit using a variety of methods.[2]

We employ a novel 3-step technique to back out the retirement costs implied by retired generators and map those costs onto active generators. Those retirement costs allow us to model retirement decisions for active generators based on the behavior of generators that have retired during our sample. We then use this technique to determine the effect of two key policies (carbon tax and coal fuel subsidy) on the timing of retirement. These results illustrate the relative importance of market versus regulatory drivers of coal-fired generator retirements and adds to the broader literature on the role of environmental regulation on plant entry and exit (e.g., Ryan, Shapiro and Walker, Suzuki 2013).

In step 1, we develop a real options model of power plant retirement decisions. The real options approach (Dixit and Pindyck, 1994) treats retirement as an investment option and captures the uncertainty and irreversibility in the retirement decision. We model the evolution of fuel and electricity prices for each coal-fired generator in the country that was active in 2010. Looking at actual retirement dates for generators that retired between 2010 and 2017, we back out the retirement costs consistent with the observed electricity and fuel prices. While several papers use real options to model entry and exit of electric power plants (Kumbaroglu, Madlener, and Demirel, 2008; Davis and Owens, 2003; Rothwell, 2006; Yang, Blyth, Bradley, Bunn, Clarke, and Wilson, 2008).[3], we use real options to impute unobservable retirements costs from observable retirement decisions.

In step 2, We impute retirement costs for active coal-fired generators based on the estimated retirement costs from the real options model using a machine learning approach. We take the estimated retirement costs from the real options model for the plants that have retired during our sample period as well as a wealth of data on the generators and use LASSO and Regression Forest algorithms to model the retirement costs at our sample of retired plants. Our models explain about 90% of the variation in retirement costs. We then take the model fit to retired plants and use it to predict the retirement costs at currently active plants. We find that these imputed retirement costs are correlated with plant characteristics in ways that are consistent with intuition.

In step 3, we use the real-options model and the mapped retirement costs to predict the retirement date for every coal generator in the U.S. that was active in 2017. To test the validity of our approach, we check our predictions out of sample. We correctly predict 48 of the 69 (70%) generators that retired in 2018 and 2019.

[1]The literature also suggests lower rates of exit in capital intensive industries (Dunne and Roberts, 1991; Rosenbaum and Lamort, 1992; Siegfried and Evans, 1992) and in industries with larger average firm size (Dunne and Roberts, 1991).

[2]Rosenbaum and Lamort (1992) accounted for sunk costs of exiting a market using the ratio of rental payments to assets and primary product specialization ratio, which measures the percentage output actually belonging to an industry. Siegfried and Evans (1992) used the proportion of book value of assets that had not yet been depreciated to reflect sunk costs.

[3]Cesena, Mutale, and Rivas-Davalos (2013) provides a comprehensive review of real options theory applications to investments in electricity generation projects.

We also are able to correctly predict which plants remain online after 2020. The real options model predicts that XXX plants would remain active given the observed fuel and electricity prices since 2018. XXXX of those plants are still online in 2020.

We conduct two counterfactual analyses, modelling a carbon tax of $40 and a fuel cost subsidy. We find that a carbon tax brings forward the average retirement age by more than six years and that a fuel cost subsidy would have to cover more than half the cost of delivered fuel to extend the average life of a coal generator by 5 years. Few papers have attempted to estimate the impact of environmental regulation and market forces on coal-fired power plant retirements. Linn and McCormack (2019) finds that slow demand growth and displacement by natural gas generation had reduced coal plant profits. Our analysis supports the findings of Linn and McCormack (2019) and suggests that there is little scope for policy to change the retirement dates of coal fired power plants.

There is a significant literature evaluating the environmental impacts of energy policy using marginal emissions rates. These papers typically estimate marginal emissions rates over space or time and use those rates to examine the environmental impacts of electric vehicles (zivin2014spatial), bulk energy storage (Holladay and LaRiviere (2018)), wind power (Cullen (2013) and many others. Linn and McCormack (2019) indicate that the long-run environmental impacts may be smaller than these studies suggest. One possible explanation is that marginal emissions rates are sensitive to the fuel mix of generators (**?**). This highlights the importance of the transitional dynamics presented in our paper when assessing the environmental implications of prices and policies in the electricity sector.

2 Coal-Fired Generator Retirement Model

The following illustrates a case where an electricity generation firm (utility) generates electricity with an existing coal-fired generator.[4] The firm receives a flow payoff:

$$\pi(P_E, P_C) = \Big(P_E(t)q_E(t) - P_C(t)q_C(t) - VC(q_E(t)) - FC\Big) \tag{1}$$

where $P_E(t)$ is the wholesale electricity price, $P_C(t)$ is the price of coal, $VC(q_E)$ is the variable operating and maintenance costs, and FC is the fixed levelized capital cost of the generator.[5] q_E is the quantity of electricity supplied by the generator with $\frac{\partial q_E}{\partial P_E} > 0$. q_C is the quantity of coal used to generate q_E with

[4]A generator contains all the equipment needed to produce electricity and typically operates independently. Electric power plants can include multiple fuel generators which can use different fuels. For that reason, we conducted our analysis at the generator level. We use the term "generator" and "unit" interchangeably.

[5]Following Baumol and Willig (1981), we define "sunk costs" as costs that cannot be eliminated even by total cessation of production. In contrast, "fixed costs" are costs that are not reduced by decreases in output so long as production is not discontinued altogether. Thus, not all sunk costs are fixed and not all fixed costs are sunk.

$\frac{\partial q_C}{\partial q_E} > 0$. This relationship between q_C and q_E captures the generation technology of a specific generating unit with newer and more fuel efficient units requiring less fuel to generate an additional unit of electricity.

2.1 Price Uncertainty

While current electricity prices and coal prices are known with certainty, future prices are unknown. For example, prices paid for the firm's on-site stock of coal are known, but future prices of coal paid by the firm are unknown. Future coal price uncertainty can be somewhat mitigated by purchasing coal on long-term contracts. Thus, firms that utilize long-term coal contracts will be less exposed to coal price uncertainty than firms that purchase coal on the spot market. Pindyck (1999) determined that energy (coal, crude oil, and natural gas) prices are mean-reverting by testing a century's worth of data. Following Pindyck (1999), future coal prices were treated as random variables and assumed to evolve according to geometric mean reversion (GMR), $\mathrm{d}P_C = r_{P_C}(\bar{P_C} - P_C)P_C\mathrm{d}t + \sigma_{P_C}P_C\mathrm{d}z_{P_C}$. Here, r_{P_C} is the rate of reversion to the mean coal price level, $\bar{P_C}$ is the long-run mean coal price level, and σ_{P_C} is the standard deviation rate. $\mathrm{d}z_{P_C} = \epsilon(t)\sqrt{\mathrm{d}t}$ is the increment of the standard Weiner process, where $\epsilon(t)$ is a standard normal variate. By not reaching 0 in any finite time (Karlin and Taylor, 1981), GMR prevents any negative coal prices. The rate of reversion to the mean, the long-run mean coal price level, and the standard deviation rate are all allowed to vary by generator to capture differences in the types of coal used and the cost of transporting coal to different units. The Weiner process does not vary across generators. This specification assumes coal market uncertainty is universal but responses to coal market shocks are idiosyncratic to the generator. In short, the stochastic differential equation above describes the time-variant coal price distribution facing a particular generator.

Unexpected shifts in supply and demand also influence the prices generators receive for the electricity they produce. For example, electricity demand is sensitive to weather conditions since weather variations lead to large variations in heating and cooling demand. Electricity supply is subject to uncertainty surrounding entry of new and exit of old generating capacities. Uncertainty in both the supply and demand side of electricity markets introduces volatility into electricity prices (Joskow, 2007). To capture this uncertainty, electricity prices evolve randomly around a long-run mean following a geometric mean-reverting process $\mathrm{d}P_E = r_{P_E}(\bar{P_E} - P_E)P_E\mathrm{d}t + \sigma_{P_E}P_E\mathrm{d}z_{P_E}$. The mean reverting process captures the flat load demand in many parts of the country in recent years. Similar to the coal price process, the parameters that govern the stochastic electricity price process are allowed to vary by generator to capture regional differences in regulated and deregulated electricity markets.

2.2 The Optimal Timing of Retirement

Based on the expectations of future coal and electricity prices, a coal-fired generator will have a nonzero probability of operating at a loss (i.e. generating negative profits for the firm). At some point in the future, a firm may choose to retire a coal-fired generator when losses become too frequent. Retirement instantly eliminates the flow payoff $\pi(P_E, P_C)$ at some sunk retirement cost, K. Retirement costs can vary widely depending on the level of decommissioning. For instance, retirement costs may be minimal if the generator can be retired while the site is maintained in its current condition with little cleanup needed to meet environmental compliance and ensure safety.[6] In contrast, full decommissioning requires substantial sunk costs associated with dismantling all equipment, demolishing structures, and site clean up including wet and dry disposal areas and coal yards. When a generator is retired, there is no legal requirement to demolish the infrastructure. As long as environmental and safety regulations are followed, the site can remain intact and no decommissioning is required. While the generator is operating, it produces electricity $q_E(t) > 0$ using coal $q_C(t) > 0$ with costs $VC(q_E(t)) > 0$ and $FC > 0$, which generates a flow of profits. If the generator is retired, it produces no electricity and uses no fuel but incurs a sunk retirement cost: $q_E(t) = 0$, $q_C(t) = 0$, $VC(q_E(t)) = 0$, $FC = 0$, and $K > 0$.

The objective of a risk-neutral utility is to determine whether and when to retire an electricity generator, t_R, to maximize the generator's expected discounted profits net of any sunk retirement costs. Using traditional discounted cash-flow analysis, the firm would retire the generator when the expected net present value of generation profits is less than the cost to retire the generator. However, since the costs associated with retirement are sunk, there is an incentive (an option value) to delay retirement longer than suggested by discounted cash-flow analysis. This option value captures the economic value to a firm in being able to respond to new information about coal and electricity markets. The size of this option value is key to determining the timing of coal-plant retirements and will vary by generator depending on coal and electricity market conditions, the efficiency of the coal-fired generation technology currently being utilized, and the sunk costs required to retire.

At each instant in time, the firm must determine whether to continue operating the coal generator or retire it given that all future retirement decisions are made optimally. Given the discount rate δ, the optimal retirement time satisfies the following:

$$V(P_{E_0}, P_{C_0}) = \max_{t_R} \mathrm{E}_0 \left[\int_0^{t_R} \pi\Big(P_E(t), P_C(t)\Big) e^{-\delta t} dt + \Big\{ V\Big(P_E(t_R), P_C(t_R)\Big) - K \Big\} e^{-\delta t_R} \right] \tag{2}$$

[6] Environmental compliance includes adhering to the EPA's Interstate Air Pollution Transport Rule, National Emissions Standards for Hazardous Air Pollutants regulations, the industrial waste rule for fossil fuel combustion waste, the Cooling Water Intake Structures rules of the National Pollutant Discharge Elimination System, the Steam Electric Power Generating Effluent Guidelines and Standards, and the most recent Mercury and Air Toxics Standards.

subject to dP_E, dP_C, $P_E(0) = P_{E_0}$, and $P_C(0) = P_{C_0}$. The evaluation at each instant in time maximizes the expected profits from the coal-fired generator from that point forward by making a choice to continue to generate electricity using coal (whose payoff is defined as V) or to retire and incur K.

Following Dixit and Pindyck (1994), the retirement decision can be specified as an optimal stopping problem. Treating retirement as an optimal stopping problem will ensure that the retirement decision maximizes the value of the coal-fired generation asset. While the generator is operating, it not only provides a flow of profits $\pi(P_E, P_C)$, but it also means the firm holds an option - the option to retire the generator when market conditions deteriorate $V(P_E, P_C)$. This option value represents the value of delaying retirement to gain more information about the profitability of coal-fired electricity generation. When the firm retires the generator, sunk retirement costs K are incurred and the option value is terminated - making it an additional opportunity cost of retirement. It is this opportunity cost that causes a more cautious response by the firm in the face of uncertainty.

The generator's unknown value function can be found by employing stochastic dynamic programming with the following Hamilton-Jacobi-Bellman (HJB) equation

$$\begin{aligned}\delta V(P_E, P_C) \geq \pi(P_E, P_C) + r_{P_E}(\bar{P_E} - P_E)P_E\frac{\partial V(P_E, P_C)}{\partial P_E} + r_{P_C}(\bar{P_C} - P_C)P_C\frac{\partial V(P_E, P_C)}{\partial P_C} + \\ \frac{1}{2}\sigma^2_{P_E}P_E^2\frac{\partial^2 V(P_E, P_C)}{\partial P_E^2} + \frac{1}{2}\sigma^2_{P_C}P_C^2\frac{\partial^2 V(P_E, P_C)}{\partial P_C^2} + \sigma_{P_E}\sigma_{P_C}\rho P_E P_C\frac{\partial^2 V(P_E, P_C)}{\partial P_E \partial P_C}\end{aligned} \quad (3)$$

ρ is the correlation coefficient between the two stochastic processes $P_E(t)$ and $P_C(t)$: $\rho = corr(\mathrm{d}z_{P_E}, \mathrm{d}z_{P_C})$. In financial terms, the firm faces an obligation to a flow of profits and option value before retirement. The obligation is treated as an asset whose value $V(P_E, P_C)$ must be optimally managed (i.e. maximized). The left-hand side of (3) is the return the manager would require to delay retiring the generator over the time interval dt. The right-hand side of (3) is the expected return from delaying retirement over the interval dt based on expectations of future coal and electricity prices. This equation acts as an equilibrium condition ensuring a willingness to delay prior to retirement.

The HJB equation in (3) is a non-homogenous, second order partial differential equation necessitating numeric methods to approximate the coal retirement value function (Miranda and Fackler, 2002). Like any differential equation, the HJB equation must be solved subject to a boundary condition. We approximated the solution to the HJB equation by using the well-known value matching condition (Dixit and Pindyck, 1994)

$$V(P_E, P_C) = K \tag{4}$$

which acts as a boundary condition between the region of the state space where it is optimal to continue operating the coal-fired generator and the region of the state space where it is optimal to retire. We approximated $V(P_E, P_C)$ over a subset of the state space using piecewise linear basis functions (Marten and Moore, 2011). The approximation procedure solves for the 2 x n x m basis function coefficients, which satisfy (3) and (4) at a set of $n = 50$ and $m = 150$ nodal points spread evenly over the two-dimensional state space.[7]

The combination of (3) and (4) ensures that the firm holds the option to retire the generator until the value of the option to retire is equal to the cost of retirement. The solution to the HJB equation (3) and value matching condition (4) can be characterized by a retirement threshold $P_{E_R}(P_C)$ that separates the state space where retirement should occur. Specifically, the retirement curve is the set of points where conditions (3) and (4) are met. Based on expectations of future profits, the firm optimally retires a generator when a decrease in electricity prices or an increase in fuel prices crosses the threshold curve $P_{E_R}(P_C)$.

A closed-form solution for the retirement threshold only exists in the simplest models where the HJB equation is an ordinary differential equation. With uncertainty in both the coal and electricity prices, the HJB equation becomes a partial differential equation.

2.3 Data and Parameter Estimation

The optimal stopping problem described in (3) and (4) is solved for each coal-fired generator that retired in the U.S. between 2009 and 2017. To identify coal-fired generators that retired as of 2017, we utilized the EIA's 860 data on existing generators. This data identifies generators that retired each year on a cumulative basis. Most retirements in the U.S. as of 2017 occurred in the Northeast, Southeast, and parts of the Midwest (Figure 1). We targeted generators that used some form of coal as their stated main energy source for generation.[8]

Estimates of the parameters included in (3) and (4) are required to approximate $V(P_E, P_C)$ over a subset of the state space. We started with a total of 528 retired coal generators as of 2017, restricting retired coal generators to those that have a primary purpose North American Industry Classification System (NAICS) code of 22: electric power generation, transmission, and distribution. For example, the U S Alliance Coosa

[7]Upwind finite difference approximations were used to construct a linear spline, which approximates the unknown value function. We used Matlab along with the CompEcon Toolbox and the smoothing-Newton root finding method to solve the resulting complementarity problem. The approximated state space ranges from 0 to 15 in the P_C dimension with $n = 50$ nodal points and from 0 to 150 in the P_E dimension with $m = 150$ nodal points. Extending the state space in either the $P_C(t)$ or $P_E(t)$ dimension or increasing the number of nodal points beyond 50 and 150 did not alter our general results.

[8]Coal includes anthracite, bituminous, lignite, sub-bituminous, waste, refined, and coal-derived synthesis gas.

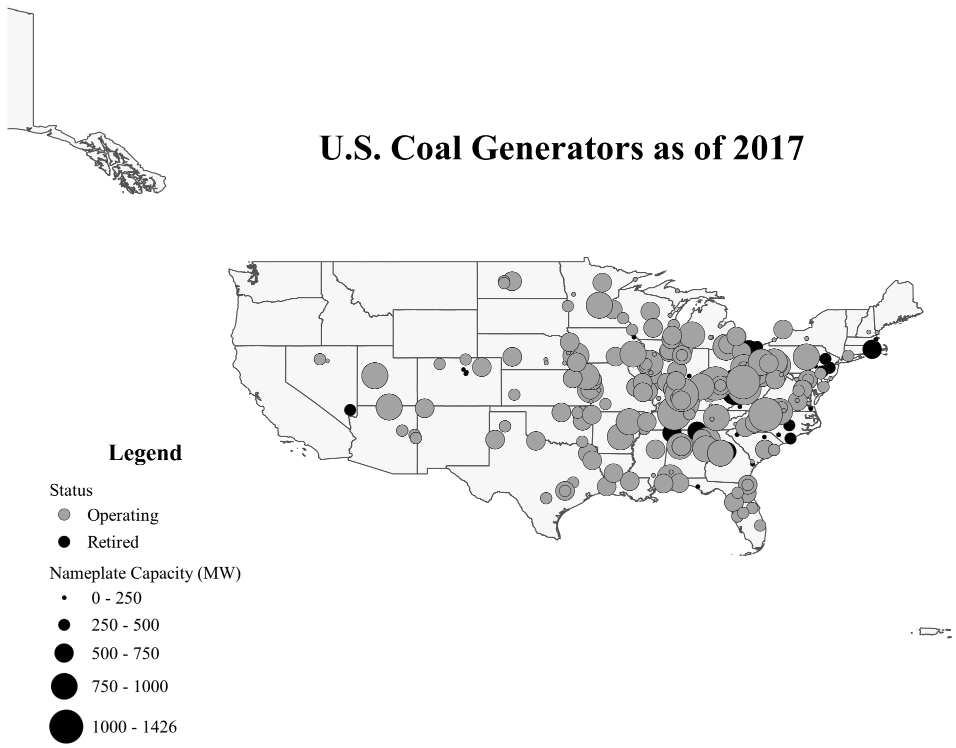

Figure 1: Retired vs. Operable Coal Generators by Generating Capacity

Pines coal plant retired in 2008, but its sole purpose was to provide electricity for a paper plant. This restriction eliminated generators that produced electricity but were never connected to the electricity grid and dropped the total number of retired coal generators by 74 generators. Uncertainty in electricity and coal prices would play a much smaller role in the decision to retire coal generators not connected to the grid.

A total of 454 coal generators with a primary purpose of electric power generation were reported as retired in the EIA 860 data as of 2017. While estimating the parameters found in (3) and (4), we restricted our analysis to coal-fired generators that retired after 2005, because we did not have electricity price data for years prior. Of those 454 coal generators, over 88% occurred after 2005 and over 38% happened after 2014. The sample was restricted further due to lack of coal price data and operating and maintenance data. In the end, 267 coal generators retired between 2009 and 2017. Of that total, 71 were in deregulated electricity markets, and the other 196 were in a regulated market. We calibrated the model for each individual coal generator. We followed Jenkin, Feldman, Kwan, and Walker (2019) and assumed firms used an 8.5% discount rate, $\delta = 0.085$.

Table 1: Average Retired Coal-Fired Generator Parameters by Market Type

Description	Parameter	Regulated	Deregulated
Coal Price Rate of Reversion	r_{P_C}	9.92%	13.36%
		(6.41)	(16.50)
Coal Price Long-Run Mean	$\bar{P_C}$	\$3.05 per MMBtu	\$2.84 per MMBtu
		(0.89)	(0.60)
Coal Price Volatility	σ_{P_C}	9.75%	12.02%
		(5.23)	(7.30)
Electricity Price Rate of Reversion	r_{P_E}	1.83%	1.46%
		(0.81)	(0.39)
Electricity Price Long-Run Mean	$\bar{P_E}$	\$11.27 per MMBtu	\$14.11 per MMBtu
		(2.58)	(1.98)
Electricity Price Volatility	σ_{P_E}	18.61%	25.44%
		(5.87)	(4.10)
Correlation Coefficient	ρ	-9.49%	-12.16%
		(32.20)	(27.29)
Quantity of Electricity	q_E	$q_E = 14,094P_E$	$q_E = 15,010P_E$
		$(19,766)$	$(20,676)$
Quantity of Coal	q_C	$q_C = 3.38q_E$	$q_C = 3.21q_E$
		(0.73)	(0.54)
Variable Costs	$VC(q_E)$	$VC = 0.05q_E$	$VC = 0.20q_E$
		(0.08)	(0.76)
Fixed Costs	FC	$FC = 9.32\bar{q_C}$	$FC = 9.32\bar{q_C}$
		(-)	(-)
Discount Rate	δ	8.50%	8.50%
		(-)	(-)
Sample Size	-	196	71

2.3.1 Electricity and Coal Prices

A critical step in solving the optimal stopping model is defining firm expectations over electricity prices (P_E) and coal prices (P_C). Coal price and quantity data for plants in our analysis came from the EIA's 923 database, which includes information on monthly fuel receipts such as the quantity and price of fuel delivered to a plant. Since an electric power plant can have more than one fuel delivery per month, we used a weighted average of the fuel-specific quantity and price delivered each month to compile monthly fuel prices (\$/MMBtu). Delivered fuel quantities do not accurately describe the fuel used in electricity generation on a monthly basis. Often plants have coal stockpiles, so the quantity of coal delivered to the plant is not all used in the month in which it was delivered.

Real options results critically depend on choosing the correct stochastic process (Tsekrekos, 2010). That process may be one where the price follows a random walk with drift, like geometric Brownian motion, or one where the price reverts back to a trend line, like GMR. Prices that statistically follow GMR would not be consistent with geometric Brownian motion. Unit root tests provide a platform to determine whether a time series follows a random walk. An augmented Dickey Fuller test was used to check the GMR assumptions for coal prices. We rejected the null hypothesis that the price process follows geometric Brownian motion for each generator in the study (see the Appendix).

With empirical support for our geometric mean reversion assumption, we follow Pachamanova and Fabozzi (2011) to estimate r_{P_C}, $\bar{P_C}$, and σ_{P_C} using data from 2002-2017.[9] The coal-fired generators that retired in a regulated market experienced, on average, less volatile but higher coal prices than those in deregulated markets. The average long-run mean coal price is \$3.05 per MMBtu for regulated coal generators and \$2.84 for the retired coal generators in deregulated markets. Average coal price volatility is 9.75% for regulated markets and 12.02% for deregulated markets. The average rate of reversion to the mean coal price across regulated coal generators is 9.92% and 13.36% for deregulated coal generators. Coal generators that retired in a deregulated market have coal prices that revert back to the long-run mean over 34% faster than their counterparts in a regulated market.

Firm expectations of electricity prices are calculated using historic data on wholesale electricity prices (\$ per MMBtu). For all other retired coal generators, we used FERC Form 714 hourly system lambda electricity prices by balancing authority area aggregated to the monthly level to estimate parameters for $\mathrm{d}P_E$. Again, we checked the GMR assumption for electricity prices with unit root tests and rejected that they follow Brownian motion for each balancing authority in our analysis (see the Appendix). Parameters for r_{P_E}, $\bar{P_E}$, and σ_{P_E} were estimated following Pachamanova and Fabozzi (2011) for the calculations (see the

[9] Pachamanova and Fabozzi (2011) rewrite $\mathrm{d}P$ as a linear function and use ordinary least squares to estimate the coefficients r_{P_C}, $\bar{P_C}$, and σ_{P_C} (see the Appendix).

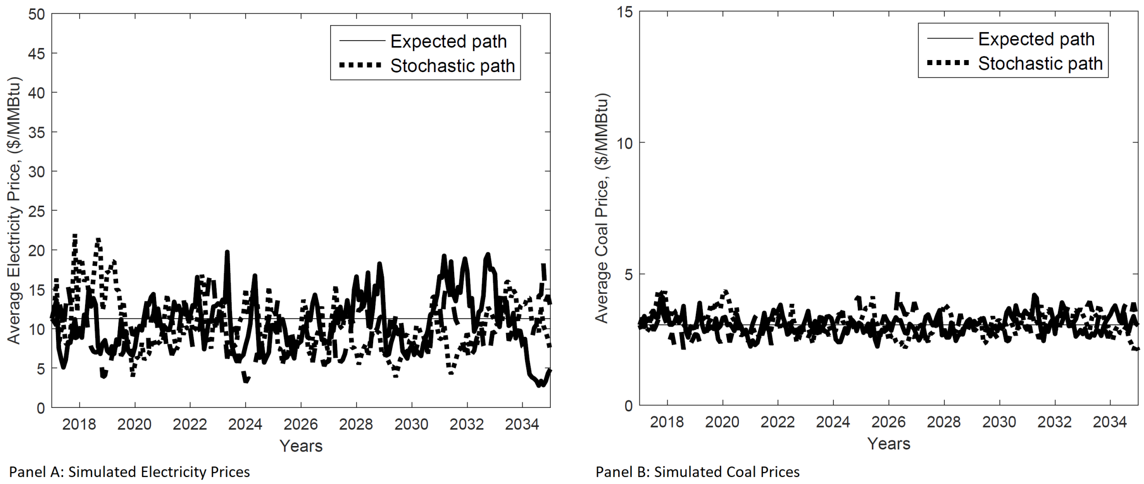

Figure 2: Expected and Stochastic Paths for Coal and Electricity Prices. The expected and stochastic paths are simulated using the average $\mathrm{d}P_E$ (Panel A) parameters and $\mathrm{d}P_C$ (Panel B) for retired regulated generators.

Appendix). Since the electricity price data were either at the balancing authority level, the parameters for $\mathrm{d}P_E$ are the same for generators within the same balancing authority. Unlike coal prices, the generators that retired in a regulated market experienced, on average, lower but less volatile electricity prices than those in deregulated markets. Retired coal-fired generators received average long-run electricity prices of \$11.27 per MMBtu in regulated markets and \$14.11 per MMBtu in deregulated markets. The average volatility of electricity prices experienced by generators in regulated and deregulated markets is 18.61% and is 25.44%. On average, electricity prices revert back to the long-run mean in regulated markets over 25% quicker than their counterparts in deregulated markets.

Electricity prices are much more volatile than coal prices (Figure 2).[10] The thin horizontal line represents the expected path of electricity prices (Panel A) and coal prices (Panel B) for regulated coal generators given the parameters described above. The dashed, dotted, and thick line represent three possible price paths over the next 18 years (2017-2035). The real options approach captures the effect of the volatility shown in both graphs on the retirement option value.

Coal's major use is electricity and coal-fired generators are often the marginal producers. This creates a clear link between coal and electricity prices. To capture the relationship between stochastic processes, we used the cross-correlation function of the two time series to calculate the correlation between coal prices and electricity prices. The average correlation coefficient for retired coal generators in regulated markets is -0.09. In deregulated markets, it is -0.12.

[10]Using the same y-axis for both panels makes the difference between coal price volatility and electricity price volatility appear more extreme.

2.3.2 Electricity and Coal Quantities

The quantity of electricity generated and fuel consumed is available at the generator level from the EIA's 923 database. Electricity and coal quantity data are reported in MMBtus. We collected monthly EIA 923 data from 2009 to 2017.

Electricity supplied by a generator q_E follows a simple linear supply function: $q_E = \beta_{P_E} P_E + \epsilon$. We use OLS to estimate the slope of each generator's supply curve from the quantity and price data in EIA 923. We suppress the constant so a generator would not supply electricity if the price of electricity is zero. The average β_{P_E} for retired coal generators in a regulated market is 14,094, and in a deregulated market, the average β_{P_E} is 15,010. To put those coefficients in context, it takes around 0.064 MMBtus to increase the temperature in a 1600 square foot home with 10 feet ceilings by 50 degrees Fahrenheit. If that house were located in Knoxville, Tennessee, during the winter, it would take only 0.064 MMBtus to heat the home to 70 degrees from 20 degrees. This means that the average quantity of electricity supplied by the retired coal generators in our analysis is highly responsive to changes in electricity prices.

Because coal units are used to generate baseload electricity, owners of coal-fired units seek to minimize the cost of supplying electricity. This means the quantity of electricity generated and the efficiency of the generation technology determines the amount of fuel used. To capture generator technology and efficiency, we used a simple ordinary least squares regression with EIA 923 data on electricity quantity and fuel quantity (in MMBtus): $q_C = \beta_{q_E} q_E + \epsilon$. Suppressing the constant eliminates the ability of a generator to create electricity without using any coal. This relationship is representative of an inverse production function. The higher β_{q_E}, the less efficient is the unit. For example, a coefficient of 3 for β_{q_E} means that a 1 MMBtu increase in electricity requires 3 additional MMBtus of coal, whereas a coefficient of 1 describes a generator that produces the same additional MMBtu of electricity with 2 fewer MMBtus of coal.

An important factor for coal generator retirements is the age of the generator. We were able to account for the effect of age on retirement through β_{q_E} (Figure 3). The slope of the inverse production function for each generator mapped against the first operational year of the generator. The negative slope shows that more efficient generators (small β_{q_E}) are associated with younger generators. The average efficiency for retired coal generators in our analysis is 3.38 for those in a regulated market and 3.21 in deregulated markets.

2.3.3 Fixed and Variable Costs

Electricity producers face fixed costs (e.g., capital costs, financing costs) and variable or O&M costs (e.g., fuel, labor, maintenance). The EIA's 923 database contains financial information at the plant level. Specif-

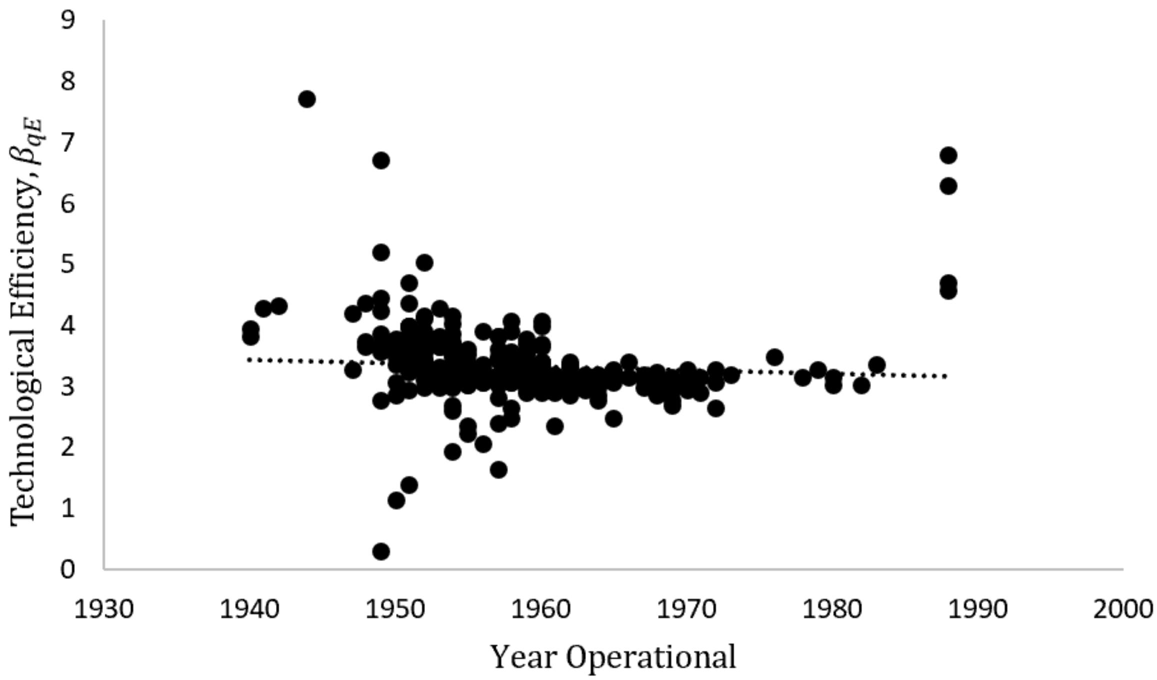

Figure 3: Retired Coal Generator Technological Efficiency by Operational Year.

ically, O&M expenses associated with environmental abatement are reported by year.[11] We combine this information with the electricity generation data mentioned above from 2009-2017 to calculate a measure of O&M expenditures in dollars per MMBtu of electricity generation. We use the average O&M cost at a plant for each coal generator housed there, meaning two different coal generators at the same plant will have the same O&M costs. The average O&M costs for regulated generators is \$0.05 per MMBtu and \$0.20 for deregulated generators. The total variable costs associated with regulated generators is $VC = 0.05q_E$ and $VC = 0.20q_E$ for deregulated generators in our model.

We assume that fixed costs are equal to the levelized capital costs. Levelized capital costs are calculated by taking the total capital costs and dividing by the total life of the generator in terms of megawatts of electricity generated. A coal-fired generator that is still operating incurs these levelized capital costs over its productive life.

According to Logan, Marcy, McCall, Flores-Espino, Bloom, Aabakken, Cole, Jenkin, Porro, Liu, Ganda, Boardman, Tarka, Brewer, and Schultz (2017), the fixed O&M for coal generators that utilize pulverized coal averages \$9.32 per MMBtu. Each generator's fixed costs are found by multiplying the per unit fixed O&M cost by the average quantity of coal (in MMBtus) used in electricity generation from the EIA 923 data

[11] There are additional variable costs faced by a generator other than those related to environmental abatement (e.g., labor and maintenance). Thus, our measure of O&M costs are a lower bound that vary by plant.

described above: $FC = 9.32\bar{q_C}$.

3 Retirement Cost Analysis: Retired Coal-Fired Generators

In this section we describe our procedure for backing out retirement costs at coal-fired generators that have exited the market.

We are not aware of any systematic data on the financial costs of retiring a coal-fired generator (K). That data is proprietary and neither deregulated nor regulated utilities are required to report it. More importantly, retirement costs will vary with the level of decommissioning the firm chooses for a site (U.S. Environmental Protection Agency, 2016a; Henson, 2004). News articles and technical reports have published estimates of retirement costs ranging from $4.1 million (Henson, 2004) to $150 million (Gearino, 2013).

If a single generator is being retired at a plant that will remain active, then decommissioning costs will relatively small. If the land is being released for other uses, full remediation is necessary. The cost and extent of cleanup of hazardous materials will vary with the anticipated new use of the site and the type and location of hazardous materials stored or disposed on the property. If this is the only or last coal generator at the plant, then plants with onsite coal ash ponds or solid waste landfills must follow federal and state permit requirements for closure. The firm can also choose to leave the generator intact and simply maintain environmental permits. Retired generators can be removed and used at other locations the firm owns or sold as scrap.

If a firm's retirement decisions are consistent with real options theory, we can back out an estimate of the retirement costs that justified retiring each generator in our analysis. To do this, we find the electricity price and coal price for the month in which each generator retired (or from the last fuel delivery made to the plant) and compare it to retirement thresholds representative of different retirement costs.[12] Because coal prices do not vary considerably over our time frame, it is unlikely that using the last coal price for each generator will lead to unreliable estimates of K.

The retirement thresholds, $P_{E_R}(P_C)$, for two randomly selected coal-fired generators are the solution to the optimal stopping problem in (3) and (4) with specific parameter values (Figure 4, Table 2). These thresholds represent retirement rules for each coal generator in our analysis whose retirement costs are fully sunk and expectations of electricity prices and coal prices are based on historic data. If the current electricity

[12]One drawback in using the month and year in which a generator retired for coal prices is that plants often receive their last coal delivery months before the generator is technically considered retired. Therefore, it may not be possible to observe the coal price for the month in which a generator retires. If a plant operates multiple coal generators and only a subset of these retire, this problem is moot; that plant still receives coal deliveries after the retired generator(s) drops out, meaning the coal price for coal delivered to the plant for use in the remaining operational generators, including the month in which the retired generator(s) leaves the sample. We can use either the coal price from the last delivery made to a plant or the coal price for the month in which a generator retired.

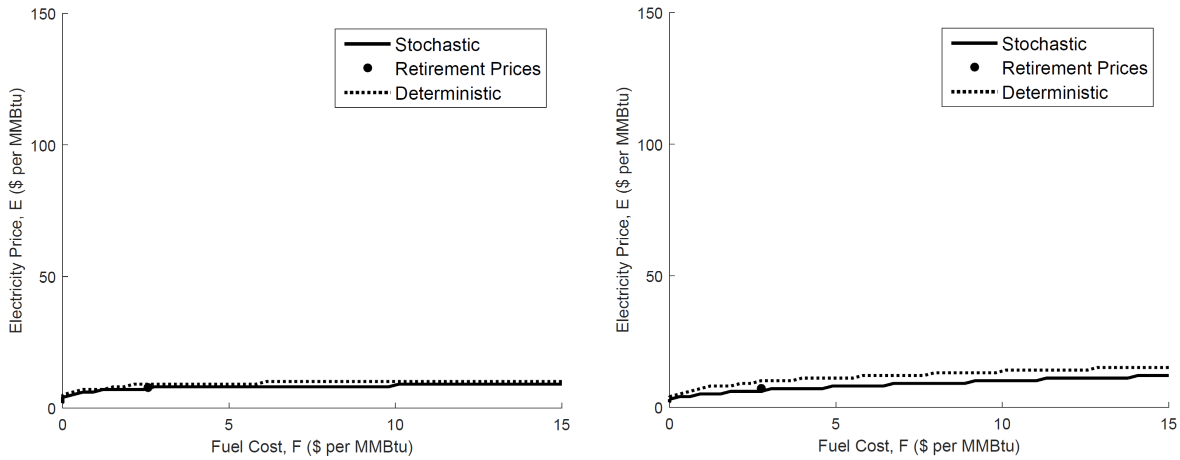

Figure 4: Retirement Thresholds for Two Random Retired Generators with and without Electricity and Coal Price Volatility

and coal prices are below a retirement threshold, it is optimal for a firm to retire that generator. Every combination of prices above the threshold is in the continuation region, that is the point at which it is optimal for the generator to remain operational. As expected, retirement becomes optimal when electricity prices fall. How far electricity prices must fall to trigger retirement depends on how much the firm is paying for coal at that generator. Firms will choose to retire a unit at a higher electricity price if they are faced with higher coal prices (a rightward movement along the x-axis); thus, retirement thresholds are upward sloping. A critical retirement cost, $K = K^*$, is found when a retirement threshold goes from being below the electricity-coal price pair from the month-year in which a generator retired (where the firm would find it optimal to continue operating the generator) to just above the electricity-coal price pair (where the firm would find it optimal to retire the unit). To identify K^* for each generator, we varied retirement costs from \$0 to \$400 million in \$500 thousand increments.

The lines in Figure 4 represent retirement thresholds when electricity and coal price volatility is present. The dashed lines are the retirement thresholds without any volatility. The points identify the prevailing electricity prices and coal prices for the month in which each generator retired. Accounting for electricity and coal price uncertainty is important in determining the optimal retirement decision (Figure 4). Eliminating the two sources of uncertainty that generators face (the dashed line), a firm would retire the right generator at electricity and coal prices much higher than is prescribed by real options theory (the solid line). Because there is a positive benefit to delaying retirement (the option value) due to price uncertainty, the solid line lies below the dashed line. Real options theory tells us that larger drops in electricity prices are necessary to drive a generator into retirement when those prices are uncertain. The estimates of implied retirement costs

Table 2: Parameter Values for Two Randomly Selected Retired Generators

Description	Parameter	Left Graph	Right Graph
Coal Price Rate of Reversion	r_{P_C}	97.73%	12.18%
Coal Price Long-Run Mean	$\bar{P_C}$	\$2.38 per MMBtu	\$3.19 per MMBtu
Coal Price Volatility	σ_{P_C}	34.50%	12.54%
Electricity Price Rate of Reversion	r_{P_E}	2.72%	1.18%
Electricity Price Long-Run Mean	$\bar{P_E}$	\$8.94 per MMBtu	\$18.45 per MMBtu
Electricity Price Volatility	σ_{P_E}	12.74%	30.89%
Correlation Coefficient	ρ	20.83%	-51.27%
Quantity of Electricity	q_E	$q_E = 7,419P_E$	$q_E = 17,226P_E$
Quantity of Coal	q_C	$q_C = 2.23q_E$	$q_C = 2.95q_E$
Discount Rate	δ	8.50%	8.50%
Variable Costs	$VC(q_E)$	$VC = 0.05q_E$	$VC = 0.06q_E$
Fixed Costs	FC	$FC = 9.32\bar{q_C}$	$FC = 9.32\bar{q_C}$
Critical Sunk Cost	K^*	\$19 million	\$70.5 million

are inclusive of the retirement option value. If we did not consider electricity and coal price uncertainty in our analysis, we would over-estimate the implied critical retirement costs. Electricity and coal price volatility have heterogeneous effects across generators (Figure 4). The generator's retirement thresholds on the left show minimal changes in the presence of price volatility compared to no volatility.

Figure 5 shows the frequency distribution of implied retirement costs across the 267 retired coal-fired generators in our study. The gray (black) bars show the number of deregulated (regulated) coal generators that have retirement costs in each bin. Almost 85% of the 267 retired coal generators have critical K values less than \$100 million. One coal generator has critical retirement costs equal to \$0. The barriers to retirement were the lowest for this coal generator; plant managers had to pay nothing to pull this generator out of the market. For the most part, coal generators that retired in a regulated market have the highest retirement costs. We estimated that the two coal generators with the largest retirement costs in our sample, totaling \$335 million when they retired, were part of the Paradise coal plant in Kentucky that has since transitioned to a combined cycle natural gas plant.

The distribution of critical K values is skewed to the left and takes the general shape of a logistic distribution for generators that retired in either market type. The K^* distribution for regulated coal generators that retired between 2009 and 2017 appears to have a long right tail, and a large mass at \$0-\$60 million. The retirement cost distribution for coal generators that retired in a deregulated electricity market is similar to that of their counterparts, but the highest frequency of retirement costs occurs between \$0-\$80 million. The weighted average of retirement costs for all coal generators in our analysis is \$61.63 million. The retirement costs are a barrier to exit for coal-fired generators and hold potentially informative signals about the likelihood of coal-fired generator retirements (Figure 5). In the following section, we estimate the retirement costs for active coal generators utilizing the retirement cost estimates for retired coal generators.

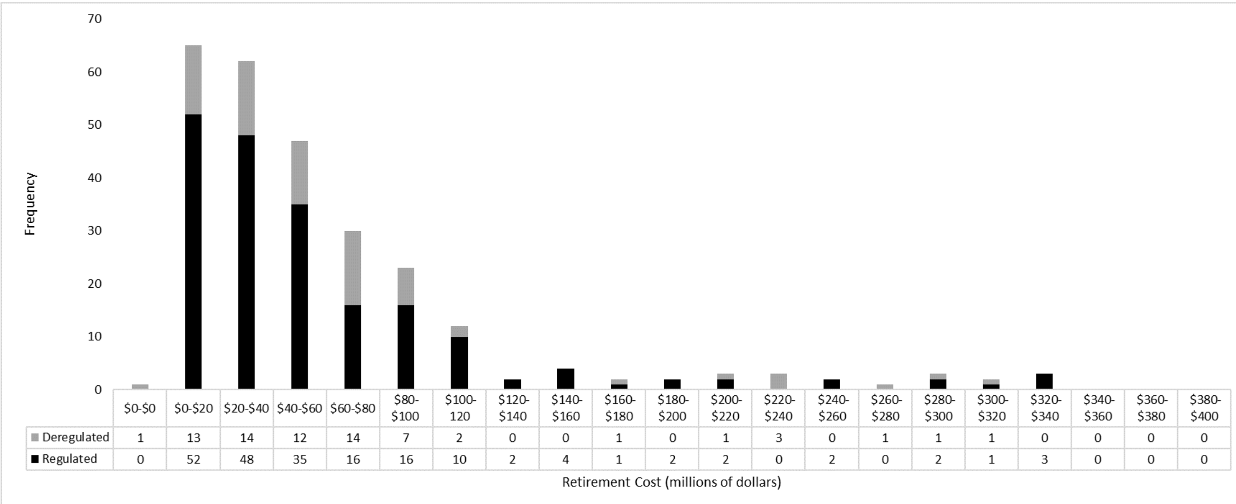

Figure 5: Retirement Cost Distribution for Retired Coal Generators by Market Type

4 Predicting Retirement Costs for Active Generators

In this section we introduce a machine learning approach to estimate the retirement costs at active coal-fired generators based on the imputed retirement costs from the real options model described above. Machine learning algorithms are well suited for this task. We want to produce the most accurate possible predicted retirement costs, but we will not be interpreting the relationships between our explanatory variables and retirement costs. Machine learning techniques essentially provide lower variance in the fitted values for the ability to interpret unbiased coefficients on the relationship between explanatory variables and the dependent variable.

In this section we briefly describe the data needed to estimate retirement costs, then the machine learning algorithms we use to produce the estimates. We use LASSO and Regression Forest algorithms on our generator level dataset using a cross-validation approach to pick the appropriate parameters. We then demonstrate that the estimated retirement costs are reasonable and that those retirement costs can be an important input in to the analysis which plants are likely to retire.

4.1 Data

We make use of several different data sets to estimate the probability of retirement for active coal plants. We take generator characteristics from EIA-860 data, local area information from the Census, pollution emissions from EPA CEMS data and fuel prices from the EIA to build an electric generator by month panel.[13] The EIA-860 data contains unique identifiers that can be used to link EIA and EPA generator records. Both

[13]Most coal-fired power plants are made up of multiple electric generators. The mean plant in our sample has 3.4 coal-fired generators. Generators can be retired separately or taken down in groups as part of a plant retirement.

data sets include the geographic location of the generator which allows us to combine generator characteristic data with data from the 2000 Census on the socioeconomic characteristics of the communities where the generators are located. We use the 2000 Census, rather than contemporaneous data, so that we do not conflate the impacts of a generator retirement on the socio-demographic status of a county with the impact of the county's socio-demographics on the retirement decision.

4.2 Estimating retirement costs for active plants

In this section we describe the machine learning algorithm we use to estimate retirement costs for active generators from the real option model's imputed retirement cost for retired plants. We employ two different methodologies: LASSO and regression forest.

Least absolute shrinkage and selection operator (LASSO) is a regression algorithm that shrinks coefficients in a regression model towards zero. Coefficients that are small in magnitude are set to zero and the remaining coefficients are closer to zero than in OLS. LASSO selects a subset of variables for inclusion in the model and estimates the coefficients for those variables jointly. The researcher chooses a shrinkage parameter, λ that governs how much coefficients are reduced.

LASSO produces improved precision of the fitted values, particularly in the presence of correlated independent variables. LASSO produces biased coefficients on the variables selected into the model, but our focus is only on the estimated retirement costs so this is an acceptable trade off in this application.

We take the dataset set described in section 4.1 and create polynomial and interaction terms for each variable. That produces a dataset with 368 variables and 1064 observations. We include all of those variables in the analysis and allow LASSO to select the subset which best predicts retirement costs on retired coal-fired generators. We choose the λ through cross-validation. We split the dataset of retired plants into a training sample consisting of 80% of the observations and a testing sample of the remaining 20%. We step through potential λ values randomly sampling 80% of the observations each time to estimate the model. We evaluate the model fit on the testing sample that was not used in estimating the model and choose the λ with the best out of sample fit.

To assess the robustness of our results to the LASSO method, we also use a regression forest algorithm to predict retirement costs for active generators. Regression forests are an ensemble method that work by randomly choosing a subset of the variables and creating multiple decision trees. Each tree works by splitting the sample into two groups using the splits that minimize the variance in each sub-sample. This process is repeated up to the maximum tree depth, a parameter specified by the researcher. The sample mean at each leaf at the bottom of the repeated split sample is the fitted value for observations in that bin. Averaging

Table 3: Data Dictionary

Variable	Description	Source
Retire	Indicator for retired generators	EIA 860
Retirement Costs	Imputed cost to retire generator	Imputed from real options model
Efficiency (100%)	Generator efficiency when operating at full capacity	EIA 860
Nameplate Cap (MW)	Generator reported nameplate capacity	EIA 860
Regulated	Indicator for generator in a regulated wholesale electricity market	EIA 860
Ash Impoundment	Indicator for a generator with an ash impoundment	EIA 860
Nonattain	Indicator for a generator located in a county that is regulated under the CAA for any pollutant	EIA 860
Mercury Control	Indicator for a generator with mercury abatement technology installed in 2009	EIA 860
Median Income	Median county income	2000 Census
Pop. Density	Population Density (pop/sq. mile)	2000 Census
Unemployment Rate	County unemployment rate	2000 Census
Male Higher Ed.	Fraction of males with at least a college degree	2000 Census
Female Higher Ed.	Fraction of females with at least a college degree	2000 Census
Lagged Gas Price	Annual average natural gas price lagged 12 months	¡¡BD¿¿

over multiple trees reduces the tendency to overfit in a single tree.

The researcher must choose the max tree depth. This parameter governs how many times the sample is split before sample means are calculated. Deeper trees provide more accurate fitted values in sample, but have a tendency to overfit and predict poorly out of sample. We choose the maximum tree depth through cross validation, following a pattern similar to the one to used to select the optimal λ.

We find that both machine learning algorithms give qualitatively similar fits. The mean square error of the out of sample fit of the models ranges from 8.8% to 11.1% of a standard deviation across algorithms and ranges of λ and tree depths. The estimated retirement costs produced using these algorithms are also highly correlated with each other, with correlation coefficients ranging from X to Y.

The LASSO model's out of sample residual prediction error is 8% of a standard deviation or around $127 thousand dollars. We select this as our baseline algorithm and re-run it for the full sample of both active and retired generators. We will use these predicted retirement costs in our counterfactual analysis in section5.

4.3 Analysis of Retirement Costs

We have two sets of retirement costs. For generators that have retired we have estimated retirement costs backed of the real options model. For active generators we have predicted retirement costs based on the machine learning algorithm described in the previous section. We conduct a simple analysis to confirm that those estimated and predicted retirement costs are reasonable and consistent with intuition.

We began by estimating the probability of retirement for each generator in the sample as a function of imputed retirement costs, generator and local area characteristics. The estimating equation is:

$$Logit[Prob(Retire_{ist} = 1)] = \alpha RetireCost_i + \beta GenChar_{it} + \gamma Census_i + \theta Gas_t + \delta_s + \zeta_m + \eta_y + e_{ist} \quad (5)$$

The dependent variable is an indicator for whether a specific generator (i) is retired in time period t. $RetireCost_i$ is the imputed retirement cost described above, and the parameter α measures how changes in retirement cost affect the probability of retirement all else equal. $GenChar$ is a matrix of generator characteristics including the efficiency of the generator and the nameplate capacity. The $GenChar$ matrix also includes separate indicators for whether the generator is in a regulated wholesale electricity market, has an ash impoundment, has mercury controls, and is in a county that is in non-attainment for any criteria pollution for any year during the sample period. Local sociodemographic conditions may play a role in an operator's decision to retire a generator. For that reason we also included a matrix of variables from the census measured at the county level. The variables include median income, population density, unemployment rate, and the percentage of males and females with bachelor's degrees or higher. We also included lagged natural gas prices to measure competition from other fuel types. Last, we included state (δ_s), month of year (ζ_m), and year of sample (η_y) fixed effects.[14]

The results of this estimation are presented in Table 4. Column 1 reports a univariate regression of our imputed retirement costs on the probability of retirement. We found that increases in retirement cost are associated with a reduced probability of retirement, which is evidence that the imputed retirement costs are good proxies for actual retirement costs. In Column 2 state, month, and year fixed effects add to the specification to control for unobserved spatial, seasonal, and time confounds. The estimated impact of retirement costs on the probability of retirement does not change. Column 3 adds generator characteristics; larger generators with ash impoundments are less likely to be retired. Including these generator characteristics moderates the relationship between the imputed retirement cost measure and the probability of retiring somewhat. Column 4 adds county level sociodemographics to the estimation. We found no strong relationship between county characteristics and the probability of generator retirement. Finally, Column 5 adds natural gas prices lagged six months to the estimation. The coefficient is negative, statistically significant, and large in magnitude, suggesting that low natural gas prices increase the probability of retirement.

The impact of estimated retirement costs on the probability of retirement are consistently negative and statistically significant, which holds true even after including a number of the controls used in the estimation of the retirement costs at active generators. This can be considered to be reassuring suggestive evidence that the estimated retirement costs are a good proxy for the actual retirement costs faced by the owners of coal-fired electric generators. Including logit coefficients means the marginal effects of changes in retirement costs are not obvious. To facilitate interpretation of these coefficients, the estimated marginal effect of retirement costs on the probability a coal-fired generator retires is reported across the support of estimated retirement costs (Figure 6). The marginal effects of increases in retirement costs are negative across the distribution, but less precisely estimated at low levels of retirement costs. There is a small reduction in the magnitude of the marginal effect of increases in retirement costs on the probability of retirement, but any single ten million increase is not significantly different from the previous. The mean retirement cost across the sample is \$71.6 million and the standard deviation is \$45 million. A one standard deviation increase in retirement cost from the mean is associated with an increase in retirement probability of 0.2%.

The logit estimation described above can generate a predicted probability of generator retirement for all the remaining active generators in the sample. We captured predicted values from the logit estimation presented in Column 3, which includes generator characteristics and fixed effects but not county sociode-

[14] We did not include the same variables used in the propensity score matching procedure due to multicollinearity issues. For example, the total cost of existing abatement technology and the indicator for having an ash impoundment measure a generator's level of abatement and are highly correlated.

Table 4: Estimates of the Probability of Retirement

	1	2	3	4	5
Retirement Costs	-0.02***	-0.02***	-0.01**	-0.01**	-0.01**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Efficiency(100%)			16.60	13.05	13.11
			(13.57)	(14.92)	(14.94)
Nameplate Cap (MW)			-0.01***	-0.01***	-0.01***
			(0.00)	(0.00)	(0.00)
Regulated			-0.59	-0.76	-0.76
			(0.61)	(0.61)	(0.61)
Ash Impoundment			-2.41***	-1.89***	-1.89***
			(0.58)	(0.55)	(0.55)
Nonattainment			0.29	-0.05	-0.05
			(0.43)	(0.47)	(0.47)
Mercury Control			0.42	0.78*	0.78*
			(0.47)	(0.47)	(0.47)
Median Income				0.00	0.00
				(0.00)	(0.00)
Pop. Density				-142.13	-139.82
				(877.47)	(878.27)
Unemployment Rate				17.18	17.12
				(46.43)	(46.45)
Male Higher Ed.				0.00	0.00
				(0.00)	(0.00)
Female Higher Ed.				-0.00	-0.00
				(0.00)	(0.00)
Lagged Gas Price					-0.15**
					(0.06)
Constant	-1.12***	-0.93	-11.55	-10.62	-9.76
	(0.19)	(0.85)	(11.73)	(12.70)	(12.68)
Month FE	No	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes	Yes
State FE	No	Yes	Yes	Yes	Yes
Pseudo R^2	0.086	0.342	0.462	0.470	0.471
Observations	43,056	43,056	43,056	43,056	43,056

Note: The dependent variable in each regression is an indicator equal to 1 if the coal fired generator was retired in that month. Retirement costs were imputed from a real options model as described above. Nameplate capacity through mercury control are generator characteristics collected from EIA 860 data. Median income through female higher education are for the generator's county collected from the 2000 census. Lagged gas prices are monthly national average electric power prices for natural gas as reported by EIA lagged by one year. Standard errors, clustered at the plant level are reported in parentheses below the coefficients. * p<.10, ** p<.05, *** p<.01.

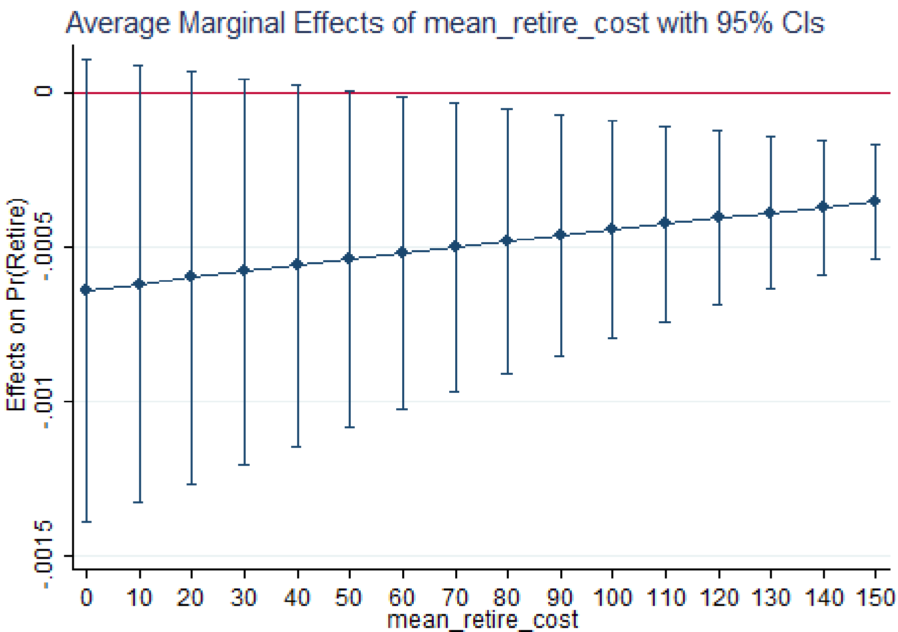

Figure 6: Marginal Effect of Retirement Costs

Note: This figure reports the marginal effect of an increase in estimated retirement costs on the probability of retirement as estimated from the logit reported in Column 5 of Table 4. The vertical lines represent the 95% confidence interval for that marginal effect estimate. The horizontal axis is measured in millions of dollars. The vertical axis is the marginal effect of retirement costs on the probability of retirement. This is estimated at every $10 million interval from $0 to $150 million.

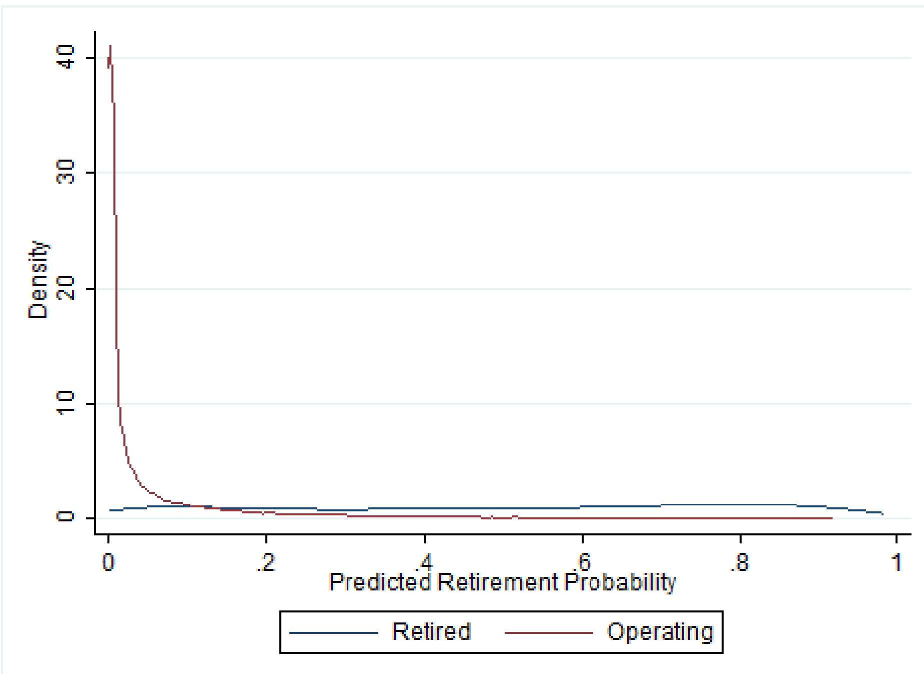

Figure 7: Predicted Retirement Density

Note: This figure reports two separate density regressions of the predicted value of retirement: one for generators that retired during our sample period, and a second for generators that survived throughout the sample. All generators that retired during the sample are included in the density of the probability of retirement for retired generators. The operating generators are limited to those that match across observables to generators that retired during the sample period described above.

mographics or fuel prices. The density for retired generators includes all generators that retire during our sample period (Figure 7). The operating generator density was estimated on a sample of generators identified as similar to the retired generators across the observable generator and county characteristics reported in the previous section. The operating generators' retirement probabilities are clustered at low probabilities. The distribution of retired plants is much more even with a tail of probabilities that are beyond the highest estimated probability for active generators.

As a test of the model's predictive power, we estimated a logit regression using all the variables included in the propensity score matching analysis, as well as retirement costs, and captured the residuals. The residuals represent the predicted probability of retirement for each generator in our analysis. As expected, the generators with the highest predicted probability of retirement are those that in fact retired. More important are the active coal generators with the highest predicted probability of retirement. We compared the top 20 active coal generators with the highest predicted probability of retirement to EIA Form 860's preliminary data on electricity generator retirements for 2016. Of the top 20 retirements predicted based on this analysis, almost half of them did in fact retire in 2016.

To isolate the predictive power of retirement costs, we estimated the same logit regression without retirement costs and collected the residuals. Using those predicted probabilities of retirement, we found the top 20 active coal generators with the highest predicted probability of retirement. They are not the same as those when we included retirement costs, which leads to the conclusion that predicted retirements for active coal generators predicts even fewer actual retirements that occurred in 2016.

5 Counterfactual Analysis

Using the estimated retirement costs for active coal generators, we can conduct counterfactual policy analysis to simulate how different policy options affect the retirement date of the remaining active coal generators. We consider two policies. First, a carbon tax imposed on coal generators through a fuel-used or electricity-generated surcharge. Second, we model how different levels of a fuel reliability subsidy alter retirement time.

5.1 Benchmark

Before we examine the impact of either policy, we must find the benchmark retirement time for each active coal generator utilizing the coal-fired generator retirement model in Section 2 and the estimated retirement costs found in Section 4. The parameters for active coal generators used in the retirement model were estimated using the same data and method as those found for retired coal generators in Section 2. Table

Table 5: Average Active Coal-Fired Generator Parameters by Market Type

Description	Parameter	Regulated	Deregulated
Coal Price Rate of Reversion	r_{P_C}	8.52%	9.12%
		(7.37)	(9.12)
Coal Price Long-Run Mean	$\bar{P_C}$	\$2.50 per MMBtu	\$2.66 per MMBtu
		(0.70)	(0.78)
Coal Price Volatility	σ_{P_C}	7.34%	7.48%
		(4.51)	(3.78)
Electricity Price Rate of Reversion	r_{P_E}	2.22%	1.61%
		(1.07)	(0.57)
Electricity Price Long-Run Mean	$\bar{P_E}$	\$10.39 per MMBtu	\$13.15 per MMBtu
		(2.79)	(2.22)
Electricity Price Volatility	σ_{P_E}	17.09%	23.42%
		(5.22)	(5.54)
Correlation Coefficient	ρ	1.26%	-12.75%
		(29.71)	(29.78)
Quantity of Electricity	q_E	$q_E = 66,510 P_E$	$q_E = 57,287 P_E$
		$(50,531)$	$(41,089)$
Quantity of Coal	q_C	$q_C = 3.13 q_E$	$q_C = 3.19 q_E$
		(0.64)	(0.60)
Variable Costs	$VC(q_E)$	$VC = 0.05 q_E$	$VC = 0.07 q_E$
		(0.05)	(0.08)
Fixed Costs	FC	$FC = 9.32 \bar{q_C}$	$FC = 9.32 \bar{q_C}$
		(-)	(-)
Discount Rate	δ	8.50%	8.50%
		(-)	(-)
Retirement Costs	K	\$189 Million	\$191 Million
		(144)	(143)
Sample Size	-	396	96

5 displays the average parameters by market type for active coal generators. The differences between the model parameters of active generators by market type are similar to those for retired generators. Active coal generators that operate in a regulated electricity markets face lower and less volatile coal and electricity prices on average than those that operate in deregulated electricity markets. The estimated retirement costs for regulated generators average \$189 million compared to the \$191 million paid by deregulated generators at the time of retirement.

With this information, we can find the retirement threshold $P_{E_R}(P_C)$ for all 492 active coal generators. As of 2017, the firms that manage those coal generators found it optimal to continue operating them. However, as coal and electricity prices change over time, that retirement threshold remains the same. If electricity prices drop low enough or if coal prices get high enough, the combination of electricity and coal prices will cross the retirement threshold, making it optimal for the firm to retire the generator. Future electricity and coal prices are unknown. However, firms form expectations over the price processes. Recall, firm expectations of electricity prices and coal prices are found to follow geometric mean reversion (Section

2.3.1): $\mathrm{d}P_E = r_{P_E}(\bar{P_E} - P_E)P_E\mathrm{d}t + \sigma_{P_E}P_E\mathrm{d}z_{P_E}$ and $\mathrm{d}P_C = r_{P_C}(\bar{P_C} - P_C)P_C\mathrm{d}t + \sigma_{P_C}P_C\mathrm{d}z_{P_C}$.

We jointly simulate electricity and coal prices out 50 years on a monthly time step (January 2018 - December 2067), for a total of 600 time steps, utilizing the parameters estimated in $\mathrm{d}P_E$ and $\mathrm{d}P_C$ for each generator. We then compare the simulated electricity and coal price combination against the retirement threshold. The month that the electricity-coal price combination first crosses the retirement threshold determines the retirement time for a generator. Figure 8 shows the retirement threshold for a random active coal generator in our sample. Each dot on the graph represents an electricity-coal price combination that was part of the first simulation. This graph only shows the initial 51 time steps instead of the total 600 time steps, because this particular run of the simulated prices determined that the generator would retire in the 51st month, which translates to March 2021. Notice that the price combination on the far right of the graph lies below the threshold curve. This is the price combination that occurred in the 51st month of the simulation. This is the first price combination that was below the retirement threshold, which signals to the firm to optimally retire the coal generator. The remaining price combinations that were simulated after March 2021 do not matter for this round of the simulation - the generator already retired. This process is completed 10,000 times for each generator. The average retirement time over the 10,000 simulations is recorded as the benchmark retirement time for the generator.

The predicted retirement time for the 492 active coal generators is shown in Figure 9. We found that over 200 (41.1%) coal generators have a predicted retirement time in 2018. Around 80% of active coal generators in our sample are predicted to retire by the end of 2027. That is a total of 190.5 thousand megawatts of generating capacity expected to retire within a ten year window. The remaining generators expected to be operating after 2027 can produce 29.9 thousand megawatts of electricity, which is an 84% reduction in the current coal generating fleet. Figure 10 shows the geographic dispersion of the predicted retirement for the existing active coal generators in our sample.

The benchmark results assume there are no changes in federal regulation that impact the electric power sector during a 50 year period. Further, price expectations are fixed at 2017. As coal generators retire, coal and electricity prices can not be expected to remain the same. Factors like the existing mix of generating capacity and reduced demand for coal will influence these prices. The current model does not account for such feedback. However, the stochastic nature of coal prices plays a very small role in the model. In addition, it is not clear whether the feedback will increase or decrease electricity prices, which play a larger role in our model. To determine how electricity prices would change due to retirements, we would need to expand the current model to account for the entire electricity market a generator operates in. We leave that analysis to future work.

Assuming this feedback between retirement and prices is a long-term issue (greater than 15 years) and

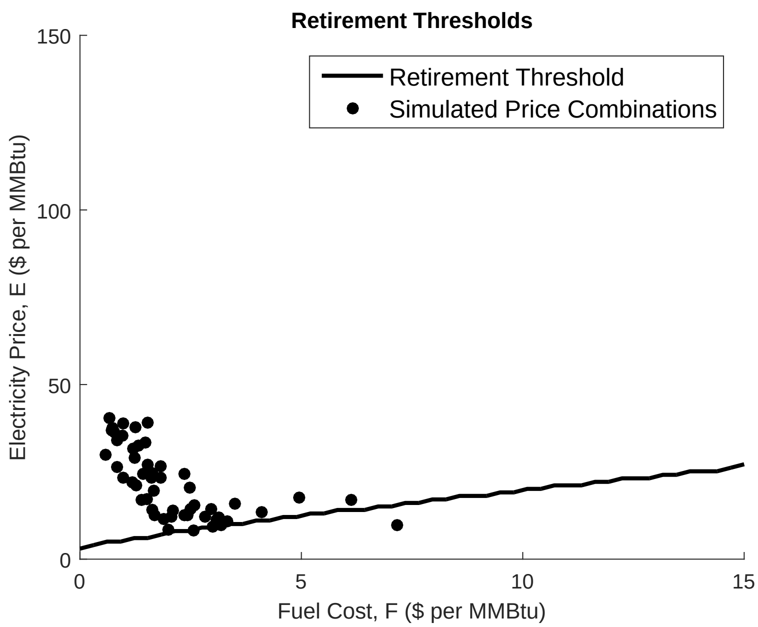

Figure 8: Retirement Threshold with Simulated Electricity and Coal Prices for a Random Active Coal Generator

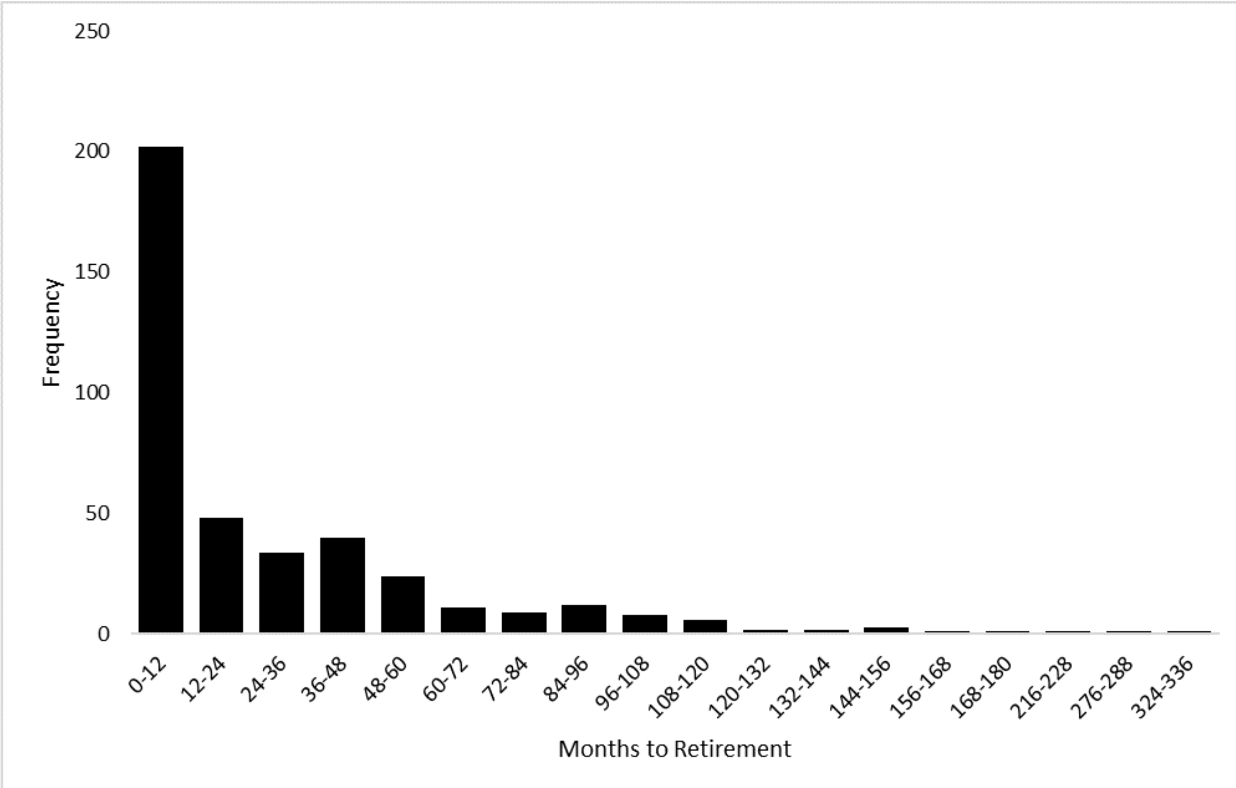

Figure 9: Benchmark Predicted Retirement Time for Active Coal Generators: January 2018-December 2067

that coal demand decreases due to coal generator retirements, we examine the impact of this feedback on coal generators that retire after 15 years by 1. Reducing the long-run mean coal price level by 20% for each generator continuing with the assumption of a GMR price process and 2. Changing the price process to GBM, which represents a market that is in the process of updating better than GMR, and lowering the drift parameter by 20% (see Appendix for details). Completing 1., we find that PUT THAT SENSITIVITY ANALYSIS HERE! Completing 2., we find that PUT THAT SENSITIVITY ANALYSIS HERE!

5.2 Carbon Tax

The EPA lists carbon dioxide as one of the greenhouse gases that traps heat in the Earth's atmosphere, which is a leading cause of global warming (U.S. Environmental Protection Agency, 2020). Carbon emissions are externalities that require government intervention to correct spillover costs. One such intervention is a carbon tax that would force emitters to internalize the full cost associated with carbon dioxide emissions. Currently, there is no federal carbon tax in the U.S. However, eleven states have active carbon-pricing programs.[15] Ten

[15] Those eleven states are California, Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Rhode Island, and Vermont. Washington state has a market-based climate policy that is currently suspended as it progresses in the court system (Ye, 2020).

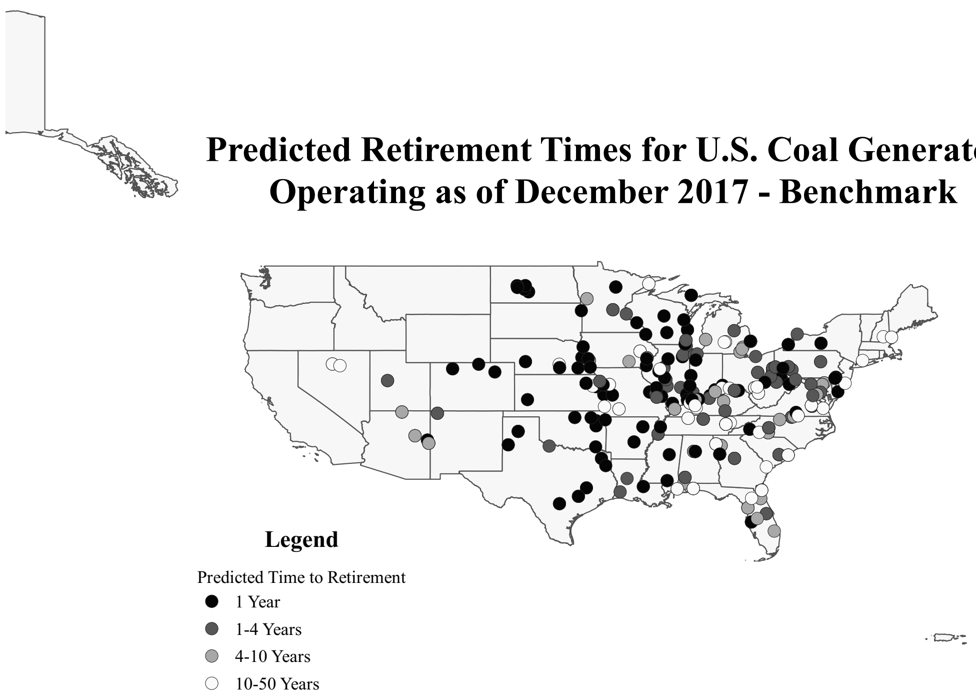

Figure 10: Benchmark Predicted Retirement Time for Active Coal Generators: January 2018-December 2067

of those states participate in the Regional Greenhouse Gas Initiative (RGGI), which is a cap-and-trade program that limits carbon dioxide from the power sector (Ye, 2020). Since its establishment, the RGGI carbon dioxide allowances have sold at prices less than \$2 to prices over \$10. However, we would like to consider a carbon tax policy that covers the entire U.S. electric power sector, instead of one region.

In 2019, three carbon tax bills were introduced in Congress: the Climate Action Rebate Act of 2019 (CAR Act), the Stemming Warming and Augmenting Pay Act (SWAP Act), and the Raise Wages, Cut Carbon Act of 2019 (RWCC Act) (Qian, 2019). The SWAP Act and the RWCC Act propose a \$30 and \$40 per metric ton of carbon dioxide tax. These two acts then increase the tax by 5% and 2.5% each year. By contrast, the CAR Act imposes a \$15 per metric ton of carbon dioxide tax and doubles every year until a certain level. For the purposes of this study, a fixed rate tax is the simplest to model, and it provides a general understanding of the impact of a carbon tax on retirement time for active coal generators.

According to the U.S. Environmental Protection Agency (2016b), the social cost of carbon dioxide is \$42 per metric ton of CO_2, assuming a 3 percent average discount rate in 2020. This is most similar to the RWCC Act that imposes a large tax initially and increases it at a slow rate. For simplicity, we prefer to use the social cost of carbon dioxide and a fixed rate tax to understand the implications on the active coal generating fleet.

The total carbon tax amount paid by each coal generator depends on the amount of carbon dioxide emitted while generating electricity. Depending on the technology of the coal generator, the amount of carbon dioxide created in the generation process varies. Thus, we converted \$42 per metric ton of CO_2 to \$4.11 per MMBtu using the EIA's assessment that the average of all coal types contains 210.20 pounds of CO_2 per MMBtu (U.S. Energy Information Administration, 2016a).[16] A carbon tax of \$4.11 per MMBtu can be imposed on either the amount of coal (input) used to generate electricity or the amount of electricity (output) generated.

5.2.1 Carbon Tax: Input

The coal-fired generator retirement model in Section 2 can be altered to accommodate a carbon tax. The main change to the model occurs in the profit function, Equation 1. The carbon tax, P_T, adds a total cost of $P_T q_C(t)$ when imposed on fuel (input) used, making the profit function the following:

$$\pi(P_E, P_C) = \Big(P_E(t)q_E(t) - (P_C(t) + P_T)q_C(t) - VC(q_E(t)) - FC\Big) \tag{6}$$

where P_T =\$4.11 per MMBtu. As seen in (6), the input carbon tax is an additional cost associated with

[16]Coal types include anthracite, bituminous, subbituminous, lignite, and coke.

operating the coal generator. In turn, the retirement threshold $P_{E_R}(P_C)$ shifts up, meaning the firm optimally retires a generator at higher electricity prices than in the benchmark scenario where the generator does not incur the extra cost of the carbon tax. Figure 11 contains three retirement thresholds for a random active generator. The solid line is the benchmark retirement threshold, the dotted line is the retirement threshold when an input carbon tax is imposed, and the dashed line is the retirement threshold when an output carbon tax is enacted. At every coal price, the corresponding electricity price must be higher to keep the coal generator operating under the input carbon tax scenario compared to the benchmark.[17] Another way of saying this is that there are more combinations of coal and electricity prices that can occur where the firm would find it optimal to retire the generator if there is an input carbon tax imposed.

How does this impact the predicted retirement time for active coal generators? The counterfactual policy analysis requires two steps for completion for each policy. First, we find the retirement thresholds $P_{E_R}(P_C)$ for all 492 active coal generators when there is an input carbon tax. Then, we simulate stochastic electricity and coal prices out and see at what time the electricity-coal price combination crosses the new retirement threshold. The price simulation is identical to the process described in the benchmark scenario - prices are simulated out based on the GMR parameters for each generator for 600 months (or 50 years). The only difference is that the retirement threshold has changed. Theoretically, a carbon tax increases the cost of operating a coal generator. Thus, drops in the price of electricity or spikes in the cost of coal are more likely to push a firm to optimally retire the generator.

As expected, imposing an input carbon tax hastens the time to retirement for active coal generators. If such a carbon tax was implemented in January 2018, almost 87% of active coal generators would retire within the first year. A total of 33 generators would remain active over the entire 50 year time horizon, generating over 6 thousand megawatts of electricity. By contrast, if none of the coal generators in our sample retired, they could produce over 220 thousand megawatts of electricity. Figure 12 compares the predicted retirement time distribution under the benchmark, the input carbon tax, and the output carbon tax scenarios. The input carbon tax reduces the predicted life of the coal generators significantly, regardless of generator capacity (Figure 13).

5.2.2 Carbon Tax: Output

What if a carbon tax was imposed on the amount of electricity generated instead of the amount of coal used to generate that electricity? Does that change the tax incidence compared to an input carbon tax? Do generators retire quicker under an output tax or an input tax? Similar to the method used to include an input carbon tax in the coal generator retirement model, we can attach the carbon tax P_T to the quantity

[17]Discussion of the output carbon tax retirement threshold is in the next section.

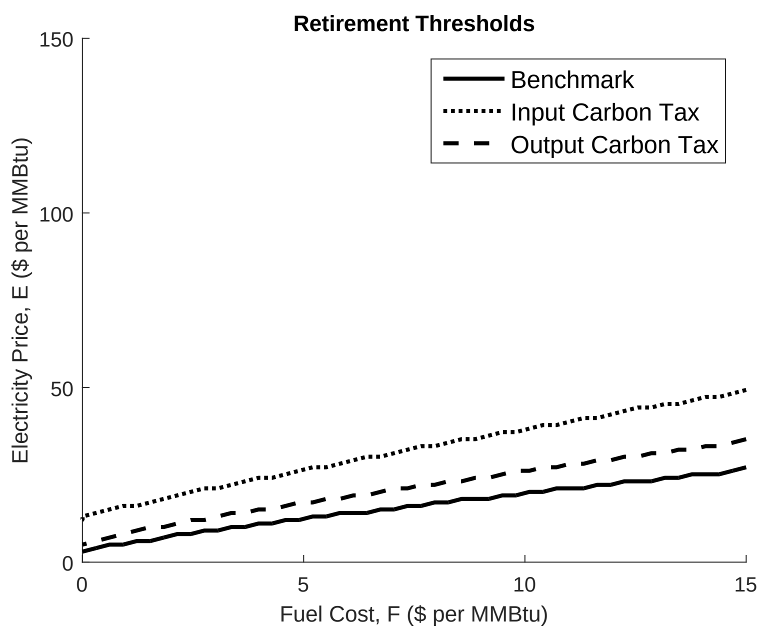

Figure 11: Benchmark Retirement Threshold vs. Input Carbon Tax Retirement Threshold vs. Output Carbon Tax

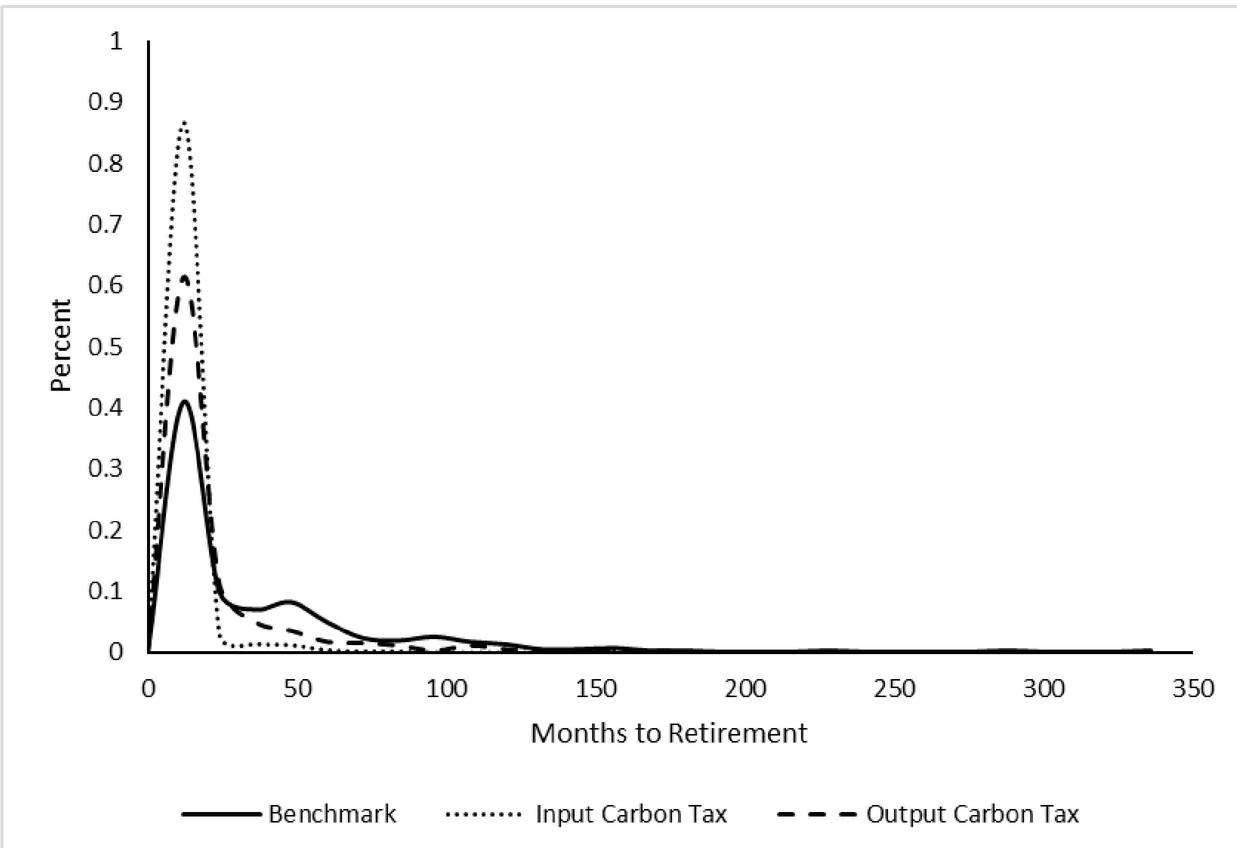

Figure 12: Predicted Retirement Time Distributions for Active Coal Generators: Benchmark vs. Input Carbon Tax vs. Output Carbon Tax

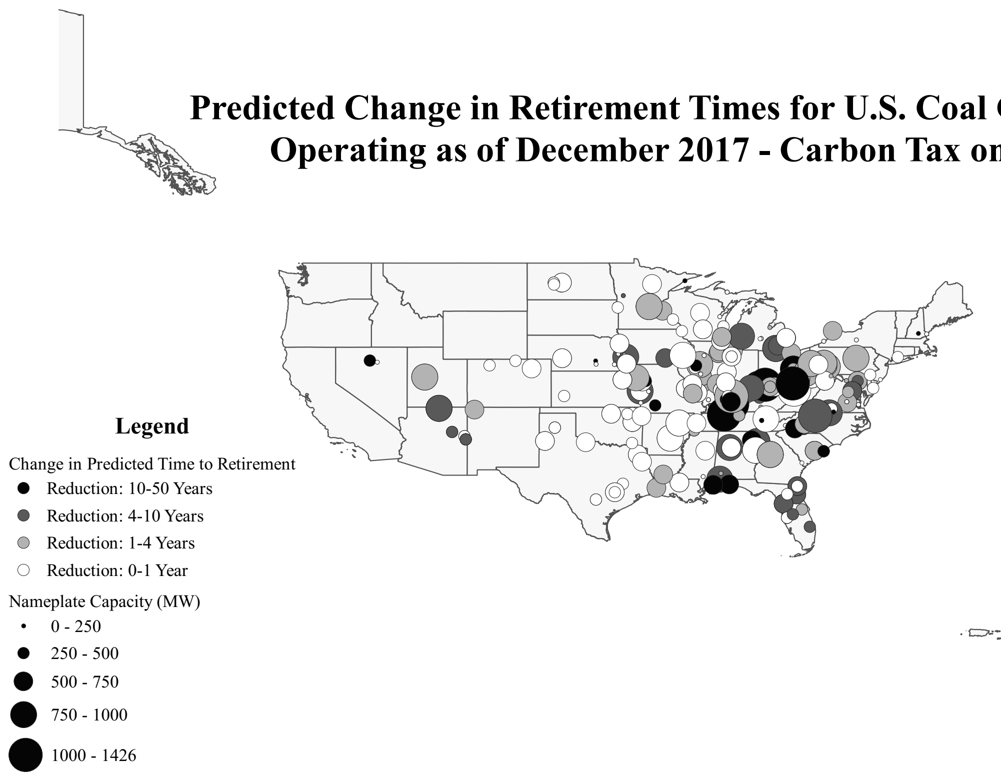

Figure 13: Change in Predicted Retirement Time by Generating Capacity for Active Coal Generators - Input Carbon Tax

of electricity generated $q_E(t)$. The profit function in Equation 1 becomes

$$\pi(P_E, P_C) = \Big((P_E(t) - P_T)q_E(t) - P_C(t)q_C(t) - VC(q_E(t)) - FC\Big) \tag{7}$$

where P_T =\$4.11 per MMBtu. The output carbon tax increases the cost of operating the coal generator, which alters the retirement threshold $P_{E_R}(P_C)$. Again, there are more electricity-coal price combinations where the firm would optimally retire the coal generator when an output carbon tax exists, though the impact on the retirement threshold is much smaller using an output carbon tax versus an input carbon tax (Figure 11).

Equipped with the retirement thresholds that account for an output carbon tax, we simulate electricity and coal prices exactly as described in the benchmark and input carbon tax scenarios. The simulated electricity-coal price combination that triggers retirement is recorded as the predicted retirement time. After 10,000 simulations, we find the average predicted retirement time for each generator. Just over 61% of coal generators would retire in the first year if an output carbon tax is enforced (Figure 12). That would be a total of over 140 thousand megawatts of electricity generation potential retired. There are 64 generators (or 13%) with almost 16 thousand megawatts of generating capacity that would not retire in a 50 year time frame, regardless of an output carbon tax imposition. Interestingly, the coal generators that are pushed to retire the quickest are not necessarily the largest capacity generators (Figure 14).

Notice in Figure 12 that there is an even higher percentage of coal generators that would retire within the first year of a carbon tax enacted if it is imposed on coal (87%) instead of electricity (61%). The per unit dollar value of the carbon tax (\$4.11 per MMBtu) is the same in both scenarios. The reason for the differences in predicted retirement times between the two policies arises from the need for coal generators to use more than 1 MMBtu of coal to generate 1 MMBtu of electricity. The coal to electricity conversion is not one-for-one. Coal generators must use several MMBtus of coal to produce just one MMBtu of electricity. This is reflected in the quantity of coal equation, which captures the generator's technology and efficiency. On average, active coal generators require over 3 MMBtus of coal to generate 1 MMBtu of electricity (Table 5). A carbon tax imposed on the quantity of coal will impact the profit function and the resulting retirement threshold more than if that same per unit dollar value carbon tax was imposed on the quantity of electricity due to the technological nature of coal generators.

5.3 Fuel Reliability Subsidy

While coal generators are a significant polluter for the electric power sector, an argument has been made in their defense. In 2017, the Secretary of Energy, Rick Perry, wrote a letter to the Federal Energy Regulatory

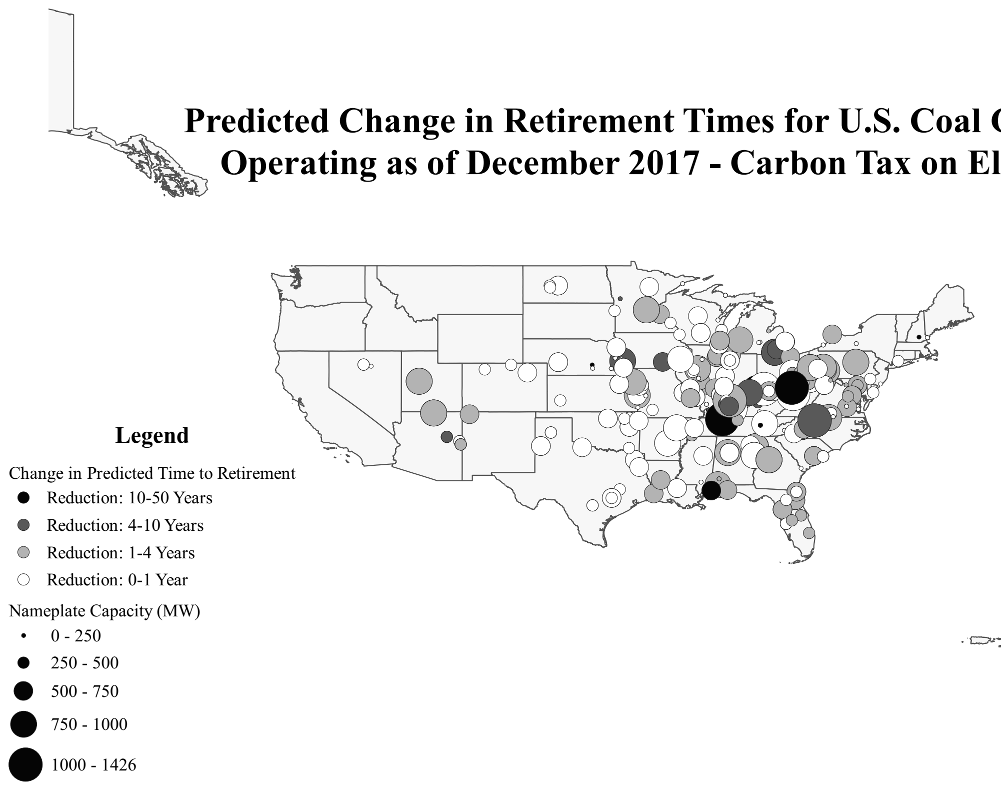

Figure 14: Change in Predicted Retirement Time by Generating Capacity for Active Coal Generators - Output Carbon Tax

Commission explaining that coal and nuclear generating resources were under-valued. This under-valuation stems from distorted price signals that do not reflect the reliability and resiliency benefits of coal and nuclear generators, which keep fuel supplies on-site and have fuel storage capabilities (Department of Energy, 2017). Secretary Perry's proposed rule would allow "for the full recovery of costs of certain eligible units physically located within the Commission-approved organized markets." FERC rejected the proposed rule by stating, "There is no evidence in the record to suggest that temporarily delaying the retirement of coal and nuclear generators would meaningfully improve the resilience of the grid" (Duke SciPol, 2018). While a federal level reliability subsidy for coal does not seem likely given FERC's response, one state agreed with Secretary Perry and took action.

The Creates Ohio Clean Air Program was signed into law July 2019 and became effective in October 2019. The program ensures FirstEnergy Solution's Davis-Besse and Perry nuclear power plants an estimated $150 million a year from 2021-2027 to keep the reactors in service (The Ohio Legislature, 2019). It also provides around $50 million a year through 2030 to keep two Ohio Valley Electric Corp. (OVEC) coal plants (Kyger Creek plant in Ohio and Clifty Creek plant in Indiana) operating while simultaneously scaling back Ohio's renewable-energy goal from a maximum of 12.5 percent by 2027 to 8.5 percent by 2026. Specifically, the subsidy for OVEC's coal plants is not a fixed credit; instead, it is a rider that passes a pro-rated percentage of the financial losses onto Ohio utilities and thus ratepayers (Seryak and Worley, 2020). This type of subsidy is quite specific to these two Ohio coal plants, making it difficult to model across all active coal generators.

The Brattle Group published a study in response to Secretary Perry's letter to FERC. They explore the costs associated with preventing the retirement of coal and nuclear plants across the U.S. and the potential policies that would qualify as a "fair rate of return" for such generators (Celebi, Chupka, Oh, and Sweet, 2018). First, they examine a uniform capacity payment of $50 per kW applied to all coal and nuclear generators. The annual payment was scaled to the size of each generator. With this subsidy, between 128 GW and 154 GW of electricity generating capacity would switch from an operating margin deficit to surplus, though the authors mention that these numbers do not "represent this as the estimated effect of the payments on deterring retirements" (Celebi et al., 2018). The second policy option only targets coal and nuclear generators that were experiencing operating deficits in 2017. The policy would see generators receive payments equal to the annual operating loss minus going-forward costs. This equates to an average annual subsidy payment between $43 per kW and $58 per kW (Celebi et al., 2018). A uniform capacity payment is the simplest reliability subsidy to model and implement. We chose to administer a subsidy similar to The Brattle Group's first policy.

We investigate the impact on predicted retirement time when coal generators receive $5, $20, or $50 per

MWh of electricity generation.[18] These subsidy values encompass those explored by The Brattle Group. The total subsidy amount paid to each coal generator depends on the amount of electricity generated. We convert \$5, \$20, and \$50 per MWh to \$1.47, \$5.86, and \$14.65 per MMBtu.[19] The per unit subsidy is then attached to the quantity of electricity generated by each generator.

Similar to the carbon tax policies, a subsidy changes the profit function (Equation 1) of the coal-fired generator retirement model. The subsidy, P_S, adds a total cost of $P_S q_E(t)$, making the profit function the following:

$$\pi(P_E, P_C) = \Big((P_E(t) + P_S)q_E(t) - P_C(t)q_C(t) - VC(q_E(t)) - FC\Big) \tag{8}$$

where P_S=\$5, \$20, or \$50 per MWh (\$1.47, \$5.86, or \$14.65 per MMBtu). The subsidy increases the price received for operating the coal generator, which puts downward pressure on the retirement threshold $P_{E_R}(P_C)$. In this case, there would be less electricity-coal price combinations where the firm would optimally retire the coal generator. Figure 15 contains four retirement thresholds for a random active generator. The solid line is the benchmark retirement threshold, the dotted line is the retirement threshold when the subsidy is \$5 per MWh, the dashed line is the retirement threshold when the subsidy is \$20 per MWh, and the dash-dot line is the retirement threshold when the subsidy is \$50 per MWh. At every coal price, the corresponding electricity price can be even lower than the benchmark and the firm would still find it optimal to keep the coal generator operational.

Using these retirement thresholds that include the different levels of the subsidy, we simulate electricity and coal prices out 50 years on a monthly time step for a total of 600 simulated price combination months. The month that the simulated electricity-coal price combination first crosses the retirement threshold from the continuation region to the retirement region is recorded as the predicted retirement time. This simulation is completed 10,000 times, and we find the average predicted retirement time for each generator. A subsidy increases the revenue a generator receives from producing, which means electricity prices would have to drop very low or coal prices would have to become very high for a firm to find it optimal to retire a coal generator compared to the benchmark case.

As theory predicts, subsidizing coal generators for their reliability benefits extends the life of active generators. Compared to the benchmark, there are almost 22% fewer coal generators predicted to retire by the end of 2018 if a \$5 per MWh subsidy is provided to operating coal generators. The generators that retire within the first year even with a \$5 per MWh subsidy remove 67 thousand megawatts of generating

[18] For reference, \$43 per kW-year is the same as \$4.9 per MWh if we assume the generator operated 24 hours a day for 365 days a year.

[19] For simplicity, we will refer to the different subsidy levels using MWh units instead of MMBtu units.

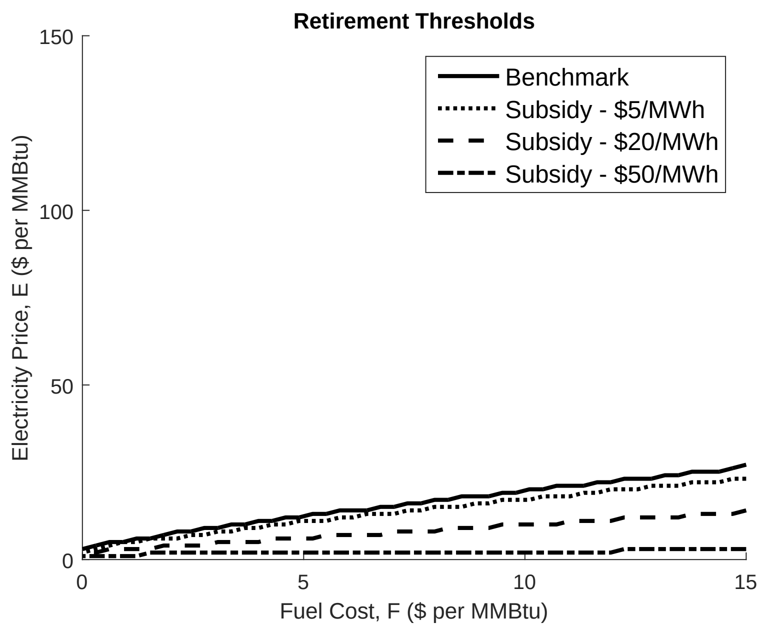

Figure 15: Benchmark Retirement Threshold vs. $5 per MWh Subsidy Retirement Threshold vs. $20 per MWh Subsidy Retirement Threshold vs. $50 per MWh Subsidy Retirement Threshold

capacity. However, in the benchmark scenario, over 89 thousand megawatts retire within the first year of the simulation, meaning the $5 per MWh subsidy prolongs the life of 22 thousand megawatts of generating capacity past the first simulated year. If the subsidy is increased to $20 per MWh or even $50 per MWh, over 61% or 90% fewer coal generators are predicted to retire within the first simulated year compared to the benchmark, keeping 64 or 85 thousand megawatts of capacity online at least an additional year. Figure 16 compares the predicted retirement time distribution under the benchmark scenario and the three reliability subsidies of differing amounts. A total of 97 (20%), 110 (22%), or 188 (38%) coal generators do not retire in the 50 year time frame given a $5, $20, or $50 per MWh subsidy. Without a subsidy, 86 (17%) coal generators remain operational. Figure 17 shows the impact of varying reliability subsidies on the change in predicted retirement time by generating capacity. Notice the subsidies affect generators relatively the same regardless of generating capacity.

A reliability subsidy aimed at extending the life of the remaining active coal generating fleet would need to pay owners more than $5 per MWh of generation. Increasing the subsidy to $20 per MWh or more drastically stretches the predicted retirement time.

This counterfactual analysis provides insight into the effects of carbon taxes as well as reliability subsidies on the predicted retirement time of the operating coal generating fleet as of December 2017. A carbon tax equivalent to the social cost of carbon hastens the retirement time while a reliability subsidy extends the life of these coal generators.

6 Conclusion

Coal-fired generators have been retiring at an increasing rate since 2009. These retirements have shifted the electricity generation portfolio across the country. In 2015, natural gas surpassed coal as the primary fuel source for electricity generation. This shift has multiple consequences for the economy (increases in unemployment in part of the country reliant on coal production) and the environment (reduction in carbon dioxide emissions produced by the electricity sector). Given the impacts of these shifts in electricity generation, economic analyses are needed to understand the drivers of these retirement decisions. This study explored how the sunk costs associated with retiring a coal generator impact the probability of retirement.

The costs associated with retiring electricity generators are scarcely reported and not publicly available. Instead of proxying for retiring costs with variables used in the past, like the amount of non-depreciated capital, we estimated the retirement costs for almost 200 coal generators that retired between 2009 and 2015 by using a real options model. We matched imputed retirement costs for retired coal generators to active generators using propensity score matching on observable generator and sociodemographic characteristics

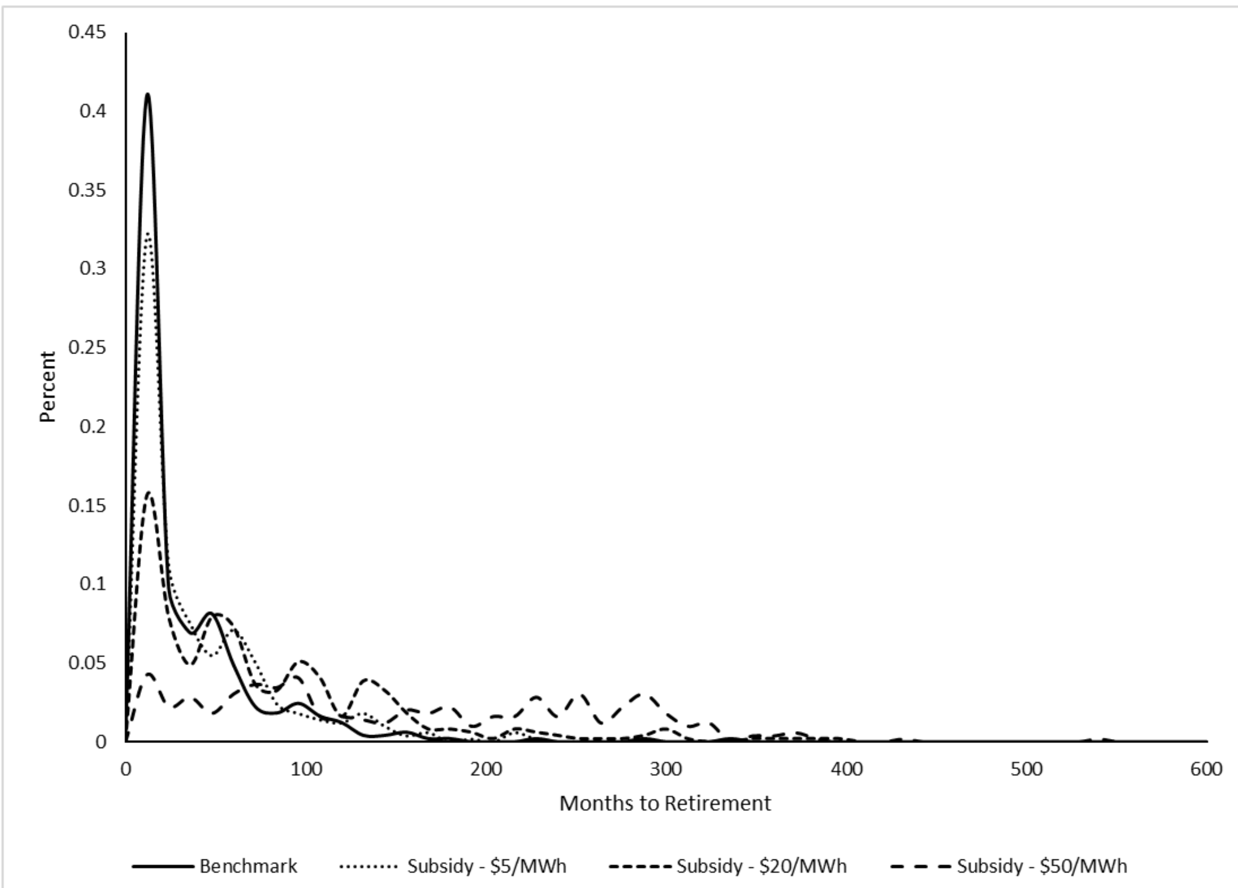

Figure 16: Predicted Retirement Time Distributions for Active Coal Generators: Benchmark vs. \$5, \$20, and \$50 per MWh Subsidy

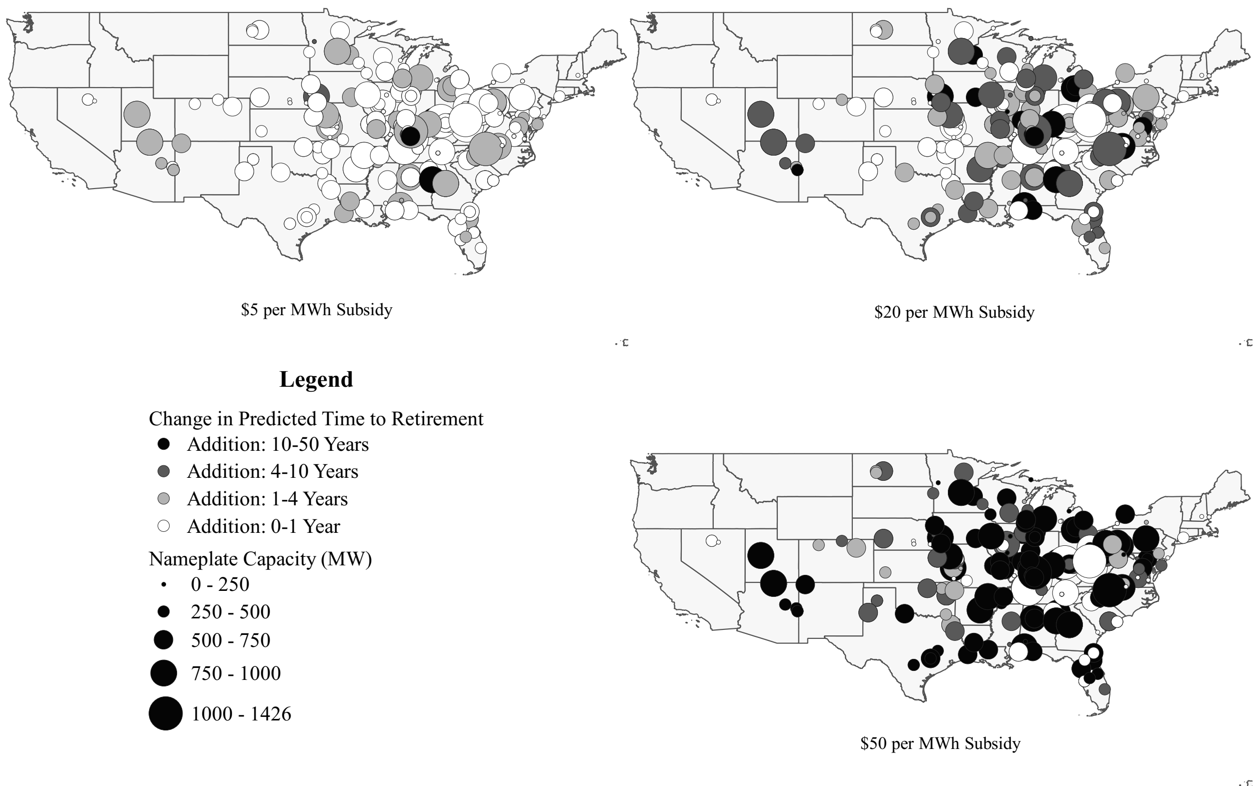

Figure 17: Change in Predicted Retirement Time by Generating Capacity for Active Coal Generators - Reliability Subsidies

and then estimated a logistic regression to determine the impact of sunk retirement costs on the probability of retirement. Our results indicate the significant impact retirement costs have in determining market exit for coal generators in the U.S. Furthermore, accounting for retirement costs improves our ability to accurately predict coal generator retirements that occurred in 2016.

To the best of our knowledge, this is the first study to back out estimates of sunk costs of market exit. The retirement decision of the average generator in our study is consistent with retirement costs greater than $37 million. Retirement costs vary across regulated and deregulated markets, across plants within the same market, and within plants. Matching retirement costs of retired generators to active generators indicates that the existing coal fleet still operating in the U.S. have retirement costs greater than $100 million.

Future applications of our real options model and subsequent empirical analysis include incorporating changes in environmental regulation. For example, our framework provides us the opportunity to add environmental regulation like a carbon tax and determine the set of existing coal generators that have the highest risk of being pushed out of the electricity market. Another application would be to create a counterfactual world in which the hydraulic fracturing boom never happened in order to determine whether the coal generators that have already retired would still be operating today. We leave these topics to future research.

References

BAUMOL, W. J. AND R. D. WILLIG (1981): "Fixed Costs, Sunk Costs, Entry Barriers, and Sustainability of Monopoly," *Quarterly Journal of Economics*, 96, 405–31.

BLACK, D., T. MCKINNISH, AND S. SANDERS (2005): "The Economic Impact of the Coal Boom and Bust," *The Economic Journal*, 115, 449–76.

BORENSTEIN, S. (2012): "Effective and Equitable Adoption of Opt-In Residential Electricity Pricing," *Review of Industrial Organization*, 42, 127–60.

CAVES, R. E. AND M. E. PORTER (1976): "Barriers to Exit," in *Essays on Industrial Organization in Honor of Joe S. Bain*, ed. by J. S. Bain, R. T. Masson, and P. D. Qualles, Ballinger Publishing Company, 36–69.

CELEBI, M., M. CHUPKA, K. OH, AND R. SWEET (2018): "The Cost of Preventing Baseload Retirements," Tech. rep., The Brattle Group.

CESENA, E. M., J. MUTALE, AND F. RIVAS-DAVALOS (2013): "Real Options Theory Applied to Electricity Generation Projects: A Review," *Renewable and Sustainable Energy Reviews*, 19, 573–81.

CICALA, S. (2015): "When Does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation?" *American Economic Review*, 105, 411–44.

CONRAD, J. M. (1997): "Global Warming: When to Bite the Bullet," *Land Economics*, 73, 164–73.

CULLEN, J. (2013): "Measuring the Environmental Benefits of Wind-Generated Electricity," *American Economic Journal: Economic Policy*, 5, 107–33.

CULLEN, J. A. AND E. T. MANSUR (2017a): "Inferring carbon abatement costs in electricity markets: A revealed preference approach using the shale revolution," *American Economic Journal: Economic Policy*, 9, 106–33.

——— (2017b): "Will Carbon Prices Reduce Emissions in the U.S. Electricity Industry? Evidence from the Shale Gas Experience," American Economic Journal: Economic Policy, forthcoming.

DAVIS, G. A. AND B. OWENS (2003): "Optimizing the Level of Renewable Electric R&D Expenditures using Real Options Analysis," *Energy Policy*, 31, 1589–1608.

DEPARTMENT OF ENERGY (2017): "Sectretary Perry Urges FERC to Take Swift Action to Address Threats to Grid Resiliency," `https://www.energy.gov/articles/`

secretary-perry-urges-ferc-take-swift-action-address-threats-grid-resiliency, [Online; accessed 5-November-2017].

Dixit, A. (1989): "Entry and Exit Decisions under Uncertainty," *Journal of Political Economy*, 97, 620–38.

Dixit, A. K. and R. S. Pindyck (1994): *Investment under Uncertainty*, Princeton, NJ: Princeton University Press.

Duke SciPol (2018): "FERC Response to DOE's Proposed Grid Resiliency Pricing Rule," https://scipol.duke.edu/content/ferc-response-does-proposed-grid-resiliency-pricing-rule, [Online; accessed 17-July-2020].

Dunne, T. and M. J. Roberts (1991): "Variation in Producer Turnover Across U.S. Manufacturing Industries," in *Entry and Market Contestability*, ed. by P. A. Geroski and J. Schwalbach, Oxford: Basil Blackwell, 187–203.

Fell, H. and D. T. Kaffine (2018): "The fall of coal: Joint impacts of fuel prices and renewables on generation and emissions," *American Economic Journal: Economic Policy*, 10, 90–116.

Fleischman, L., R. Cleetus, J. Deyette, S. Clemmer, and S. Frenkel (2013): "Ripe for Retirement: An Economic Analysis of the U.S. Coal Fleet," *The Electricity Journal*, 26, 51–63.

Forsyth, M. (2000): "On Estimating the Option Value of Preserving a Wilderness Area," *Canadian Journal of Economics/ Revue canadienne d'Economique*, 33, 413–434.

Gearino, D. (2013): "AEP to Close Coal-Fired Power Plant Instead of Converting," http://www.dispatch.com/content/stories/business/2013/07/11/aep-to-close-coal-fired-power-plant-near-beverly.html, [Online; accessed 27-February-2017].

Henson, L. (2004): "Decommissioning Handbook for Coal-Fired Power Plants," Tech. rep., Electric Power Research Institute.

Hepbasli, A. (2008): "A Key Review on Exergetic Analysis and Assessment of Renewable Energy Resources for a Sustainable Future," *Renewable and Sustainable Energy Reviews*, 12, 593–661.

Hoag, J. H. and M. Wheeler (1996): "Oil Price Shocks and Employment: The Case of Ohio Coal Mining," *Energy Economics*, 18, 211–220.

Holladay, J. S. and J. LaRiviere (2017): "The impact of cheap natural gas on marginal emissions from electricity generation and implications for energy policy," *Journal of Environmental Economics and Management*, 85, 205–227.

——— (2018): "How Does Welfare from Load Shifting Electricity Policy Vary with Market Prices? Evidence from Bulk Storage and Electricity Generation," *The Energy Journal*, 39.

Holladay, J. S. and S. Soloway (2016): "The Environmental Impacts of Fuel Switching Power Plants," *The Energy Journal*, 37, 187–213.

Insley, M. (2002): "A Real Options Approach to the Valuation of a Forestry Investment," *Journal of Environmental Economics and Management*, 44, 471–92.

Jenkin, T., D. Feldman, A. Kwan, and B. J. Walker (2019): "Estimating the Impact of Residual Value for Electricity Generation Plants on Capital Recovery, Levelized Cost of Energy, and Cost to Consumers," Tech. rep., U.S. Department of Energy.

Johnsen, R., J. LaRiviere, and H. Wolff (2019): "Fracking, coal, and air quality," *Journal of the Association of Environmental and Resource Economists*, 6, 1001–1037.

Joskow, P. L. (2007): "Competitive Electricity Markets and Investment in New Generating Capacity," in *The New Energy Paradigm*, ed. by D. Helm, Oxford University Press.

Kaffine, D. T., B. J. McBee, and J. Lieskovsky (2013): "Emissions Savings from Wind Power Generation in Texas," *The Energy Journal*, 34, 155–75.

Karlin, S. and H. M. Taylor (1981): *A Second Course in Stochastic Processes*, New York: Academic Press.

Knittel, C. R., K. Metaxoglou, and A. Trindale (2017): "Are we Fracked? The Impact of Falling Gas Prices and the Implications for Coal-to-Gas Switching and Carbon Emissions," Oxford Review of Economic Policy, forthcoming.

Kumbaroglu, G., R. Madlener, and M. Demirel (2008): "A Real Options Evaluation Model for the Diffusion Prospects of New Renewable Power Generation Technologies," *Energy Economics*, 30, 1882–1908.

Linn, J. and K. McCormack (2019): "The roles of energy markets and environmental regulation in reducing coal-fired plant profits and electricity sector emissions," *The RAND Journal of Economics*, 50, 733–767.

Linn, J. and L. Muehlenbachs (2018): "The heterogeneous impacts of low natural gas prices on consumers and the environment," *Journal of Environmental Economics and Management*, 89, 1–28.

Logan, J., C. Marcy, J. McCall, F. Flores-Espino, A. Bloom, J. Aabakken, W. Cole, T. Jenkin, G. Porro, C. Liu, F. Ganda, R. Boardman, T. Tarka, J. Brewer, and T. Schultz (2017): "Electricity Generation Baseline Report," Tech. rep., U.S. Department of Energy.

Lu, X., J. Salovaara, and M. B. McElroy (2012): "Implications of the Recent Reductions in Natural Gas Prices for Emissions of CO2 from the U.S. Power Sector," *Environmental Science and Technology*, 46, 3014–21.

Marten, A. L. and C. C. Moore (2011): "An Options Based Bioeconomic Model for Biological and Chemical Control of Invasive Species," *Ecological Economics*, 70, 2050–2061.

Miranda, M. J. and P. L. Fackler (2002): *Applied Computational Economics and Finance*, Cambridge, MA: MIT Press.

O'Brien, J. and T. Folta (2009): "Sunk Costs, Uncertainty and Market Exit: A Real Options Perspective," *Industrial and Corporate Change*, 18, 807–33.

O'Brien, J. N., J. Blau, and M. Rose (2004): "An Analysis of the Institutional Challenges to Commercialization and Deployment of IGCC Technology in the U.S. Electric Industry," Tech. rep., The National Association of Regulatory Utility Commissioners.

Pachamanova, D. A. and F. J. Fabozzi (2011): "Modeling Asset Price Dynamics," in *The Theory and Practice of Investment Management: Asset Allocation, Valuation, Portfolio Construction, and Strategies*, ed. by F. J. Fabozzi and H. M. Markowitz, John Wiley & Sons, Inc., chap. 6, 125–58, 2 ed.

Pettersson, F., P. Soderholm, and R. Lundmark (2012): "Fuel Switching and Climate and Energy Policies in the European Power Generation Sector: A Generalized Leontief Model," *Energy Economics*, 34, 1064–73.

Pindyck, R. S. (1991): "Irreversibility, Uncertainty, and Investment," *Journal of Economic Literature*, 29, 1110–48.

——— (1999): "The Long-Run Evolution of Energy Prices," *The Energy Journal*, 20, 1–27.

Pindyck, R. S. and D. L. Rubinfeld (1998): *Econometric Models and Economic Forecasts*, New York, NY: Irwin/McGraw-Hill.

Qian, C. (2019): "Three Carbon Tax Bills Introduced in Congress," `https://taxfoundation.org/carbon-tax-bills-introduced-congress/`, [Online; accessed 14-July-2020].

Rosenbaum, D. I. and F. Lamort (1992): "Entry, Barriers, Exit, and Sunk Costs: An Analysis," *Applied Economics*, 24, 297–304.

Rothwell, G. (2006): "A Real Options Approach to Evaluating New Nuclear Power Plants," *The Energy Journal*, 27, 37–53.

Schiavo, J. G. and R. Mendelsohn (2019): "The Effect Of Domestic Air Pollution Mitigation And Fracking On Retirements Of Coal Power Plants," *Climate Change Economics*, 10, 1950008.

Seryak, J. and P. Worley (2020): "FERC's December 2019 Order on State Subsidies," Tech. rep., The Ohio Manufacturers' Association.

Siegfried, J. J. and L. B. Evans (1992): "Entry and Exit in United States Manufacturing Industries from 1977 to 1982," in *Empirical Studies in Industrial Organization: Essays in Honor of Leonard W. Weiss*, ed. by D. W. Audretsch and J. J. Siegfried, Dordrecht: Kluwer Academic Publishers.

——— (1994): "Empirical Studies of Entry and Exit: A Survey of the Evidence," *Review of Industrial Organization*, 9, 121–55.

The Ohio Legislature (2019): "House Bill 6," `https://www.legislature.ohio.gov/legislation/legislation-summary?id=GA133-HB-6`, [Online; accessed 20-July-2020].

Tsekrekos, A. E. (2010): "The Effect of Mean Reversion On Entry and Exit decisions under Uncertainty," *Journal of Economic Dynamics & Control*, 34, 725–42.

U.S. Energy Information Administration (2014): "Levelized Cost and Levelized Avoided Cost of New Generation Resources in the Annual Energy Outlook 2014," `https://www.eia.gov/outlooks/archive/aeo14/pdf/electricity_generation_2014.pdf`, [Online; accessed 07-September-2016].

——— (2016a): "Carbon Dioxide Emissions Coefficients," `https://www.eia.gov/environment/emissions/co2_vol_mass.php`, [Online; accessed 27-January-2020].

——— (2016b): "Coal made up more than 80% of Retired Electricity Generating Capacity in 2015," `https://www.eia.gov/todayinenergy/detail.php?id=25272`, [Online; accessed 14-March-2017].

——— (2016c): "Electric Power Annual 2015," `https://www.eia.gov/electricity/annual/pdf/epa.pdf`, [Online; accessed 07-September-2016].

——— (2016d): "Total Electricity Sales Fell in 2015 for 5th Time in Past 8 Years," `https://www.eia.gov/todayinenergy/detail.php?id=25352`, [Online; accessed 30-April-2017].

——— (2017a): "Net Generation by Energy Source: Total - All Sectors," https://www.eia.gov/electricity/data/browser/, [Online; accessed 30-April-2017].

——— (2017b): "U.S. Natural Gas Electric Power Price," https://www.eia.gov/dnav/ng/hist/n3045us3A.htm, [Online; accessed 20-April-2017].

U.S. Environmental Protection Agency (2016a): "Coal Plant Decommissioning Financing Cleanup and Redevelopment," https://nepis.epa.gov/Exe/ZyNET.exe/P100OWFH.TXT?ZyActionD=ZyDocument&Client=EPA&Index=2011+Thru+2015&Docs=&Query=&Time=&EndTime=&SearchMethod=1&TocRestrict=n&Toc=&TocEntry=&QField=&QFieldYear=&QFieldMonth=&QFieldDay=&IntQFieldOp=0&ExtQFieldOp=0&XmlQuery=&File=D%3A%5Czyfiles%5CIndex%20Data%5C11thru15%5CTxt%5C00000020%5CP100OWFH.txt&User=ANONYMOUS&Password=anonymous&SortMethod=h%7C-&MaximumDocuments=1&FuzzyDegree=0&ImageQuality=r75g8/r75g8/x150y150g16/i425&Display=hpfr&DefSeekPage=x&SearchBack=ZyActionL&Back=ZyActionS&BackDesc=Results%20page&MaximumPages=1&ZyEntry=1&SeekPage=x&ZyPURL, [Online; accessed 26-April-2016].

——— (2016b): "EPA Fact Sheet Social Cost of Carbon," https://www.epa.gov/sites/production/files/2016-12/documents/social_cost_of_carbon_fact_sheet.pdf, [Online; accessed 27-January-2020].

——— (2020): "Report on the Environment," https://www.epa.gov/report-environment/greenhouse-gases, [Online; accessed 20-July-2020].

Yang, M., W. Blyth, R. Bradley, D. Bunn, C. Clarke, and T. Wilson (2008): "Evaluating the Power Investment Options with Uncertainty in Climate Policy," *Energy Economics*, 30, 1933–50.

Ye, J. (2020): "U.S. State Carbon Pricing Policies," https://www.c2es.org/document/us-state-carbon-pricing-policies/#:~:text=Those%20states%20are%20California%20and,Greenhouse%20Gas%20Initiative%20(RGGI)., [Online; accessed 13-July-2020].

Zivin, J. S. G., M. J. Kotchen, and E. T. Mansur (2014): "Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies," *Journal of Economic Behavior & Organization*, 107, 248–268.

Table A.1 : Unit Root Test for Delivered Coal Prices at a Random Retired Plant

Unrestricted regression		
Coefficient	Estimate	Std. Error
β_0	0.588	0.209
β_1	0.065	0.090
β_2	-0.001	0.000
β_3	-0.185	0.047
Restricted regression		
β_0	0.004	0.010
β_1	0.006	0.093
$N = 117$		$F = 8.219$
Prob > F = 0.000		

A Appendix

A.1 Unit Root Tests

Given that real options results critically depend on choosing the correct stochastic process, we tested data for consistency with Brownian motion instead of the assumption of mean reversion. This is typically completed using an augmented Dickey Fuller test (Conrad, 1997; Forsyth, 2000; Insley, 2002; Pindyck and Rubinfeld, 1998). Geometric Brownian motion (GBM) assumes P is log-normally distributed.[1] The logged price level $p = \ln(P)$ is normally distributed and follows an arithmetic Brownian motion (ABM) $\mathrm{d}p = \mu \mathrm{d}t + s\mathrm{d}z$. If p is consistent with ABM, Ito's Lemma ensures P must be consistent with GBM. To test that p are consistent with ABM, we ran a restricted regression

$$(p_t - p_{t-1}) = \beta_0 + \beta_1(p_{t-1} - p_{t-2}) + \epsilon_t \tag{A.1}$$

and an unrestricted regression

$$(p_t - p_{t-1}) = \beta_0 + \beta_1(p_{t-1} - p_{t-2}) + \beta_2 t + \beta_3 p_{t-1} + \epsilon_t \tag{A.2}$$

The null hypothesis that corresponds with p being ABM is $\mathrm{H}_0 : \beta_2 = \beta_3 = 0$. This null hypothesis is rejected at the 1% or 5% level for all coal generators in the analysis, retired or active. This is true for coal prices and electricity prices. Tables A.1 and A.2 provide an example of the augmented Dickey Fuller tests for a random retired coal generator.

[1] P is for either P_E or P_C

Table A.2 : Unit Root Test for Wholesale Electricity Prices at a Random Retired Generator

Unrestricted regression		
Coefficient	Estimate	Std. Error
β_0	1.512	0.423
β_1	-0.157	0.084
β_2	-0.001	0.000
β_3	-0.264	0.067
Restricted regression		
β_0	-0.006	0.012
β_1	-0.290	0.081
$N = 142$		$F = 7.868$
Prob > F = 0.001		

A.2 Geometric Mean Reversion Parameter Estimation

After determining that Brownian motion is inappropriate for modeling both coal prices and electricity prices, we turned to geometric mean reversion (GMR) $\mathrm{d}P = r_P(\bar{P} - P)P\mathrm{d}t + \sigma_P P \mathrm{d}z_P$. The geometric mean reversion model can be written as the following:

$$P_{t+1} = P_t + r_P(\bar{P} - P_t)P_t + \sigma_P P_t \epsilon_t \tag{A.3}$$

Here, ϵ_t is a standard normal random variable. To estimate the parameters r_P, $\bar{P}$, and σ_P, we used the data described in Section 2.3.1. We assumed that the parameters of the geometric mean reversion remain constant during the time period of estimation to rewrite the equation for GMR as

$$\frac{P_{t+1} - P_t}{P_t} = r_P \bar{P} - r_P P_t + \sigma_P \epsilon_t \tag{A.4}$$

This equation bears characteristics of a linear regression model, with the percentage price change $\frac{P_{t+1} - P_t}{P_t}$ as the dependent variable and P_t as the explanatory variable.

According to Pachamanova and Fabozzi (2011), the estimate of r_P is obtained as the negative of the coefficient in front of P_t. One way to check whether GMR is a consistent assumption for prices is to determine whether the coefficient in front of P_t is positive since the rate of reversion r_P cannot be a negative number. The estimate for the long-run mean of prices $\bar{P}$ is obtained as the ratio of the intercept term estimated from the regression and the negative of the slope coefficient in front of P_t. The last estimate for volatility σ_P is obtained as the standard error of the regression.

Two other methods for determining if GMR is suitable are the following: (a) the p-value for the coefficient in front of P_t should be small, preferably less than 0.05, and (b) the points in a scatter plot of P_t versus $\frac{P_{t+1} - P_t}{P_t}$ should vary around a straight line with no visible cyclical or other patterns. Tables A.3 and A.4

Table A.3 : Geometric Mean Reversion Estimates for Delivered Coal Prices at a Random Retired Plant

Coefficient	Estimate	Std. Error
$r_P\bar{P}$	0.172	0.052
r_P	-0.043	0.013
$N = 118$		$r_P = 0.043$
$\bar{P} = \$4.03$		$\sigma_P = 0.118$

Table A.4 : Geometric Mean Reversion Estimates for Wholesale Electricity Prices at a Random Retired Generator

Coefficient	Estimate	Std. Error
$r_P\bar{P}$	0.123	0.045
r_P	-0.011	0.004
$N = 143$		$r_P = 0.011$
$\bar{P} = \$11.60$		$\sigma_P = 0.163$

display the results of the regressions that determine the parameter estimates for coal and electricity prices for a random coal generator in our analysis. Coal and electricity prices satisfied all three checks described by Pachamanova and Fabozzi (2011) for all generators. Our results were simulated under the assumption that coal and electricity prices follow GMR.