

Coal-Fired Power Plant Retirements in the U.S.

Rebecca J. Davis J. Scott Holladay Charles Sims*

May 2020

Abstract

We summarize the history of U.S. coal-fired generator retirements over the last decade, describe planned future retirements, and forecast the remaining operating life for every operating coal-fired generator. Nearly one-third of the coal fleet retired during the 2010s and a quarter of the remaining capacity has announced plans to retire. We summarize the generator technology and location trends that are correlated with the observed retirements. We then describe a theoretical model of the retirement decision coal generator owners face. We use retirements from the last decade to quantify the relationships in the model for retired plants. We use a machine learning algorithm to predict the decisions of active generators based on the relationship observed in retired units. We simulate future coal and electricity prices and identify which generators retire at what time in a business-as-usual policy scenario, under a coal subsidy, and under a carbon tax. Our model predicts that three-quarters of coal generation capacity will retire in the next twenty years, with most of that retirement concentrated in the next five years. Policy has limited ability to affect retirement times. A \$20 per MWh electricity subsidy extends the average life of a generator by six years. A \$51 per ton carbon tax brings forward retirement dates by about two years. In all scenarios, a handful of electricity generators remain on the grid beyond our twenty-year forecast horizon.

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

JEL Codes: Q4, L1, L5, H4

1 Introduction

In 2010, coal-fired generation accounted for over half of electricity produced in the United States. A decade later, that had fallen to around a quarter of total generation. This shift has been driven primarily by the retirement of existing coal-fired generators and has already had wide reaching effects. Regions of the country that produce coal are likely to struggle economically as consumption falls. Coal generators are 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 to perpetuate this shift away from coal generators (Cullen and Mansur, 2017). Subsidies to keep coal plants on the grid have also been proposed for reliability

*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. Thanks to Jim Stock and Tatyana Deryugina for helpful feedback. Kate Martella provided outstanding research assistance. This research 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. Holladay and Sims also gratefully acknowledge funding from the Alfred P. Sloan Foundation grant no. G-2019-11399 and the Tennessee Valley Authority.

reasons. When and where such policies will alter the trajectory of coal generator retirements remains an unanswered question.

Instead of a postmortem analysis of the drivers of retirement, we consider the effect of policies that may exacerbate or alleviate future coal-fired power plant retirements in the U.S. We predict the expected retirement time for every coal-fired generator in the country and then evaluate how environmental regulation and efforts to enhance grid reliability would affect those retirement dates. This generator-level analysis gives us a unique perspective on the relative influence of market forces and policy on the composition of the generating fleet. In our no-policy baseline, we find that nearly three-quarters of coal-fired generation capacity retires by 2040, the end of our simulation. Retirements are concentrated in the upper Midwest, the Ohio Valley, and the southeastern U.S. A \$20 per MWh of electricity generated production credit, more than half the cost of delivered fuel on average, extends the median retirement date by six years. A carbon tax set to recent estimates of the social cost of carbon, \$51 per ton, pulls the median retirement forward by about two years.¹

A significant fraction of the remaining coal generating capacity is forecast to retire in the next five years in the no policy baseline. Because so much capacity is forecast to retire soon, the ability of carbon taxes to speed-up retirements is limited in much of the U.S. There is more scope for reliability subsidies to extend the operations of generators than for carbon taxes to drive retirements. However, the magnitude of the subsidies required to maintain a large coal fleet are enormous. There are likely less expensive ways to ensure reliability.

In this paper we do not explore entry of new coal plants. Given the electricity and coal prices we observe over our sample period, our model finds that entry of new coal plants is not economic. This is borne out by the data. The EIA's 2021 Energy Outlook estimates the levelized cost of energy (LCE), the all in cost of generating from a particular fuel type, as \$73 per MW of coal-fired capacity. That is more than double the LCE of solar and wind capacity and nearly double the LCE of combined-cycle natural gas. As of 2019, EIA reports 135,000 MWs of proposed capacity additions across all fuel types. Of those additions, only 17 MWs are coal-fired.² Building new coal plants would require new technological developments or energy policy changes that are hard to foresee at the present time.

We start by briefly summarizing the history of recent retirements of coal-fired generators. Using publicly available data on electric generators and power plants, we show that retirements have increased and that those retirements are coming from smaller, older generators. We also identify differences in retirements

¹Our model does not produce standard errors or other traditional measures of uncertainty. Simulation of coal and electricity prices, as well as forecasting retirement costs, could produce an analog of sampling variation. In Davis, Holladay, and Sims (2021) we examine those individual sources of variation.

²The University of Alaska Fairbanks has proposed replacing its existing coal-fired combined heat and power system with a new coal-fired system. <https://www.uaf.edu/heatandpower/background.php>

across regions of the U.S.

Next, we describe a new data set collecting all information on scheduled and planned retirements for every active coal-generator on the grid. We collect scheduled retirement dates from the U.S. Energy Information Administration (EIA) and supplement that with media and financial database searches to identify announced retirements that have not yet been formally published in official sources.

Finally, we describe a novel 3-step technique to back out the unobserved retirement costs implied by retired generators and map those costs onto active generators first proposed in Davis et al. (2021). 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 (2012), Shapiro and Walker (2018), Suzuki (2013)). 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.

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 every coal-fired generator in the United States that was active in 2009. Looking at actual retirement dates for generators that retired between 2009 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, 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.

We conduct two counterfactual analyses, modelling a carbon tax of \$51 per ton of carbon dioxide and a fuel cost subsidy of \$20 per MWh of generation. 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.

Our ability to generate counterfactual retirement dates across different electricity and fuel price levels allows us to model a variety of policy scenarios of interest to policy makers and the energy industry. We can identify specific retirement times for each existing coal plant in the country. The procedure could allow us to identify reliability issues, local environmental or economic impacts of plant retirements as well.

2 Coal-Fired Generator Retirements

Power plants typically include multiple separate generators each of which can be fueled separately. The average coal plant active on the grid in 2010 had more than five generators of which 3 were coal-fired and the others, typically much smaller in capacity, were oil or gas fired. The smaller units are mostly used to help start the coal-fired generators and manage small fluctuations in production or demand. Because each generator at a power plant can use different fuels and can be retired separately, we use electric generators as the unit of analysis in this study. Of the retirements we observe over the last decade, around sixty percent retire all coal units at a plant at the same time.

Coal generating capacity has fallen steadily since its peak in 2011. Since that time environmental regulations have increased, renewable generation capacity has exploded and natural gas prices have fallen, driving peak electricity prices lower. Given these headwinds it is not particularly surprising that coal generating capacity has shrunk. Figure 1 displays the location of the coal-fired generators that have retired between 2010-2019. They represent 473 generators with a nameplate capacity of nearly 80 thousand MWs. At the end of 2019, 632 coal-fired generators with a cumulative nameplate capacity of 244 thousand MWs remained active on the grid.³

Retiring generators were smaller, older than those that remained operating at the end of the decade. Generators that retired between 2010 and 2019 had a median capacity of 115 MWs compared to 350 MWs for surviving generators. The capacity-weighted average operating date⁴, for retired coal-fired generators was

³An additional 32 generators with 3,200 MWs in capacity were on standby, a likely precursor to retirement. Those generators are not included in this analysis.

⁴The industry refers to the date that a unit first produced electricity for the grid as its operating date. Each generator at a

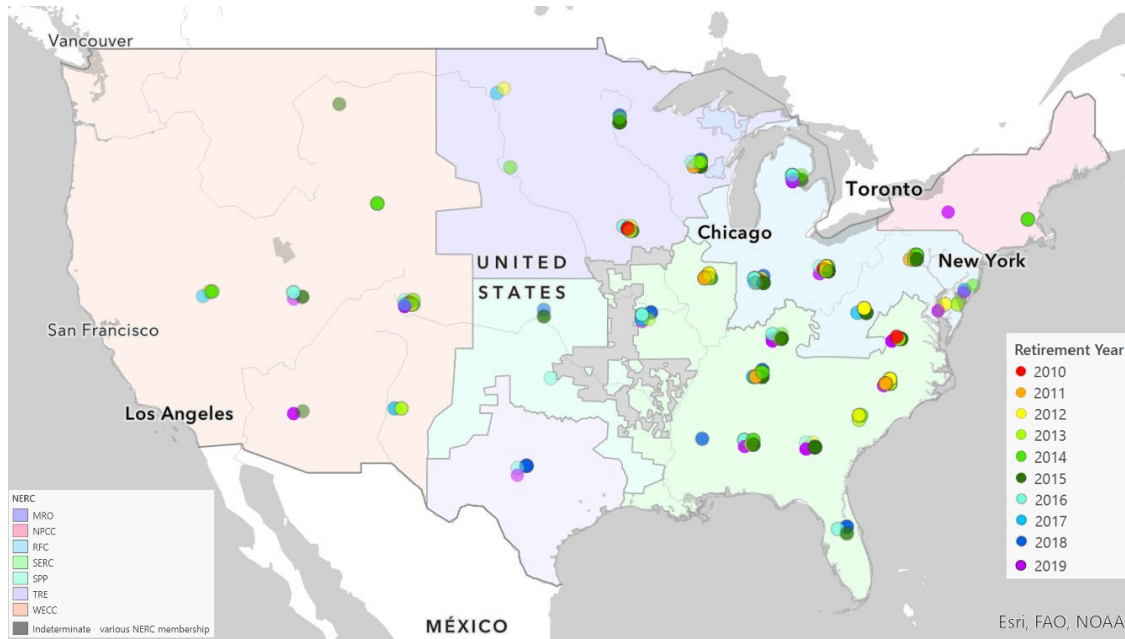


Figure 1: Coal-Fired Generators Retirements 2010-2020

Note: The figure shows the location of the 473 coal-fired generators that have retired between 2010-2019. The color of the circle represents year the generator retired. A small amount of jitter is added so that multiple generators at the same plant are all visible. Source: Authors' mapping of EIA Form 860 Data for 2019.

1965 while active generators have a capacity weighted average operating date of 1978. Retired generators were much more likely to use diesel oil, rather than natural gas, as their start up fuel source and less likely to have super-critical technology which is more expensive to install, but allows for more efficient generation.⁵

The North American Reliability Corporation (NERC) breaks the country into six reliability regions. These regions represent well-connected portions of the grid where electricity can flow from generators to load with a minimum of congested choke points. The NERC's goal is to ensure generation capacity can meet demand within and across these regions.⁶ In absolute terms, most of the retirements were in RFC (Midwest) and SERC (southeast) with around 30 thousand MWs of coal-fired generating capacity retiring over the last decade. Relative to coal capacity, nearly one-third of coal in NPCC (the New England states) retired. In MRO (upper mid-west down to Oklahoma) only around ten percent of coal-fired generating capacity retired.

plant can have a different operating date.

⁵Less than ten percent of retired generators reported super-critical generation ability, but nearly 90 percent of active generators have that technology. Super-critical boilers heat steam to such high temperatures that it exhibits properties of both liquid and gas phases allowing it to spin generators more efficiently.

⁶The six regions are the Northeastern states (NPCC), the Southeastern states (SERC), the Midwest (RFC), the middle of the country from the upper Midwest down to Oklahoma (MRO), Texas (TRE) and the entire western part of the country from New Mexico to Washington state (WECC). See Figure A1 for a map of NERC regions.

3 Announced Retirements

The wave of retirements we have witnessed over the last decade is likely just the beginning. Many operating coal generators have announced their intention to retire. In this section, we use EIA data on announced retirements to describe the coming wave of retirements. We supplement that data with hand-collected media reports to identify coal-fired generators that have publicly announced planned retirement dates, but not yet reported that information to the EIA.⁷ To our knowledge, this data set represents the most complete list of planned coal generator retirements at this time.

We performed a series of Google and Lexis-Nexus searches for the names of power plants with active coal-fired generators. We found press releases, media stories and official filings that mentioned these plants and looked for any mention of retirement plans. In some cases we found specific retirement dates. More common were announcements of plans to retire a generator before a specific date, (i.e. “by 2025” or “later this decade”).⁸ The full list of active coal generators with announced retirement dates and links to the primary sources we used will be available on our websites.

Table 1: Planned Coal Generator Retirements

Year	MW Scheduled to Retire	Generator Count
2020	9,709	35
2021	3,626	14
2022	8,437	23
2023	5,985	17
2024	2,697	7
2025	2,847	8
2026	621	1
2027	1,895	3
2028	2,286	3
2029	905	2
2030	3,875	6
2030+	9,046	16
Total	51,929	135

Note: Planned retirements reported in EIA’s Form 860 data in 2019. The form asks generators to report any retirement planned in the next five years. Some generators report retirements beyond that date by choice. The total row reports the sum of all retirements reported in the EIA 860 across each year. The EIA 860 data reports 463 active generators with a capacity of 192 thousand megawatts in 2019.

Table 1 reports the planned coal-fired generator retirements by year. Planned retirements are front-loaded with much of the planned retirements occurring in the next three years. The EIA 860 questionnaire asks respondents to report the date of planned retirements for each generator.⁹ More than ten percent of

⁷The EIA requires that any generator with a planned retirement date in the next five years to report this information. In practice, a number of generators with later retirement dates report that information as well.

⁸Because these unofficial retirement announcements are non-binding we do not use them in the summary statistics that follow.

⁹Specifically, the latest version of the EIA 860 survey, on line 8 says, “If this generator will be retired in the next ten years, what is its estimated retirement date?” Previous version of the survey only asked for retirement dates in the next five years.

Table 2: Planned Coal-Generator Retirements by NERC Region

NERC Region	Num Gen	Total Cap (MWs)	Median Planned Retirement Year
MRO	21	7,438	2022
NPCC	2	1,055	2020.5
RFC	49	16,376	2022
SERC	28	9,008	2021.5
TRE	9	6,239	2030
WECC	28	12,747	2025

Note: Planned retirements reported in EIA’s Form 860 data in 2019. The form asks generators to report any retirement planned in the next five years. Some generators report retirements beyond that date by choice. The EIA 860 data reports 463 active generators with a capacity of 192 thousand megawatts in 2019.

coal generating capacity plans to retire in the next three years. Including all announced retirements with a specific date, there are 50 thousand MWs of capacity planning to retire, more than a quarter of the total. Of course, additional capacity may retire before 2030. Over the last decade the median time between announcement and actual retirement has been three years. Generators that may retire in the second-half of the 2020’s have not yet made that decision. For that reason, the 25 thousand MWs of planned retirements beyond 2024 should be considered a lower bound on the actual capacity retired.

Table 2 reports planned retirements by NERC region. Most of the announced retirements are in RFC (Midwest) and WECC (Western United States). Conditional on reporting a retirement date, the planned retirements in RFC are much sooner with a median retirement date of 2022. The median retirement date for the 28 generators in WECC with an announced retirement date is 2025. NPCC (New England) has only twelve active generators and two of those are scheduled to retire in the near future.

After this wave of planned retirements, coal generation will become largely a regional phenomena. RFC (mid-Atlantic states) and SERC (southeast) will have around three quarters of coal generating capacity. MRO (upper Midwest) will have less than twenty percent and the northeast (NPCC), Texas (TRE) and western U.S. (WECC) will have less than five percent total.

Figure 2 displays the coal-fired generators that have not yet reported a retirement date. These coal generators are expected to be online for at least five years and potentially much longer. These generators are concentrated in the eastern half of the United States, with Florida and Georgia in the Southeast retaining a cluster of generators. The generators are (mostly) exposed to higher electricity prices helping them maintain a profit as other coal-fired generators are squeezed out. Another swath of large generators runs from Appalachian coal country in West Virginia and Kentucky through Ohio over to Indiana. These generators have access to relatively inexpensive coal which keeps them competitive even in the face of lower electricity

Respondents can enter a month, but are told to enter June if they are uncertain which month. We choose not to use the month portion of the response because most generators report planned retirements of June or December while actual retirements typically occur during the lower demand spring or fall.

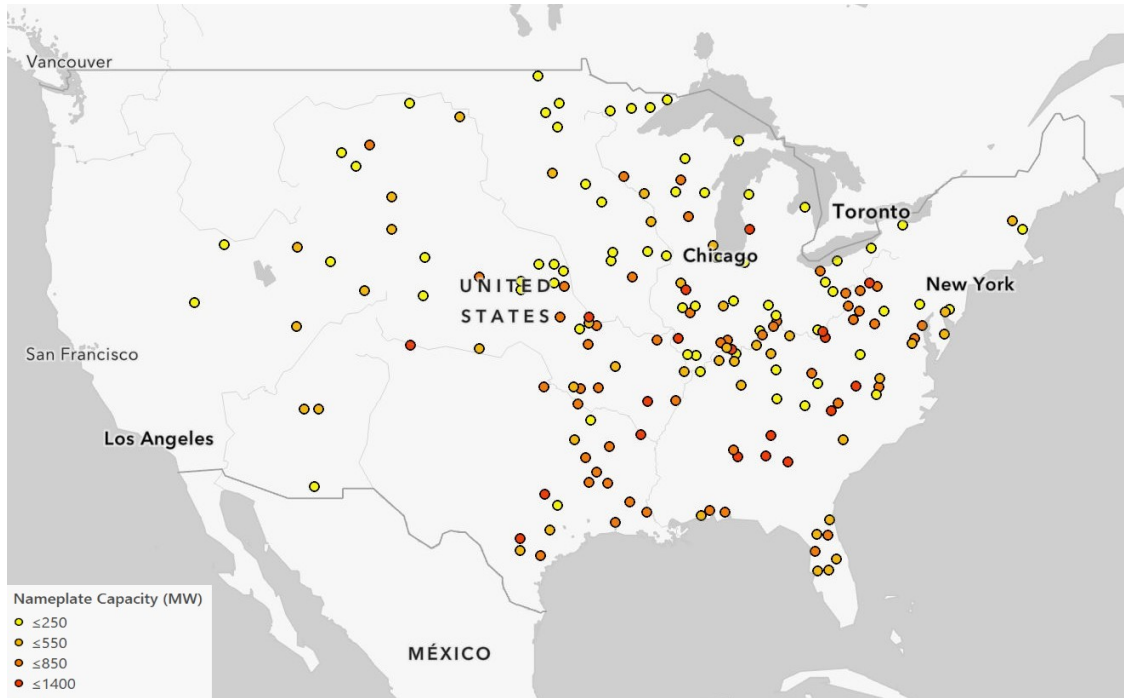


Figure 2: Active Coal-Fired Generators with No Announced Retirement Date

Note: The figure shows the location of the approximately 300 coal-fired generators that are currently operating that have not announced a retirement date. The color indicates the nameplate capacity of the generator. Source: Authors’ mapping of EIA Form 860 Data for 2019.

prices that cause generators elsewhere to choose to exit.

4 Model of Coal Generator Retirement

The details of the model are described in Davis et al. (2021). Here we briefly describe the process, but focus on discussing future coal plant retirements and the implications of the retirements we are predicting. To predict future retirements of coal-fired power plants we need a model of power plant owner’s decision-making process. We choose a real options model. We then use a machine learning algorithm to map characteristics of generators that retired over the last ten years to the conditions that caused those generators to retire. We can use that mapping to predict future retirements of the remaining active generators on the grid. This section briefly describes each of those two steps.

4.1 Real Options Model

Dixit and Pindyck (1994) define an investment as “the act of incurring an immediate cost in the expectation of future rewards.” When these future rewards are uncertain and the immediate cost is sunk, investment

rules based on traditional methods such as discounted cash flow analysis are biased. When an investment decision can be postponed and this delay will alleviate the uncertainty in future returns, there is an incentive, an option value, to delay this decision. The delay avoids the irreversibility of the sunk cost and allows one to respond to new information.

By incorporating the option value created by uncertainty and irreversibility, real options theory focuses on the optimal timing of an irreversible investment decision by accounting for the value of being able to postpone an uncertain investment. Traditionally, investment decisions ask “if” an investment should be made. In the presence of uncertainty and irreversibility and the option value they create, the investment decision becomes “when” the investment should be made.

Coal-fired power plant retirements face similar incentives. Retiring coal-fired generating units incurs sunk costs in the form of scrapping machinery, decommissioning sites, and selling suitable land. These sunk retirement costs (K) 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). There are also significant uncertainties associated with future fuel costs, electricity demand, and regulations. The combination of sunk costs of retirement and uncertain future profits from coal-fired generation create an incentive for firms to delay retiring these units.

Our approach treats the decision to retire a coal-fired generating unit as a divestment option. If you own a coal-fired generator, you can remain online and earn profits based on local electricity and coal prices or you have the option of retiring the unit. If you exercise the option of retiring the unit, you incur a sunk retirement cost but forego all future losses that the unit might incur. If you choose to keep the unit online, you defer the retirement cost and the possibility of future profits remains open.

For each retired coal-fired generator, we observe coal prices, electricity prices, and the date of retirement. We use that data to create a threshold of electricity and coal prices that would induce the utility to pay a retirement cost and retire the generator. We identify the threshold by examining each possible electricity price, coal price combination and calculate whether the plant would be profitable at those prices. The firm receives a flow payoff:

$$\pi(P_E, P_C) = \left(P_E(t)q_E(t) - P_C(t)q_C(t) - VC(q_E(t)) - FC \right) \quad (1)$$

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 for the generator, and FC is the fixed levelized capital cost of the generator.¹⁰ q_E

¹⁰Following 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.

is the quantity of electricity supplied by the generator. q_C is the quantity of coal used to generate q_E . We assume that as the price of electricity goes up, the quantity of electricity produced and coal consumed goes up as well. The 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.

At points with very low electricity prices and high coal prices the plant loses money and would choose to exit. At points with high electricity prices and low coal prices, the plant will make money and choose to operate. We search for the combinations of electricity and coal prices where the generator just breaks even and is indifferent between staying on the grid and generating or retiring. The set of pairs of electricity and coal prices where the generator breaks even is the threshold. That threshold also depends on the size of the retirement costs faced by these generators. Profitability is dependent of the sunk retirement costs. A generator that is expensive to decommission faces a high barrier to exit; thus, there would be fewer combinations of electricity and coal prices where it would be profitable to retire the generator than to keep it online. However, there are also generators that require little in the way of retirement costs. Those generators would have a lower barrier to exit, meaning there would be more combinations of electricity and coal prices where it would be profitable to retire the generator.

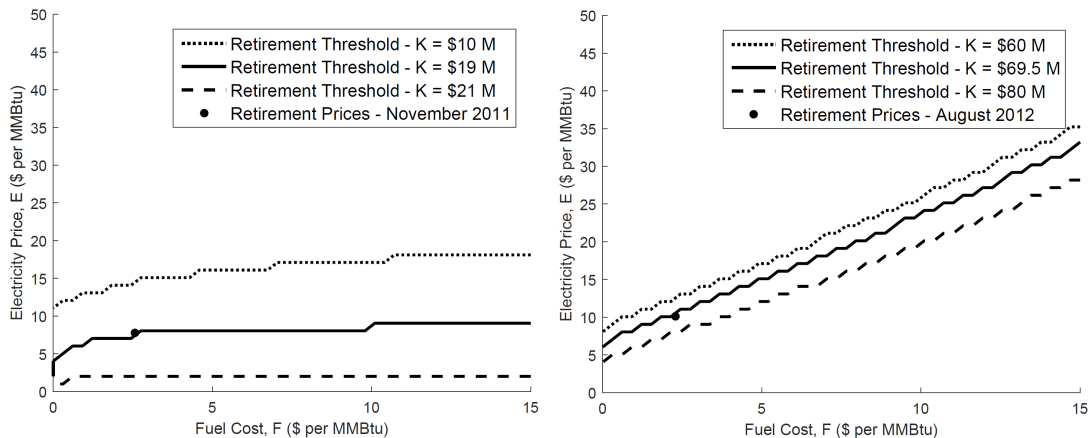


Figure 3: Estimating Implied Retirement Costs

Note: The figure shows the coal price and electricity price combinations that will result in coal generator retirements for two coal generators: ST1 at Vermilion Power Station in IL (left) and generator 7 at Crawford Generating Station in IL (right). Electricity and coal price combinations above the diagonal line(s) keep the generator online. Pairs below the diagonal(s) induce retirement. Each threshold represents a different retirement cost level. The dot represents the observed prices on the date of retirement.

To illustrate our approach, Figure 3 shows the retirement thresholds for two coal-fired generators in Illinois. On the left is ST1 at the Vermilion Power Station in Vermilion, IL with 75 MW nameplate capacity that came online in 1955. On the right is generator 7 at Crawford Generation Station in Cook County, IL.

It had a 293 MW nameplate capacity and came online in 1958. The threshold for the Vermilion generator is much flatter, suggesting it is less sensitive to fuel prices than the Crawford Generator.

The three thresholds on each graph illustrate how the level of retirement costs (K) impacts the electricity and coal price pairs for which retirement is optimal. Larger retirement costs shift the retirement threshold down indicating that retirement is desirable over a smaller set of prices. We do not observe the retirement cost paid by each retired coal-fired generator. To determine the costs associated with retiring each generator, we find the retirement cost that causes the retirement threshold to intercept the point represented by the coal and electricity prices observed at the time of retirement.¹¹ Higher retirement costs would imply the unit should have remained active after it was retired while lower retirement costs would imply the unit should have been retired far sooner than it was actually retired.

Based on that information, we can back out what retirement costs must have been at those plants - the implied retirement cost. Specifically, we can find the retirement cost that places that electricity and coal price combination point on the retirement threshold. For reference, Figure 3 displays retirement thresholds with retirement costs above and below the implied retirement costs. Notice the retirement thresholds that intersect the electricity and coal price pairs at the time of retirement, which are represented by the solid lines in Figure 3. The level of retirement costs that correspond with that retirement threshold become the implied retirement costs. For Vermilion ST1, that works out to be \$19 million. For Crawford-7 the implied retirement cost is \$69.5 million. The higher retirement cost for Crawford-7 is driven by the larger capacity.

For price combinations below the threshold, retiring the unit maximizes the discounted profits generated by the unit. Those thresholds are upward sloping; if coal prices are low then the generator can survive with low electricity prices. If coal prices rise, it takes higher electricity prices to keep the generator online.

We repeat this process for each of the 267 coal-fired generators that have retired since 2009 and use the implied retirement costs in the next step to train our machine learning algorithm to predict retirement costs for active coal-fired generators.

4.2 Predicting Electricity Prices, Retirement Costs, and Fuel Costs

In this section we describe our machine learning (ML) model for predicting retirement costs for active generators. In this paper we briefly summarize the data, the algorithm and our results. Davis et al. (2021) describes the machine learning algorithms' implementation in much more detail and presents detailed analysis of the predictions and prediction error for each algorithm as well as some empirical tests that show

¹¹Ideally we would identify the electricity and coal prices at the time the owners made the decision to retire the unit. Unfortunately, the decision date is not consistently reported so we use retirement date. This generates some measurement error when assigning retirements to observed prices. Owners have the option of ignoring the announced retirement and extending the life of the generator up until the actual retirement date, so we think the assumption that electricity and coal prices on that date drove retirement is a reasonable one.

the predictions from the algorithms are reasonable.

We take retirement costs for each retired generator and collect data on the generator and plant characteristics for all active and retired generators. We train a machine learning (ML) model on the relationship between the characteristics of the generators and plants as well as characteristics of the county where the generator is located. We use a total of 99 features in model including name plate capacity of the generator and plant, whether the plant has access to a natural gas pipeline, whether the plant has an ash impoundment and sociodemographics like education level and unemployment rates in the generator’s county.¹² Because we observe less than 300 retirements and have nearly 100 features, ML is good choice for prediction. We then use the algorithm trained on retired generator data to predict retirement costs at active coal-fired generators.

For each generator that has retired over the last decade we know the implied retirement costs from our real options model. We use that as our target variable for the ML model. We collect the generator’s nameplate generating capacity, operating date, pollution emissions, location, regulatory status and information about the pollution abatement technology at the plant. We also compile the NERC region and utility owner for each generator.

We split the sample of retired generators into two groups. We train the ML algorithms on 80% of the sample. This portion of the sample is used to non-parametrically estimate the relationship between the generator’s characteristics and the imputed retirement costs. We then use the relationship identified in the training dataset to predict the retirement costs for the remaining twenty percent of the sample. We evaluate those predictions against the imputed retirement costs from the real options model for three different ML algorithms: LASSO, Regression Forest and XGBoost. Regression Forests perform the best, with a prediction error of around ten percent of a standard deviation.¹³

We use the Regression Forest parameters estimated on the retired generators to predict the retirement costs for all active generators. This process implicitly assumes that the relationship between retirement costs and plant characteristics is the same in the set of retired and active generators. Unfortunately, there is no way to directly test this assumption, but we find that running the ML algorithm on subsets of the retired generators does not materially affect the prediction accuracy, so we are optimistic that this assumption will hold. We use those predicted retirement costs in the retirement simulations to predict the retirement date of active generators.

¹²The full list is provided in Appendix Table 7.

¹³The other algorithms performed similarly. The prediction error for XGBoost was around 13% of a standard error and LASSO was just over 15%. The correlation coefficient among prediction for the three different algorithms was over 0.8, so the choice of algorithm for learning the relationship between generator characteristics and retirement costs is unlikely to materially affect our retirement prediction results.

4.2.1 Analysis of Retirement Costs

In this section we describe a simple analysis, first reported in Davis et al. (2021), on the retirement costs we estimate from the real options model and predict from the machine learning algorithm. Because the ML procedure for forecasting retirement costs at active generators is new to the literature, we examine the relationship between predicted retirement cost and observed retirement decision. This analysis provides evidence that the predicted retirement costs are reasonable and generators respond to them in a way consistent with economic theory.

We estimate 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:

$$\text{Logit}[\text{Prob}(\text{Retire}_{ist} = 1)] = \alpha \text{RetireCost}_i + \beta \text{GenChar}_{it} + \gamma \text{Census}_i + \theta \text{Gas}_t + \delta_s + \zeta_m + \eta_y + e_{ist} \quad (2)$$

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 3. 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;

¹⁴We 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 3: Estimates of the Probability of Retirement

	1	2	3	4	5
Retirement Costs	-0.02*** (0.00)	-0.02*** (0.00)	-0.01** (0.00)	-0.01** (0.00)	-0.01** (0.00)
Efficiency(100%)			16.60 (13.57)	13.05 (14.92)	13.11 (14.94)
Nameplate Cap (MW)			-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)
Regulated			-0.59 (0.61)	-0.76 (0.61)	-0.76 (0.61)
Ash Impoundment			-2.41*** (0.58)	-1.89*** (0.55)	-1.89*** (0.55)
Nonattainment			0.29 (0.43)	-0.05 (0.47)	-0.05 (0.47)
Mercury Control			0.42 (0.47)	0.78* (0.47)	0.78* (0.47)
Median Income				0.00 (0.00)	0.00 (0.00)
Pop. Density				-142.13 (877.47)	-139.82 (878.27)
Unemployment Rate				17.18 (46.43)	17.12 (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*** (0.19)	-0.93 (0.85)	-11.55 (11.73)	-10.62 (12.70)	-9.76 (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. * $p < .10$, ** $p < .05$, *** $p < .01$.

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 suggests that the estimated retirement costs are a good proxy for the actual retirement costs faced by the owners of coal-fired generators. Including logit coefficients means the marginal effects of changes in retirement costs are not obvious. 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%.

4.2.2 Forecasting Coal and Electricity Prices

Next, we estimate a statistical model to predict the fuel price and electricity prices at each generator. We collect a decade worth of data on fuel prices and electricity prices from publicly available data sources. We collect electricity prices at the balancing authority (BA) level from Federal Energy Regulatory Commission (FERC) form 714.¹⁵ Balancing authorities represent the smallest unit of the electric grid for which we can get electricity prices nationwide. We map each generator to its balancing authority and assign it hourly marginal price for each MWh of generation.¹⁶ We collect coal prices from EIA-923 data. That data includes plant level coal prices for plants in regulated regions and state level for plants in deregulated markets. Coal is delivered to plants, not generators, so we aggregate all coal deliveries to a plant in a month and assign the weighted average price of that coal to that generator for the month.¹⁷

We follow Pachamanova and Fabozzi (2011), who use a geometric mean reversion process to predict

¹⁵See figure A2 for map of the balancing authorities.

¹⁶The FERC 714 data includes system lambda which is the zonal locational marginal price for deregulated markets and the incremental cost of providing more electricity from thermal units in regions with no market prices.

¹⁷Coal plants can store several months of fuel on site. We choose to assign coal consumed in a month the weighted average price of coal deliveries because we consider that the opportunity cost of burning coal that month. See Chu, Holladay, and LaRiviere (2017) for more discussion of the trade-offs in this decision.

electricity and coal prices. In this model, prices hover around their long-run mean. Transitory shocks occur at random intervals that move the price above or below its mean for a period of time. As the shock fades, the price returns towards its mean.

Davis et al. (2021) shows that electricity and coal prices have behaved in a manner consistent with geometric mean reversion over the last decade. That paper estimates the means, the timing and size of the shocks that move prices from their mean and the process by which prices return to their mean. All those parameters are the inputs to the electricity and coal price simulations used in the next section. In the simulations prices evolve following the parameters of the mean reversion process and do not vary with the amount of coal-fired generation on the grid.¹⁸

5 Forecasting Coal Generators Retirements

In our baseline simulations we project retirement dates for every coal-fired generator on the grid. We assume that energy policy, electricity prices and coal costs evolve as they have over the last decade with no major changes. In the next section, we will explore the impact of potential policy changes on retirement dates.

In our business-as-usual baseline, we take the predicted retirement costs from the machine learning algorithm described above and generate retirement thresholds for each active generator using our real options model. We then simulate electricity and coal prices at each generator to find the expected time these prices cross the retirement threshold - the year that simulated future profits become small enough to justify paying the estimated retirement costs and taking the generator offline. This process requires simulating 20 years of monthly fuel and electricity prices at 492 surviving generators, making it extremely computationally intensive.

We find that under a business-as-usual scenario, 38 generators representing nearly 16 thousand MWs of capacity are predicted to retire in the first year. Another 118 generators representing over 28 thousand MWs of capacity are not predicted to retire in the next twenty years. The remaining generators are predicted to retire at some point over the next twenty years, with retirements starting at a high level and dropping off as the coal fleet shrinks.

The generators that are predicted to retire do so quickly. The bulk of the exit happens in the first five years of the simulation, with nearly 150 thousand MWs of the approximately 200 thousand MWs of active capacity gone by 2025. Most of this exit happens in the first two years of the simulation with nearly a quarter of coal generators retiring in each of those years. Our model believes that a big portion of the coal-fleet is

¹⁸This implicitly assumes that coal-fired generator retirements do not affect electricity prices or coal prices. For small numbers of retirements this is likely to hold. . Grid managers are likely to ensure that retirements are matched with entry of other generators or transmission to ensure no large electricity price effects. Retirements of large numbers of coal generators will affect demand for coal and potentially reduce prices, though exports may make up for some of the difference.

Table 4: Predicted Retirements in the Baseline Scenario

	Frac. Cap. Retiring	Median Retirement Time	Mean Retirement Time	Generator Retirements (count)	Active Cap. (MW)	Active Gen. (count)
MRO	76%	18	23	78	47,338	231
NPCC	17%	14	15	4	3,896	35
RFC	62%	19	28	123	109,737	424
SERC	57%	24	36	132	119,266	405
TRE	12%	28	43	6	27,374	43
WECC	26%	32	33	21	39,164	137

Note: Predicted generator retirements in our business-as-usual baseline scenario. Each row represent a NERC region’s results in the simulation. The first column reports the fraction of coal-fired generating capacity predicted to retire during by the end of the simulation period. The second column in the median retirement time in months conditional on retiring and the third column is the mean retirement time, again conditional on retiring. The last two columns report the active coal fired generating capacity and number of generators at the beginning of the simulation.

at risk of retirement in the very near future.

The implication of these predictions is that the coal fleet will shrink by 85% over the next twenty years, but there are a handful of generators that are likely to remain on the grid for the foreseeable future. Table 4 reports the distribution of simulated retirement times across the active generators by NERC region. In three regions, more than half of all coal-fired generating capacity is expected to retire over the next twenty years. The largest share of retirements comes in the upper Midwest (MRO) region, with around three-quarters of coal-fired capacity expected to retire. Conditional on retiring, the mean retirement time is about two years. Our model predicts a large amount of coal retirement in MRO and expects it to happen relatively soon. This is generally consistent with our announced retirement data that finds 7,500 MWs of announced coal generating capacity retirements in MRO with a median planned retirement year of 2022.

Generators that expected to remain operating through the end of the simulation are clustered in SERC (southeastern states) and slightly smaller on average than the generators that are predicted to retire. The operating dates of survivors are also comparable to those that retire in the simulation. These plants are in locations that earn slightly higher electricity prices which induces them to stay online for longer and endure occasional periods of financial losses to retain the possibility of future profits.

These predictions are contingent on the current policy landscape. We assume that the underlying process that leads to changes in prices of electricity and coal remain unchanged over the next twenty years. Obviously, this assumption is unlikely to hold in reality. In practice, environmental regulations have been mostly increasing and some generators may have chosen to retire taking that upward drift in regulation into account. Renewable penetration is increasing, but the rate of increase varies by region. For all of these reason, we interpret the baseline as a tool to understand the implications of our model, not as a prediction of actual

retirements as the policy landscape changes. In the next section we explore some counterfactual simulations that allow us to forecast retirement times under different policy scenarios.

5.1 Out of Sample Validation

We take two approaches to provide out of sample validation for the predictions of our model. First we take data through the end of 2018 and use the model to predict observed coal generator retirements in 2019 and 2020. We find that our model predicts more retirements than we observe during that time period.

Our model correctly predicts 48 of the 69 coal generator retirements observed in 2018-2019. Of the remaining twenty-one, thirteen have announced retirement dates in before 2024. Six have announced retirement dates after 2025. Two generators we predict to retire remain active and have not announced a retirement date. We take this as evidence that our model does a good job of predicting which coal-fired generators are likely to retire.

In addition to those 48 units we correctly predicted to retire in the first two years of the simulation we predict 32 other units will retire that remain active. Our model produces many more false positive retirement predictions than false negatives. Of those 34 false negative predictions 19 have announced retirement dates prior to 2025, two have announced retirement dates after 2025 and the remaining thirteen have not announced any retirement plans. The generators we predict to retire that have not announced any retirement plans are smaller, more efficient and more likely to be in the upper-midwest (MRO) than the generators whose retirements we correctly predict.

There are a number of reasons why we may be overestimating retirements in the first few years of the simulation. Our model identifies the date at which observed prices make it worthwhile to pay the retirement cost and exit the market. Once generators reach that point, it may take some time to formally announce retirement and have it appear in EIA records or media reports. It is also possible that unobserved characteristics of these generators may make retiring them more difficult than our model predicts. For example, if retirement would cause reliability issues owners may wait until other generation or transmission becomes available before retiring the units. Some owners may be operating outside the real options framework, keeping coal-fired generators operating for political reasons rather than economic. All of these are realistic possibilities, but our model cannot incorporate these issues. We continue to refine the model to understand the source of the false positives to better understand the drivers of the generators retirement decisions.

Second, we compare the predictions from our model to the National Energy Modeling System (NEMS).¹⁹ NEMS does not report the retirement decisions for individual generators, but it does forecast coal generation

¹⁹We collect NEMS results from the 2021 Annual Energy Outlook published by EIA ((US) (2021)). These results are reported in Table 54. Electric Power Projections by Electricity Market Module Region, for the reference case.

capacity.²⁰ Figure 4 indexes observed capacity at the start of the simulation to 100 and compares aggregate capacity forecast in our model to the forecast from the base case of the NEMS. Our model predicts more retirement overall through 2035. Both models predict a large drop in capacity in the next five years, but our model predicts a large drop happening sooner.

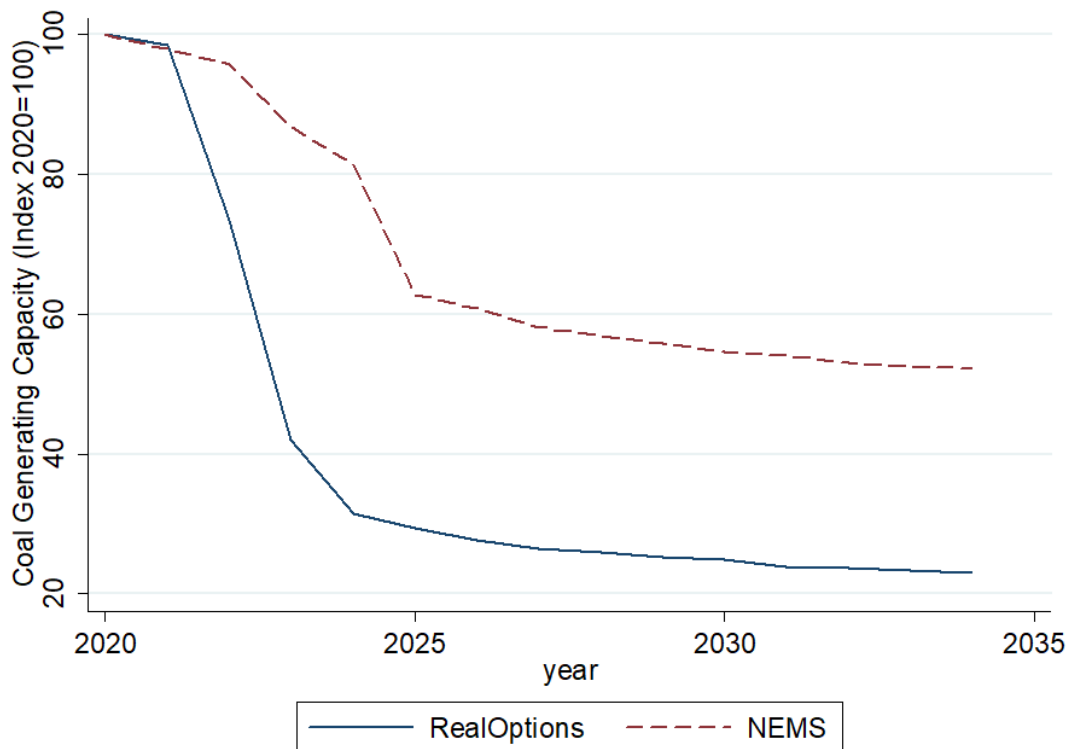


Figure 4: Comparison of Coal Generating Capacity from Real Options and NEMS Models
Note: The figure shows the forecast coal generation capacity from EIA’s National Energy Modeling System (NEMS) and our real options model. Coal capacity is indexed to 100 in 2020. Our model forecasts more retirement and that most of the retirements in the next decade will happen sooner than in the NEMS forecast.

We take these results with a grain of salt. We are predicting more retirement than the NEMS, but since the fracking boom NEMS has consistently over predicted coal generation. Forecasts produced in 2009 for coal consumption in 2019 were too high by 50%. Even forecasts produced in 2018 for 2019 coal consumption were ten percent too high. It is possible the systematic error in coal forecasts in the NEMS has been corrected and the most recent forecast will be more accurate, but our model suggests that this may not be the case.

Like our model, NEMS forecasts no coal capacity additions in the base case. Given that NEMS has over predicted coal consumption for more than a decade, if their model predicts no entry we can be reasonably

²⁰NEMS also reports capacity by “model regions”. Unfortunately, those regions do not line up well with NERC regions making it difficult to compare regional results across the two models.

confident that our model’s no entry prediction is consistent with the consensus.

6 Policy Simulations

The previous section describes predicted retirements of coal generators under a business-as-usual scenario, but policy changes are likely to have a big impact on retirement decisions. In this section we describe how two policy changes might affect which generators retire and when they choose to go offline.

Some policy makers have suggested that coal-fired generation provides important reliability benefits not rewarded by the market. Coal plants store a considerable amount of fuel on-site meaning that they are not subject to pipeline interruptions or intermittency that can affect natural gas and renewable generators respectively. Some policy makers have argued that electricity markets do not provide enough incentive to provide this type of reliability and that coal generators should be subsidized to ensure a steady supply of fuel-secure generators.

In 2018, Department of Energy Secretary Rick Perry asked FERC to “ensure financial returns” for any plant that could store at least ninety days worth of fuel on site.²¹ In 2019 the state of Ohio enacted a bill to subsidize nuclear and coal-fired generators that had been scheduled for retirement. The subsidies included \$60 million in payments through 2030 for two Midwestern coal plants owned by the Ohio Valley Electric Corporation (OVEC).²²

The rationale for proposed coal subsidies was that they would slow the retirements of coal-fired generators, but it is not clear how large of a subsidy would be needed or how long a given subsidy would delay retirement. We can use our model to understand how subsidizing coal-fired generators will affect their retirement decisions. There are lots of ways that policy makers could choose to implement a coal generator subsidy, but we model a simple output subsidy of \$20 per megawatt hour. In this simulation each coal-fired generator receives payment of the electricity they generate and an additional credit of \$20 per MWh, equivalent to around 100% of the revenue from electricity sales.

We re-run the baseline simulation above including the subsidy to coal-fired generators and identify the retirement times based on revenue from electricity sales, subsidies and costs for coal using our estimated retirement costs. The results suggest that large production credits can move the average retirement times by a few years, but are unlikely to keep coal-fired generators online far into the future.

Figure 5 reports the change in simulated retirement date from the no policy baseline to the coal subsidy. The median retirement age increases by 6 years. A number of retirements remain unchanged, either because

²¹FERC considered this request and eventually chose not to act.

²²The FBI is currently investigating the former Speaker of the Ohio House as well as other legislators for corruption related to this subsidy. It is not certain whether the plants will accept the payments or remain online.

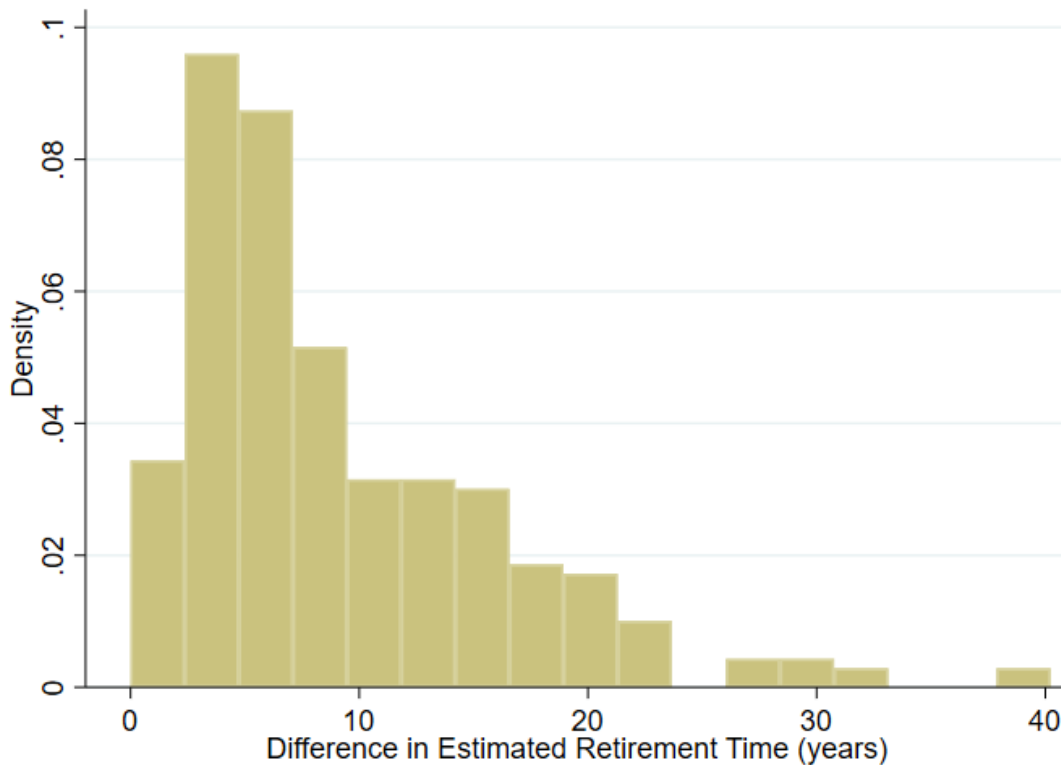


Figure 5: Difference in Retirement Times Under a \$20 Coal Generation Credit

Note: The figure shows the distribution of changes in coal generator retirement times in the no policy baseline and under a \$20 production credit. The vertical axis is the fraction of generators in a bin and the horizontal axis is how much later the generator retires under the subsidy plan.

the subsidy was not enough to prevent the plant from retiring at the beginning of the simulation or because the plant was not predicted to retire in next twenty years in the no policy baseline. The total subsidy cost in 2018 would be around \$20 billion, but as coal-fired generators retire the cost would fall.²³

In our second policy counterfactual we examine the impact of environmental regulation on coal plant retirements. The Biden Administration has recently released guidance suggesting that a social cost of carbon of \$51 per ton should be used in evaluating federal policy. We next simulate implementing that carbon tax and again find the predicted retirement date for each coal-fired generator. We weight each generators coal usage by the average carbon content for that fuel and state based on EIA’s published Carbon Dioxide Emission Factor.²⁴ The average carbon content across a million BTU of coal is a little over 200 pounds. We multiply total coal use (measured in millions of BTU) by the carbon tax rate of \$51 per ton and add that to the production cost for each generator. This shifts the retirement thresholds illustrated in Figure 3 up

²³Coal generated 1.15 billion MWhs in 2018.

²⁴We collected the data on the carbon content of fuels from EIA’s Annual Energy Outlook 2019, Table A8: Electricity Supply, Disposition, Prices, and Emissions available at: <https://www.eia.gov/outlooks/aeo/>.

Table 5: Predicted Retirements Under a Carbon Tax Scenario

NERC Region	Capacity (MWs)	Generators	Average Op. Date
MRO	381	7	1965
NPCC	500	3	1959
RFC	1300	6	1965
SERC	4796	19	1975
TRE	0	0	.
WECC	290	1	1985

Note: Predicted generator retirements under a \$51 carbon tax. Each row represent a NERC region’s capacity surviving to the end of the 20 year simulation facing a carbon tax. The first column reports the predicted coal capacity in the NERC region. The second column reports the number of generators online at the end of the simulation. The third column reports the capacity weighted average operating date for the generators that remain online through the end of the simulation.

increasing the combination of electricity and coal costs that would induce retirement.

We again simulate electricity and coal prices monthly over twenty years. We identify if those prices, inclusive of the carbon tax induce the generator to pay the retirement cost and exit the market. The carbon tax adds around six dollars to the average coal price of \$2 per MMBtu. This represents a major increase in environmental regulation and our model predicts that it will lead to large changes in retirement behavior relative to the baseline scenario.

The average generator in our sample has its retirement date moved forward around 2 and a half years by the carbon tax. The median is just over two years. Seventy-seven generators that are not predicted to retire in the baseline are predicted to retire under a carbon tax set to \$51 per ton of emissions. Those generators have an aggregate capacity of 21 thousand MWs, a little over ten percent of coal generation capacity at the beginning of the simulation. The predicted retirement time for those plants is around two years after the introduction of the carbon tax.

By the end of the simulation coal is an afterthought on the grid. Total capacity is under 1500 MWs. That is comparable to the level of synthetic gas (1326 MWs), waste heat (1089 MWs) and gas from landfills (2508 MWs) as fuel sources on the grid today. Table 5 reports the generators that remain online throughout the simulation by NERC region. All generators in TRE (Texas) have retired and only one remains in WECC (Western United States). Most of the remaining generators are in SERC (southeast).

6.1 Surviving Coal Generators

There are a handful of generators that survive through the end of the simulation in the baseline and both policy counterfactuals. These generators have some combination of low retirement costs, high electricity prices or low coal prices that induce them to stay online even in the face of significant environmental

regulation.

Table 1 in the appendix lists the generators our model predicts will remain in operating status for decades even in the face of significant environmental regulation. The list includes 22 generators with a combined capacity of around 4200 MWs compared to around 400 generators with just under 200 thousand MWs of generating capacity at the beginning of the simulation. Fifteen of those generators are in SERC reliability region in the southeastern US.

The remaining generators represent a tiny fraction of the coal generating capacity currently on the grid, but they will each produce large amounts of damaging local pollutants. These units typically have SO₂, NO_x and particulate matter controls, but a handful of units do not report any abatement equipment for one or more pollutants. Even the units with abatement technology emit hundreds, and in some cases thousands of tons, of SO₂ and NO_x per year.²⁵

Even a socially efficient carbon tax of \$51 per ton of carbon emissions is not expected to force these generators into retirement for decades. These units will also collectively generate tens of millions of tons of CO₂ per year until their retirement. Units like these represent a puzzle for policy makers concerned with reducing pollution exposure in communities around power plants and looking for ways to further reduce U.S. CO₂ emissions.

7 Conclusion

Coal-fired generators have been retiring at an increasing rate over the last decade. 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. Renewable generation will soon surpass coal-fired generation, if it has not already. This shift has consequences for regional economies, the environment, and the electric grid.

Given the importance of the coal sector to the economy, it is crucial to understand why coal generators have retired and which remaining generators are likely to retire. We have shown that coal generator retirements have increased over the decade and that smaller and older generators are the most likely to retire. Over the next five years nearly a quarter of the current coal generating capacity is planning to unplug from the grid. This exit will be concentrated in the Midwest and Western United States. Managing these retirements will require additional generation capacity and energy efficiency to replace departing base-load generation. The reduction in coal demand will continue to accelerate, leading coal producing regions into even deeper economic distress.

²⁵Northside Generating Station 2 in Duval County Florida emitted over 1600 tons of SO₂ and 774 tons of NO_x in 2019, for example.

The exit of these coal-fired generators is just the tip of the iceberg. Our analysis shows that entry is currently uneconomic for new coal-fired generators. The remaining coal-fired generators will continue to exit the market. To help predict future changes in coal generation capacity beyond announced retirements, we describe a model of coal generator retirement introduced in Davis et al. (2021). That model uses a real options framework to identify the optimal retirement time for an electric generator given electricity prices, coal prices, and retirement costs. Using the observed retirement decisions of coal-generators over the last decade, we estimate the unobserved retirement costs for each retired generator.

We then use a machine learning algorithm to map the retirement costs from the real options model on to generator characteristics. We take this mapping to the remaining active coal-fired generators to predict their retirement costs. We simulate electricity and coal prices for each generator over the next twenty years. For each surviving generator we identify whether and when it makes sense for them to retire given simulated prices and their estimated retirement costs. This creates a set of predicted retirement times for each generator currently on the grid. Around 85% of coal capacity is estimated to retire in the next twenty years. Those retirements are front loaded with nearly a quarter happening in the next three years. This is broadly consistent with the announced retirements, but our model predicts around 10 thousand MWs more retirement than has been announced. It remains to be seen whether our model is overly pessimistic about coal-fired generation prospects or whether more retirements are likely to be announced.

We use the model to estimate to policy counterfactuals. First, we estimate the impact of a production credit for coal-fired generators. State and federal governments have considered subsidizing coal plants to maintain fuel-secure base-load generation and for political reasons. Our analysis suggests that for these subsidies to have a significant impact on coal-fired generation retirement decisions, they must be very large. A \$20 per MWh production credit extends the life of the median generator in our sample by around six years. The cost of such a policy would be around five percent of total electricity expenditure.

Finally, we estimate the impact of a carbon tax on retirement dates. We estimate the carbon content of fuel for each generator on the grid and apply a \$51 a ton carbon tax. That tax rate is the current preferred estimate for the social cost of carbon (SCC), the total cost of society-wide damage from emitting a ton of carbon. Introducing that policy brings forward the median retirement time by only two years. Our model predicts that a substantial portion of the nation's coal-fired generating capacity will retire in the near future, leaving little room for policy to pull forward retirements. The carbon tax policy does have an impact on the amount of the coal-fired generation on the grid at the end of our simulation. We find that only around 1500 MWs of coal capacity remain online after twenty years under a carbon tax less than one percent of operating capacity today.

Owners of coal-fired generation have experienced increasing environmental regulation for the last two

decades. The likely expect increases in environmental regulation to continue. This background has influenced both the level of retirements we observe, but also retirement simulations in the future. Our baseline scenario assumes no change in policy. In fact, coal-fired generator owners might reasonably expect regulation to increase over the next twenty years. In that case, our baseline might over-predict retirement or more subtly generators might retire expecting environmental regulation and then regret that decision when no regulation happens. The difference hinges on whether owners know that no environmental policy changes will happen for twenty years. Some version of the carbon tax policy counterfactual might more accurately represent reality in a world where environmental regulation continues to ratchet up in the U.S.

Coal-fired generation retirements have reshaped the grid over the last decade. That process is likely to continue. Our model identifies which generators are at risk of immediate retirement and which generators are likely to remain online for the foreseeable future. We also show that there is limited scope for policy to affect retirement dates. Large subsidies extend the life of generators by a handful of years. Stringent environmental policy brings retirements forward by a couple of years. Our model suggests that no matter the policy landscape, electric grid operators, environmentalists, and policy makers should expect the retirement of coal-fired generators to continue until coal essentially disappears from the grid.

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.
- 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.
- CHU, Y., J. S. HOLLADAY, AND J. LARIVIERE (2017): “Opportunity cost pass-through from fossil fuel market prices to procurement costs of the US power producers,” *The Journal of Industrial Economics*, 65, 842–871.
- CULLEN, J. A. AND E. T. MANSUR (2017): “Inferring carbon abatement costs in electricity markets: A revealed preference approach using the shale revolution,” *American Economic Journal: Economic Policy*, 9, 106–33.
- DAVIS, R. J., J. S. HOLLADAY, AND C. SIMS (2021): “Drivers of Coal Generator Retirements and their Impact on the Shifting Electricity Generation Portfolio in the U.S.” Tech. rep., University of Tennessee Working Paper.

- DIXIT, A. K. AND R. S. PINDYCK (1994): *Investment under Uncertainty*, Princeton, NJ: Princeton University Press.
- 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.
- 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.
- RYAN, S. P. (2012): “The costs of environmental regulation in a concentrated industry,” *Econometrica*, 80, 1019–1061.
- SHAPIRO, J. S. AND R. WALKER (2018): “Why is pollution from US manufacturing declining? The roles of environmental regulation, productivity, and trade,” *American Economic Review*, 108, 3814–54.
- SHAWHAN, D. L., J. T. TABER, D. SHI, R. D. ZIMMERMAN, J. YAN, C. M. MARQUET, Y. QI, B. MAO, R. E. SCHULER, W. D. SCHULZE, ET AL. (2014): “Does a detailed model of the electricity grid matter? Estimating the impacts of the Regional Greenhouse Gas Initiative,” *Resource and Energy Economics*, 36, 191–207.
- SIEGFRIED, J. J. AND L. B. EVANS (1994): “Empirical Studies of Entry and Exit: A Survey of the Evidence,” *Review of Industrial Organization*, 9, 121–55.
- SUZUKI, J. (2013): “Land use regulation as a barrier to entry: evidence from the Texas lodging industry,” *International Economic Review*, 54, 495–523.
- (US), E. I. A. (2021): *Annual Energy Outlook 2021: With Projections to 2050*, Government Printing Office.

Table 1: Surviving Generators in Every Simulations

PlantName	Generator	County	State	Nameplate Capacity (MWs)	Operating Year
Big Bend	ST4	Hillsborough	FL	486	1985
Indian River Generating Station	4	Sussex	DE	445.5	1980
Big Bend	ST3	Hillsborough	FL	445.5	1976
Winyah	2	Georgetown	SC	315	1977
Winyah	3	Georgetown	SC	315	1980
Winyah	4	Georgetown	SC	315	1981
Northside Generating Station	1	Duval	FL	297.5	1966
Northside Generating Station	2	Duval	FL	297.5	1972
Birchwood Power	1	King George	VA	258.3	1996
Deerhaven Generating Station	2	Alachua	FL	250.7	1981
McIntosh	1	Effingham	GA	177.6	1979
B L England	2	Cape May	NJ	163.2	1964
Hammond	1	Floyd	GA	125	1954
Chesterfield	3	Chesterfield	VA	112.5	1952
Crist	4	Escambia	FL	93.7	1959
Mecklenburg Power Station	GEN1	Mecklenburg	VA	69.9	1992
Mecklenburg Power Station	GEN2	Mecklenburg	VA	69.9	1992
Schiller	4	Rockingham	NH	50	1952
Schiller	6	Rockingham	NH	50	1957
Muscatine Plant #1	7	Muscatine	IA	25	1958
Lon Wright	7	Dodge	NE	22	1963
Lon Wright	6	Dodge	NE	16.5	1957

Note: These generators do not retire any any of our policy simulations. Our model predicts these will be the final generators to retire and that they are unlikely to retire in the next ten years even if significant environmental policies are implemented.

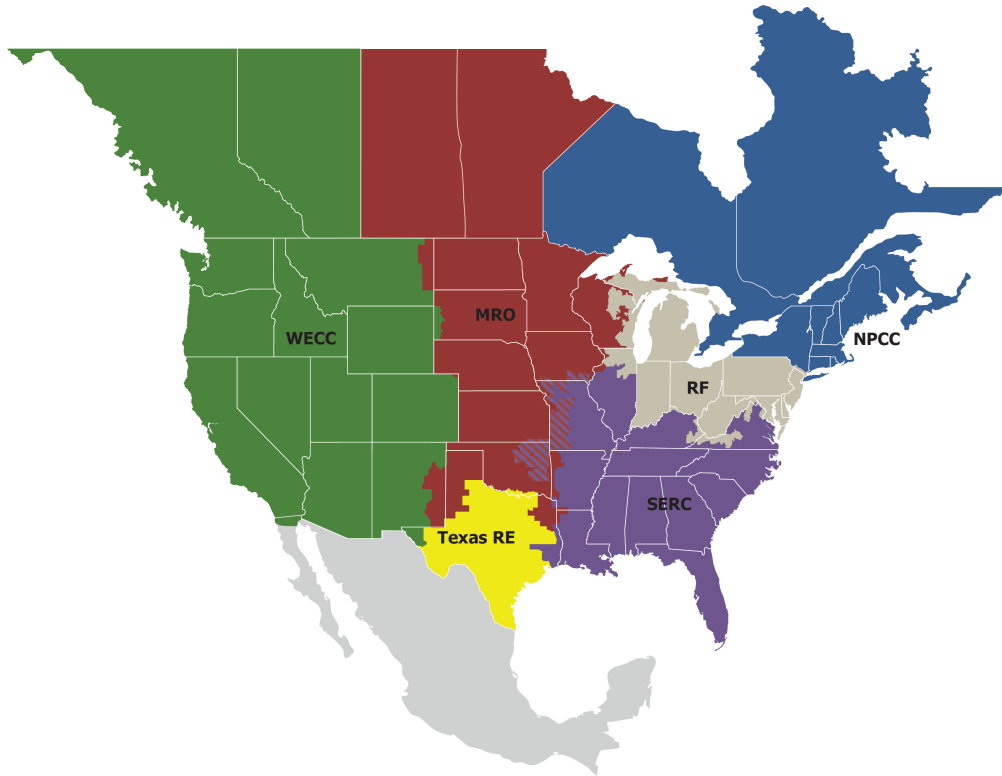


Figure A1: NERC Regions

Note: The figure shows the six mainland U.S. North American Energy Reliability Council (NERC) regions. The regions are internally well connected with transmission capacity to ensure that generation can serve load with a minimum of congestion. Source: Energy Information Administration (EIA)

Figure A1 displays the six mainland U.S. North American Energy Reliability Council (NERC) regions. NERC regions are a convenient way to summarize generating capacity geographically. NERC regional entities are charged with ensuring electrical service is efficient and reliable in their region. Electricity flows between WECC, TRE and the other four NERC regions are trivial. Flows between the four eastern regions are small, but can help balance supply and demand for electricity at the margin. NERC regions monitor generation capacity retirements and entry to ensure that there is enough generation capacity to meet demand and a safety margin to maintain service if generators or transmission lines are out of service.

Figure A2 shows the 66 balancing authorities in the U.S. Balancing authorities are responsible for managing the balance of electricity supply and demand hour-by-hour to maintain grid in their geographic region. We collect electricity market price data for these regions because they represent the most disaggregated electricity prices for which we can collect data nationwide. Some generators are in wholesale markets where we can collect electricity prices at the generator level. Unfortunately, large sections of the country are not in electricity markets and this detailed data is not available. In practice, the balancing authority data is

Regions and Balancing Authorities

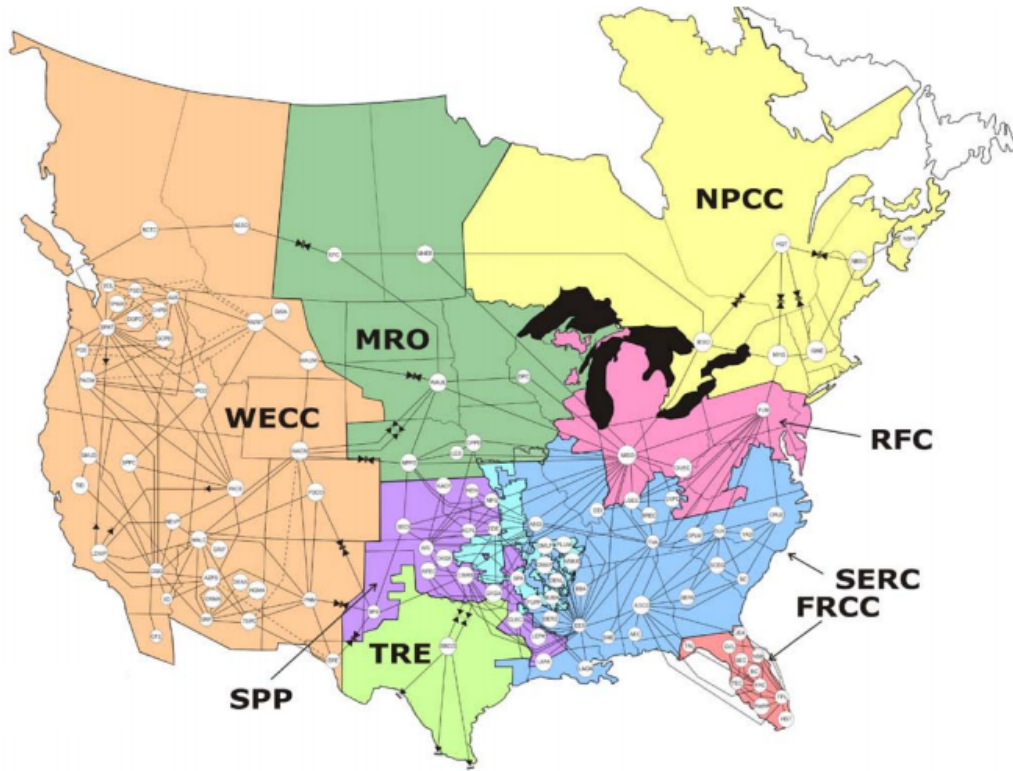


Figure A2: Balancing Authority Map

Note: The figure shows the 66 balancing authorities in the United States. Source: Energy Information Administration.

comparable to zonal prices in an ISO market. Differences between zonal and nodal prices would typically only occur during times when the grid is congested.

Variable	Source
Median Income	U.S. Census (2010)
Population Density	U.S. Census (2010)
Unemployment Rate	U.S. Census (2010)
Male Higher Edu.	U.S. Census (2010)
Female Higher Edu.	U.S. Census (2010)
Utility	EIA 830 Data
Nameplate Capacity MW	EIA 830 Data
Operating Year	EIA 830 Data
Natural Gas Pipeline Access	EIA 830 Data
Total Plant Generating Capacity	EIA 830 Data
Fraction Plant Coal Capacity	EIA 830 Data
Ash Impoundment	EIA 830 Data
Ash Impoundment Lined	EIA 830 Data
Ash Impoundment Status	EIA 830 Data
Nox Rate	Shawhan et al. (2014)
So2 Rate	Shawhan et al. (2014)
Heat Rate (efficiency)	Shawhan et al. (2014)
Dollars of Environmental Damage per MWhr	Shawhan et al. (2014)
State	EIA 830 Data
Sector	EIA 830 Data
Energy Source (Coal Type)	EIA 830 Data
NERC Region	EIA 830 Data
Regulatory Status	EIA 830 Data
Ownership (Single/Joint)	EIA 830 Data
Mean Fuel Price	Authors Calculations from EIA 923 Data
Variance Fuel Price	Authors Calculations from EIA 923 Data
Mean Electricity Price	Authors Calculations from EIA 923 Data
Variance Electricity Price	Authors Calculations from EIA 923 Data

Note: This table reports the data used in the machine learning algorithm that maps retirement costs for retired generators to active generators. After creating categorical variables, there are 99 features for 720 active and retired generators.