Does Regulation Distort Exit Decisions? Evidence from U.S. Power Plants

Lucas W. Davis   Paige Weber*

February 2026

Abstract

Hundreds of power plants have closed in the United States since 2010, including 130+ gigawatts of coal and 50+ gigawatts of natural gas. In this paper, we highlight the potential for regulation to distort this type of exit decision. Using generator-level data from 2010–2023, we show that regulated units have been 45% less likely to exit than unregulated units. For unregulated units, exit decisions are made based on wholesale electricity prices, ongoing capital costs, and other traditional economic factors. In contrast, owners of regulated units are largely insulated from these factors and, in some cases, have a strong incentive to continue operating capital-intensive equipment.

Key Words: Regulatory Bias, Averch-Johnson Effect, Electricity Markets, Coal, Natural Gas
JEL: D24, L94, Q41, Q48, Q54

*(Davis) Haas School of Business at University of California, Berkeley and Energy Institute at Haas; *lwdavis@berkeley.edu*. (Weber) Energy and Resources Group at University of California, Berkeley and Energy Institute at Haas; *paigeweber@berkeley.edu*. We are thankful for helpful comments to Josh Blonz, Jim Bushnell, Steve Cicala, Catie Hausman, Ashley Langer, Erin Mansur, Marten Ovaere, Nick Ryan, three anonymous referees, and seminar participants at UC Berkeley and Stanford. We have not received any financial compensation for this project nor do we have any financial relationships that relate to this research. The analysis relies entirely on publicly-available data and all data and code will be posted on our websites upon completion of the project.

1 Introduction

Hundreds of power plants have closed in the United States since 2010, including 130+ gigawatts of coal and 50+ gigawatts of natural gas. These exit decisions have important implications for electricity markets, grid reliability, and the environment. Holland et al. (2020), for example, calculate that the total damages from U.S. power plants decreased by more than $100 billion annually between 2010 and 2017, with plant closures being one of the largest single factors.

In this paper, we highlight the potential for regulation to distort this type of exit decision. Using generator-level data from 2010–2023, we show that regulated units have been 45% less likely to exit than unregulated units. This difference remains after controlling for observable characteristics and implies that regulated units operate 8.4 years longer than unregulated units on average.

Our analysis focuses on coal and natural gas, the two largest sources of electricity generation in the United States as of 2010. The two types of electricity generation are quite different in many ways, including a much higher overall retirement rate for coal. But in both cases, we find that regulated units have been less likely to exit than unregulated units. The effect is similar in magnitude for coal and natural gas and strongly statistically significant across specifications.

We interpret these results by pointing out that regulated and unregulated firms have very different incentives. Unregulated firms – known as "independent power producers" in electricity markets – make exit decisions based on wholesale electricity prices, fuel costs, ongoing capital expenditures, and other traditional economic factors. In contrast, regulated firms are largely insulated from these factors and instead subject to the preferences and priorities of local regulators. For investor-owned utilities, there is also a guaranteed rate-of-return, which creates a strong incentive to continue operating capital-intensive equipment.

A novel feature of our analysis is that we also separately examine investor-owned and publicly-owned utilities. The latter category has received much less attention in the economic literature, and there is reason to think that behavior by publicly-owned utilities could be quite different. However, we find that both forms of regulation are associated with significantly lower exit rates than unregulated companies, consistent with both being less exposed to market conditions.

We also report results from an alternative specification aimed at testing for market power. High exit rates for unregulated companies might alternatively reflect efforts to push up prices for remaining units owned by the same company. However, when we divide unregulated companies into those with small, medium, and large portfolios, we do not find higher exit rates for larger portfolios, nor do we find smaller exit rates for smaller portfolios. Thus, we conclude that our main results are unlikely to be driven by market power.

In addition, we test whether regulated units were less likely to exit in states that are major producers of coal or natural gas. Our hypothesis was that in states like Wyoming, West Virginia, and Kentucky, with significant coal production, coal producers would exert political pressure on state politicians and regulators to keep coal plants open. We do not find any evidence of this form of regulatory capture, though we cannot rule out the role of other local preferences and priorities that are not related to in-state coal and gas production.

Our paper does not attempt to quantify economic or environmental impacts, nor does it attempt to say whether this behavior is efficient or inefficient. Exit decisions impact wholesale electricity prices, grid reliability, emissions, future investment, and other outcomes. Quantifying these impacts would require modeling both short- and long-run equilibrium, and goes beyond the scope of our analysis.

In a closely related paper, Gowrisankaran et al. (forthcoming) builds a structural model of rate-of-return regulation during an energy transition. In their model, utilities balance lowering operating costs with maintaining coal capacity. Our paper is quite different methodologically, but both papers have at their core the idea that rate-of-return regulation insulates regulated firms from traditional market forces, and our paper provides empirical corroboration for their modeling approach.[1]

Previous studies of power plant exit decisions have tended to focus instead on wholesale electricity prices and other traditional market forces. Linn and McCormack (2019) finds that low natural gas prices and weak demand for electricity explain 80% of coal plant retirements in the U.S. eastern interconnection between 2005 and 2015,

[1] They find, for example, "Thirty years after a sudden energy transition, a cost minimizer has retired 71% more coal capacity than the regulated utility". Further afield, Borrero et al. (2023) provides a new approach for estimating ramping and operations costs for coal and other power plants and Gowrisankaran et al. (2025) estimates a dynamic oligopoly model for coal plants that recovers owners' beliefs and uncertainty about future MATS enforcement. There are also several papers that examine the short-run relationship between natural gas prices, coal plant generation, and carbon emissions. Cullen and Mansur (2017), for example, finds that a $6 per mmBTU decrease in natural gas prices causes a 10% decline in carbon emissions, almost entirely through reduced coal plant generation. Fell and Kaffine (2018) finds that natural gas prices and wind generation jointly explain most of the decline in U.S. electricity generation from coal between 2007 and 2013.

whereas nitrogen oxides emissions caps and the Mercury and Air Toxic Standards (MATS) had relatively little effect.[2] Similarly, Rebecca Davis and coauthors find that wholesale electricity prices, coal prices, generator size, and generator age are the most important drivers of U.S. coal plants retirements 2009–2017 (Davis et al., 2022).

Our paper contributes to a small empirical literature on the Averch Johnson effect. This regulatory bias has been discussed for decades (Averch and Johnson, 1962), but empirical evidence is limited, with recent papers focused on capital-intensive emission control technologies (Fowlie, 2010; Cicala, 2015) and how overall capital investments increase with the allowed rate of return (Dunkle Werner and Jarvis, 2025). Although the literature has focused on new investments, our paper emphasizes that this same regulatory distortion also affects exit decisions.

In addition, our paper contributes to the literature on the effects of electricity market deregulation. Previous studies have examined market power (Wolfram, 1999; Borenstein et al., 2002; Bushnell et al., 2008), operating efficiency (Fabrizio et al., 2007; Davis and Wolfram, 2012; Chan et al., 2017; Demirer and Karaduman, 2025), procurement practices (Cicala, 2015), and allocation of production between plants (Mansur and White, 2012; Cicala, 2022). The effect of deregulation on plant exit decisions has received little attention.

[2] Relatedly, Coglianese et al. (2020) finds that 92% of the decline in U.S. coal production 2008 to 2016 can be explained by low natural gas prices, with environmental regulations explaining an additional 6%.

2 Background

In this section, we discuss: (1) how regulation works, (2) the economic incentives for regulated plants, (3) deregulated markets, (4) the economic incentives for unregulated plants, and (5) public ownership.

2.1 How Regulation Works

Under rate-of-return regulation, the regulated firm is granted the exclusive right to operate in a particular geographic area, but, in exchange, the regulator sets rates based on the regulated firm's costs. Rate-of-return regulation is widely used in U.S. electricity markets including electricity generation in most states, as well as electricity transmission and distribution. Rate-of-return regulation is also widely used in other markets, including pipelines for natural gas, crude oil, and refined petroleum products, as well as in water distribution, sewage, trash, and recycling.

The regulator's objective is to set rates so that the regulated firm exactly breaks even, including the opportunity cost of capital investments. The first step for the regulator is to determine the firm's total costs (i.e. the "revenue requirement"), including operating expenses and capital costs. For power plants, operating expenses include fuel costs, salaries, taxes, and insurance. Capital costs include depreciation and the opportunity cost of capital. Whereas operating expenses are observed and relatively straightforward to calculate and put into rates, capital costs are more complicated.

A central concept in rate-of-return regulation is the "rate base", i.e. the current value of the firm's capital investments. As the regulated firm makes capital investments, these investments enter the rate base. Then, over time, these investments are depreciated according to multiple-year schedules determined by the regulator. The amount that any particular investment counts toward the rate base decreases every year until, eventually, the investment is "fully depreciated" at which point the investment no longer appears in the rate base.

A stylized example is helpful. Suppose a regulated firm makes a $1 billion capital investment in a new power plant. During the first year of operation, the $1 billion enters the rate base.[3] If the firm's allowed rate of return is 10%, then rates are increased so that the firm receives an additional $100 million in the first year. If the regulator uses 20-year straight-line depreciation, then $50 million of the plant's value is depreciated in each year, and the rate base would decline to $950 million in year 2, $900 million in year 3, and so on.

Regulated firms make exit decisions in consultation with the regulator. The process typically begins with the regulated firm making the decision that it wants to retire the plant. This usually needs to be approved by the state regulatory commission or other regulatory body. The exact details vary, but there is generally some review of potential impacts to reliability or other "impairment of service".[4] Thus observed

[3]A related concept in rate-of-return regulation is the "used and useful rule". In an effort to approximate what happens in unregulated markets, this is the idea that a capital investment enters the rate base only once that investment is operational. Even if it takes a decade to complete a new power plant, for example, those capital costs (including financing) do not enter the rate base until the plant begins operation. The "useful" refers to the capital investment serving some public purpose. Although it is relatively rare in practice, the regulator can decide to "disallow" a capital investment if it is not deemed "useful".

[4]See, for example, National Association of Clean Air Agencies, "Implementing EPA's Clean Power Plan: A Menu of Options," May 2015. "In traditionally regulated jurisdictions, electric gen-

exit behavior by regulated firms reflects both decisions by the regulated firm as well as the decisions made by the regulator during this review process.

2.2 Incentives For Regulated Plants

Under rate-of-return regulation, the owners of regulated power plants are largely insulated from market conditions. When operating costs increase, for example, these costs are passed along to ratepayers, with no implications for shareholders. The regulated firm also bears little risk with regard to capital investments. Whether broader economic conditions end up favorable or unfavorable, the regulated firm continues to receive a guaranteed rate-of-return.

This insulation from market conditions means that owners of regulated plants face little urgency with regard to exit decisions. Suppose that market conditions end up being unfavorable for a particular power plant because, for example, fuel costs have gone up. These unfavorable market conditions simply do not matter to the owner of the regulated plant, as they simply pass along these higher costs to ratepayers.

Owners of regulated plants are also insulated from market-based environmental policies. For example, consider a carbon tax or cap-and-trade program for carbon dioxide. These policies increase the marginal costs of coal and natural gas electricity production, with equilibrium wholesale electricity prices typically increasing as a result. However, these changes in market conditions are largely irrelevant for the

erating unit (EGU) owners' retirement decisions must be reviewed and approved by state regulatory commissions... this regulatory review is intended to examine whether or not abandoning the EGU will affect the company's service, specifically calling out "impairment of service" (i.e., reliability) as a criterion."

owner of the regulated plant, as they are passed along to ratepayers and do not impact profits.

Firms subject to rate-of-return regulation have a strong incentive to make and maintain capital-intensive investments. Probably the single hardest job of the regulator is to determine the firm's allowed rate-of-return on capital investments. In practice, the rate-of-return is typically set higher than the firm's cost of capital, creating a distortion and biasing the firm towards capital-intensive technologies (Averch and Johnson, 1962; Joskow, 1972; Dunkle Werner and Jarvis, 2025).

The literature has mostly focused on how this regulatory distortion affects new investments. But this same regulatory distortion also impacts exit decisions. When a plant is closed, this capital investment leaves the rate base and is immediately depreciated. This is undesirable from the perspective of the regulated firm because it stops receiving the guaranteed rate-of-return on the capital investment.

Once a plant is fully depreciated, the regulated firm no longer has an incentive to keep it open. At that point, the regulated firm might actually want to close the plant so that it can build something else and start this process over again with a new capital investment in the rate base. But outcomes in this case are uncertain. A regulator might decide that no replacement plant is necessary. Or, the regulator might decide that it will be an unregulated firm, not the regulated firm, that makes this new capital investment.

Another consideration is recurring capital costs. Over time, plant equipment wears out and must be replaced, requiring additional capital investments. For a regulated firm, this goes into the rate base on which the firm earns a guaranteed rate-of-return. Thus, for a regulated firm, recurring capital costs do not hurt profits and may even provide an incentive to keep the plant open. Indeed, if there are enough recurring capital costs, then a plant is unlikely to ever become fully depreciated.

Largely insulated from traditional economic factors, regulated firms are instead subject to the preferences and priorities of local regulators. For example, reporting from West Virginia Public Broadcasting describes the "unwavering support" for coal in West Virginia at all levels of state government including the governor, state legislature, and state public service commission.[5] The coal industry has a long history in West Virginia and Appalachian Power (a regulated company) has long-term contracts with coal plants in the state.

2.3 Deregulated Markets

Since at least the 1980s, economists have articulated a competitive vision for electricity markets (Joskow and Schmalensee, 1988). The standard prescription was to deregulate generation while keeping transmission and distribution regulated. Beginning in the late 1990s, hundreds of power plants in about a dozen states were sold from regulated companies to unregulated companies, also known as "independent power producers" or "merchant generators". Wholesale markets were established at the same time, with electricity buyers required to participate in these markets.

The hope was that a competitive generation market would lead to greater efficiency,

[5] "West Virginia Relies Mostly on Coal for Its Electricity. Customers Are Paying a Heavy Price" WBUR Radio, September 25, 2025.

both in the short-run and long-run. In the short-run, competition would create an incentive for firms to operate power plants efficiently and, through the newly established wholesale markets, to allocate production to low-cost plants. In the long-run, competition would impose market discipline on new investments and eliminate the regulatory distortion towards capital expenditures.

U.S. deregulation efforts were stalled by the California Electricity Crisis in 2000 and 2001. Since then, the regulatory status of U.S. power plants has been essentially frozen, with little attempt to further deregulate, but also little attempt to regulate those plants that had already deregulated. Thus today's mix of regulated and unregulated plants still reflects which states deregulated first over 25 years ago. Deregulated plants are found mainly in California, Texas, the Northeast, and parts of the Midwest, while most plants in the rest of the country remain regulated (Borenstein and Bushnell, 2015). White (1996) shows that deregulation tended to occur in states with high electricity prices, where consumers potentially had the most to gain from deregulation.

2.4 Incentives for Unregulated Plants

Owners of unregulated plants make exit decisions based on wholesale electricity prices, operating costs, and other traditional economic factors. These are profit-maximizing firms, so they will close a plant when it ceases to be profitable. How much the plant cost to build originally, or however much the firm paid for the plant is a sunk cost and, therefore, irrelevant. There is no rate base and certainly no guaranteed rate-of-return.

Recurring capital costs are a helpful example. For a regulated firm, these costs go into the rate base on which the firm earns a guaranteed rate-of-return. In contrast, for an unregulated firm, recurring capital costs come directly from firm profits. Owners of unregulated plants do not like to make this type of capital expenditure and may choose to close the plant to avoid incurring such costs.

Owners of deregulated plants are also directly impacted by market-based environmental policies. Whether it is paying a carbon tax or paying for permits as part of a cap-and-trade system, carbon pricing increases operating expenses for coal and natural gas plants. Owners of carbon-intensive units, in particular, are made worse off and may choose to exit.

Market conditions tended to be unfavorable for U.S. power plants during our study period. Previous papers have pointed to low natural gas prices, weak electricity demand, and the entry of renewables, among other factors, as hurting the economic prospects for U.S. power plants (Fell and Kaffine, 2018; Linn and McCormack, 2019). Owners of unregulated plants were fully exposed to these unfavorable conditions, so many owners made the decision to exit.

2.5 Public Ownership

A third category is public ownership. This includes municipally owned utilities like the Los Angeles Department of Water and Power (LADWP), state-owned utilities like the New York Power Authority (NYPA), and federally owned utilities like the Tennessee Valley Authority (TVA), as well as member-owned electric cooperatives like Old Dominion Electric Cooperative (ODEC), which owns power plants in Virginia and Maryland.

The governance of publicly owned utilities varies widely. LADWP, for example, has a five-member board that is appointed by the mayor of Los Angeles and approved by the Los Angeles city council. NYPA has a nine-member board appointed by the governor of New York and approved by the state Senate. TVA has a nine-member board appointed by the president and approved by the Senate. ODEC has a 10+ member board elected by its members. These boards then select chief executives, financial officers, and other executives to manage operations.

There are reasons to think that publicly owned utilities are insulated from market conditions. When operating costs increase, for example, these costs are passed along to ratepayers. There are no shareholders, nor are these utilities owned or co-owned by private individuals. The risk of capital investments is borne by the ratepayers. There is no rate base for a publicly owned utility, and thus no regulatory bias toward capital-intensive investments, but also no direct private economic consequences when economic conditions end up unfavorable for a particular capital investment.

However, at the same time, there are also reasons to think that publicly owned utilities would be less insulated. The governance structure of publicly owned utilities potentially makes them more responsive to the concerns of local ratepayers. Board members typically live in the utility service area, know and work with ratepayers and other constituents, and often view keeping rates low as an important part of their duty. Publicly owned utilities make power plant closure decisions based on the needs of the utility and the needs of ratepayers, but also influenced by local priorities and preferences. For example, a publicly owned utility might value retaining local employment in a way that an unregulated firm is not incentivized to do.

3 Data and Descriptives

3.1 Data Construction

We constructed for this analysis a comprehensive dataset describing U.S. power plants over the period 2010–2023. Our primary source is EIA-860, an annual survey of U.S. power plant operators conducted by the U.S. Department of Energy, Energy Information Administration (EIA). EIA-860 includes generator, plant, and operator-level information about all U.S. power plants over one megawatt capacity. The key variables from EIA-860 for our analysis are regulation type, name of owner, fuel type, capacity (in MW), technology type (e.g. combustion turbine), year opened, and year retired (if retired).

We restrict the analysis to coal and natural gas units. These were the two largest categories of U.S. electricity generation as of 2010, with coal and gas responsible for 45% and 24% of generation, respectively, in that year.[6] The other categories in 2010 were nuclear (20%), conventional hydro (6%), wind, solar, and other renewables (4%), and petroleum/other (1%). We perform all analyses at the generator unit level,

[6]See EIA "Electric Power Annual 2010", Table 2.1.A Net Generation by Energy Source. Coal and natural gas represented 69% of total U.S. generation in 2010. The third largest category of generation, nuclear, would have been interesting to include, but there are too few total units to support an empirical analysis.

rather than at the power plant level, as it is common for units within a plant to open and/or to retire in different years.

We include in the analysis all units that were operating as of 2010. Thus, we exclude units that retired in or before 2010, as well as units that began operation after 2010. We selected 2010 as the starting point because this is the beginning of a period of low wholesale electricity prices. By 2010, hydraulic fracturing had substantially increased the supply of natural gas, lowering the cost of gas-fired electricity generation, and decreasing wholesale electricity prices (Hausman and Kellogg, 2015). Thus, our analysis describes retirement behavior for U.S. coal and natural gas units that were operating as of 2010.

Regulation type is an important variable in our analysis, so we devote considerable attention to ensuring that this variable is measured as accurately as possible. The key question in EIA-860 is "*What type of entity is the principle owner and/or operator? i) Cooperative, ii) Investor-Owned Utility (IOU), iii) Independent Power Producer (IPP), iv) Municipally-Owned Utility, v) Political Subdivision, vi) Federally-Owned Utility, vii) State-Owned Utility, viii) Industrial (principal business is not electricity generation, ix) Commercial (principal business is not electricity generation).*[7]

We drop units that operate for industrial or commercial purposes (i.e. the last two categories). These are units for which the primary purpose is something other than

[7] The "entity type" question was not asked in EIA-860 until 2013. In addition, this information is missing for a small number of units in some years during the period 2013–2023. To address both forms of missing information, we impute using information for that same unit in other years. In a small number of cases (<2%), we use additional information about the regulation type available from a related "sector type" question in EIA-860.

selling electricity. Instead, these units typically provide electricity or heat (or both) to co-located facilities and face a very different set of constraints and incentives. These industrial and commercial units were less than 4% of capacity in 2010.

The remaining units we categorize as "regulated" or "unregulated". In particular, we categorize independent power producers (i.e. the third category above) as "unregulated", and all other remaining categories as "regulated". Later in the paper, we further break down this regulated category to distinguish investor-owned utility (i.e. the second category above), from the various types of publicly-owned utilities, cooperative, federal, municipal, municipal marketing authority, political subdivision, and state (i.e. the first, fourth, fifth, sixth, and seventh categories above).

For our analysis, we characterize units based on their regulation type as of the first available year in which that information is available. As mentioned earlier, there has been little effort anywhere in the U.S. during our study period toward deregulation, but also little effort to regulate again those units that were deregulated during earlier decades. There are a small number of cases for which the EIA-860 data show changes in regulation type over time. Some of these may be reporting errors. Moreover, it is not completely clear how to treat time-series variation in regulation type in the context of retirement. For example, would a unit sale from a regulated to unregulated company constitute a retirement for the regulated company? We instead choose to frame our analysis as describing the retirement choices of the fleet as it appears in the beginning of our study period.

We apply this same approach to other unit characteristics, describing units based on their fuel type and capacity as of the first available year in which that information is available. For inherently cross-sectional characteristics such as year opened and year retired for which different values across years reflect data errors or misreporting and not time series variation, we use the most recent data, under the assumption that reporting should improve over time. With retirements, this also likely reflects units that were expected to retire but ended up not retiring until the following year.

A notable trend during our study period was the conversion of units from coal to gas. Between 2011 and 2023, 15% of the coal units in the 2010 fleet, and close to 9% of coal capacity was converted from coal to gas. We do not treat these as retirements. Consistent with our approach of studying units based on their characteristics at the beginning of the study period, these units remain in our sample with the fuel type as measured in 2010. Our view is that these conversions are quite different from retirements. For an investor-owned utility, for example, a conversion allows the company to maintain or even increase the amount of capital included in the rate base. We leave this related question of fuel switching for future research.

3.2 Summary Statistics

The U.S. electricity generation market has several features that make it particularly suited for an empirical analysis of regulatory distortions. Most importantly, the U.S. has a mix of regulated and unregulated units, with thousands of units in each category.

Table 1 reports summary statistics. Panel (A) describes all units operating as of

2010. There are more natural gas units than coal units, but coal units tend to be larger and older. Panel (B) describes the units that retired between 2011 and 2023. More than 1500 units closed during our study period. More gas units closed than coal units, but coal units were much more likely to retire. Of the units operating in 2010, 55% of coal units retired between 2011 and 2023, compared to 18% of gas units. With both coal and gas, the units that retired tended to be older.

Figure 1(A) shows capacity by regulation type. We distinguish between investor-owned utilities (IOU), publicly-owned utilities (POU) and independent power producers (IPP). For coal, IOU units comprise the largest share of capacity, although there are also substantial shares for both POU and IPP. For gas, IOU and IPP have similar total capacity, with POU representing a smaller but still substantial share.

Figure 1(B) shows year opened by regulation type. Since this includes all units operating as of 2010, the year opened effectively illustrates the age distribution as of 2010. Coal units tend to be older, with many units open since the 1950s and 1960s. Most gas units opened after 1970, and there was a large spike in openings in the early 2000s. In terms of regulation type, IPP units tend to have later opening years, both for coal and gas, but there is considerable overlap in the distributions.

3.3 Retirements 2011–2023

Figure 2 plots U.S. coal capacity (panel A), and cumulative retired capacity (panel B), over our study period. Total coal capacity declined steadily between 2010 and

Figure 1: U.S. Power Plants by Regulation Type

(A) Capacity as of 2010

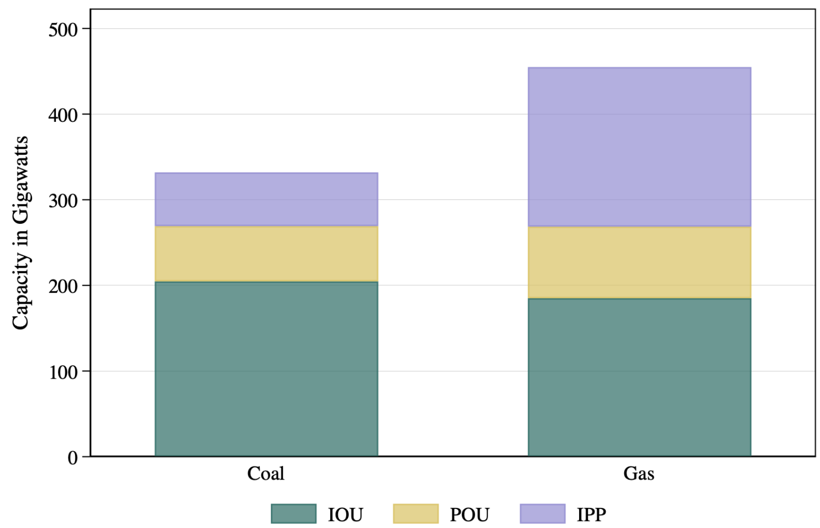

(B) Year Opened

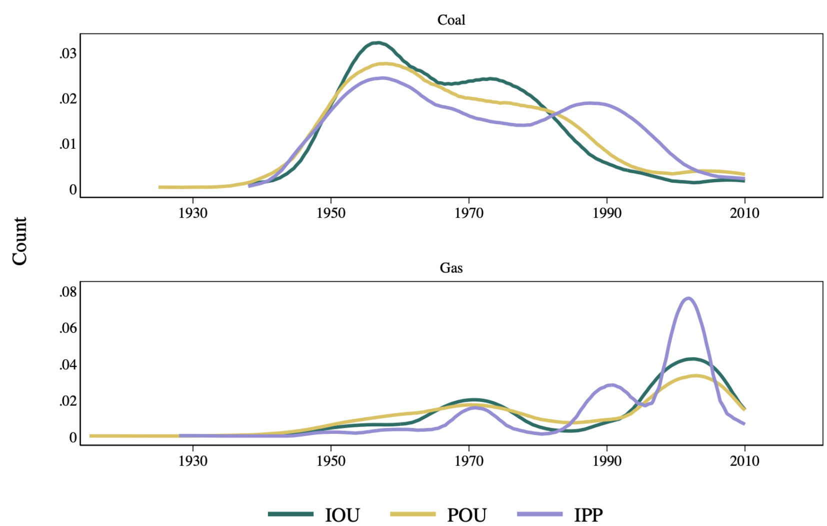

Notes: This figure describes U.S. power plants by regulation type as of 2010. We focus on coal and natural gas, the two largest categories of generation as of 2010. Regulation types include investor-owned utilities (IOU), publicly-owned utilities (POU) and independent power producers (IPP).

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max
All units:				
Coal (n = 1189, regulated = 964)				
Size (MW)	279.1	272.2	0.4	1,425.6
Year Opened	1968	15	1925	2010
Gas (n = 4960, regulated = 3303)				
Size (MW)	91.7	103.7	0.2	1,027.0
Year Opened	1989	17	1915	2010
Retired units:				
Coal (n = 652, regulated = 506)				
Size (MW)	164.3	182.6	0.4	1,425.6
Year Opened	1962	13	1925	2010
Gas (n = 898, regulated = 600)				
Size (MW)	62.6	110.0	0.2	799.2
Year Opened	1972	17	1928	2010

NOTES: This table reports summary statistics for our dataset of electricity generating units. Panel (A) describes all units in operation as of 2010. Panel (B) describes all units that retired between 2011 and 2023. For each category, we report the total number of units as well as the number of units that are "regulated", defined here as including both investor-owned utilities and publicly-owned utilities. Statistics for year opened have been rounded to the nearest integer.

2023, with total coal retirements reaching 130+ gigawatts by the end of the period. Panels C and D plot the same information for natural gas. Total gas capacity decreased steadily during our study period, but by much less than coal. By the end of our study period, total gas retirements reached 50+ gigawatts.

Table 2 reports the percent retired by fuel type and regulation type. Unregulated coal units were significantly more likely to retire during our sample period – 65% compared to 53% for regulated coal units. The difference is highly statistically significant (p-value .001). For natural gas, the percent retired is almost identical – 18% in both cases – for regulated and unregulated units.

Comparing retirement shares by regulation type is suggestive, but does not fully answer our research question. In particular, we have already seen differences in opening year for regulated and unregulated units. In any given year, we would expect older units to be more likely to retire than younger units. Thus, in the following subsection, we examine the average age at retirement.

Table 2: Percent Retired

	(1) Regulated Units	(2) Unregulated Units	(3) p-value (1) vs (2)
Coal	52.5%	64.9%	0.001
	(506 of 964)	(146 of 225)	
Gas	18.2%	18.0%	0.876
	(600 of 3,303)	(298 of 1,657)	
All	25.9%	23.6%	0.053
	(1,106 of 4,267)	(444 of 1,882)	

Notes: This table reports the percent retired by fuel type and regulation type. In each category, the denominator is the number of units operating in 2010, and the numerator is the number of units that retired over the period 2011–2023. The final column reports p-values from a t-test of the null hypothesis that retirement is equally likely for regulated and unregulated units.

Figure 2: U.S. Coal and Gas Electricity Generating Capacity, 2010–2023

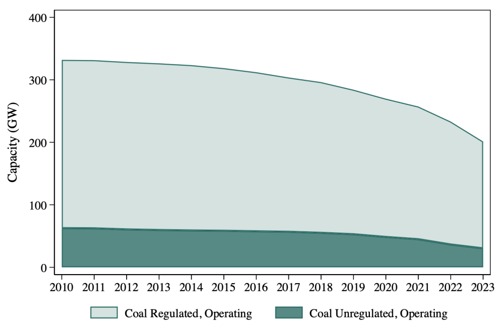

(A) Coal Operating

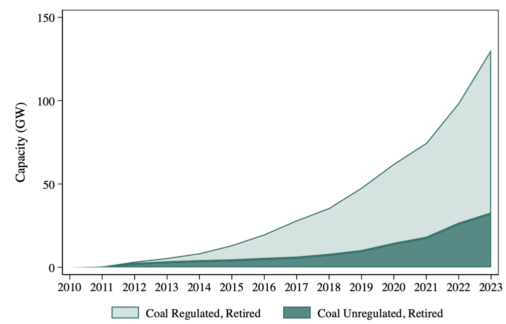

(B) Coal Retired

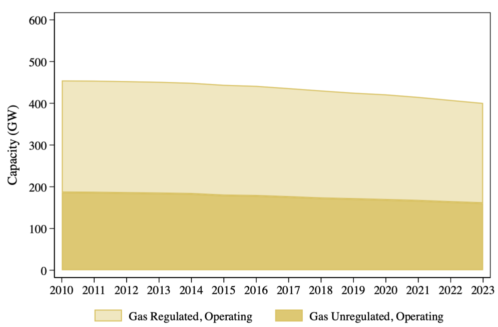

(C) Gas Operating

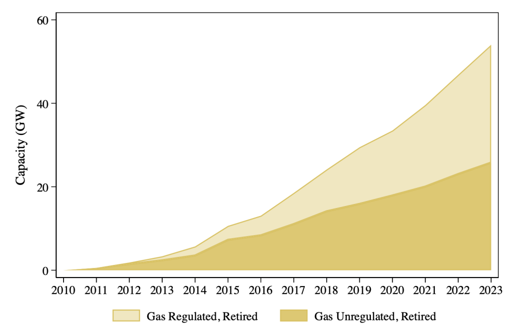

(D) Gas Retired

Notes: This figure plots total capacity and cumulative retired capacity for U.S. coal and natural gas plants. As discussed in the text, our dataset includes all generating units that were operating in the United States as of 2010. The goal of our research is to understand retirement behavior so we have excluded units that opened after 2010, and the total capacity in the figure does not match aggregate U.S. coal and gas capacity in those years.

3.4 Average Age At Retirement

Figure 3 plots the distribution of retirement ages by fuel and regulatory status for units that were online in 2010 and retired between 2011 and 2023. Among coal units, IPP retirement ages cluster around 30 and 60 years, while IOU and POU units' ages are more concentrated around age 60. For gas units, IPP retirement is similarly bimodal, with peaks around 30 and 50, whereas IOU and POU retirements are centered around age 50. Notably, only 9% of IOU units retire before age 40, compared to 40% and 26% for IPP and POU, respectively. The lack of IOU retirements before age 40 could reflect the incentives for regulated companies to continue operating capital-intensive plants as long as they remain in the rate base.[8]

Table 3 reports the average age at retirement for regulated and unregulated units. Regulated units tend to be considerably older at retirement than unregulated units. Considering both coal and gas, regulated units average 50 years at retirement, compared to only 42 for unregulated units. The difference (8 years) is strongly statistically significant with a p-value < 0.001.

The pattern goes in the same direction for coal and gas. In both cases, regulated units have a considerably higher average age at retirement than unregulated units. In both cases, the differences are strongly statistically significant with p-values < 0.001. It may initially be surprising that this table shows a difference for gas, whereas Table

[8] This lack of exit prior to age 40 for regulated units is noted by Duff (2020). Focusing on U.S. natural gas plants between 2006 and 2017, Duff (2020) finds that a $1 increase in the price spread between coal and natural gas decreases the probability that an unregulated unit shuts down by 0.375 percentage points, from a baseline of 1.2 percentage points, driven by "generators changing relative position along the dispatch order".

2 did not, but unregulated gas units tend to be newer than regulated gas units, so even though the percent retired was similar, the unregulated units were retired earlier.

The magnitude of the difference is economically and statistically significant. The pattern is consistent with owners of regulated units being insulated from traditional economic factors and facing less pressure to exit during unfavorable economic times and, in some cases, having an incentive to continue operating capital-intensive equipment. In the following section, we turn to a hazard model that allows us to continue to compare regulated and unregulated units, while also controlling for unit capacity, plant technology, vintage, and regional economic trends.

Table 3: Average Age at Retirement

	(1) Regulated Units	(2) Unregulated Units	(3) p-value (1) vs (2)
Coal	55	49	0.000
	(n=506)	(n=146)	
Gas	47	39	0.000
	(n=600)	(n=298)	
All	50	42	0.000
	(n=1106)	(n=444)	

Notes: This table reports the average age at retirement in years for regulated and unregulated units. We calculated these ages using our dataset of U.S. electricity generating units that were operating in 2010 and retired between 2011 and 2023. The final column reports p-values from a t-test of the null hypothesis that the ages at retirement are equal for regulated and unregulated units.

Figure 3: Age at Retirement

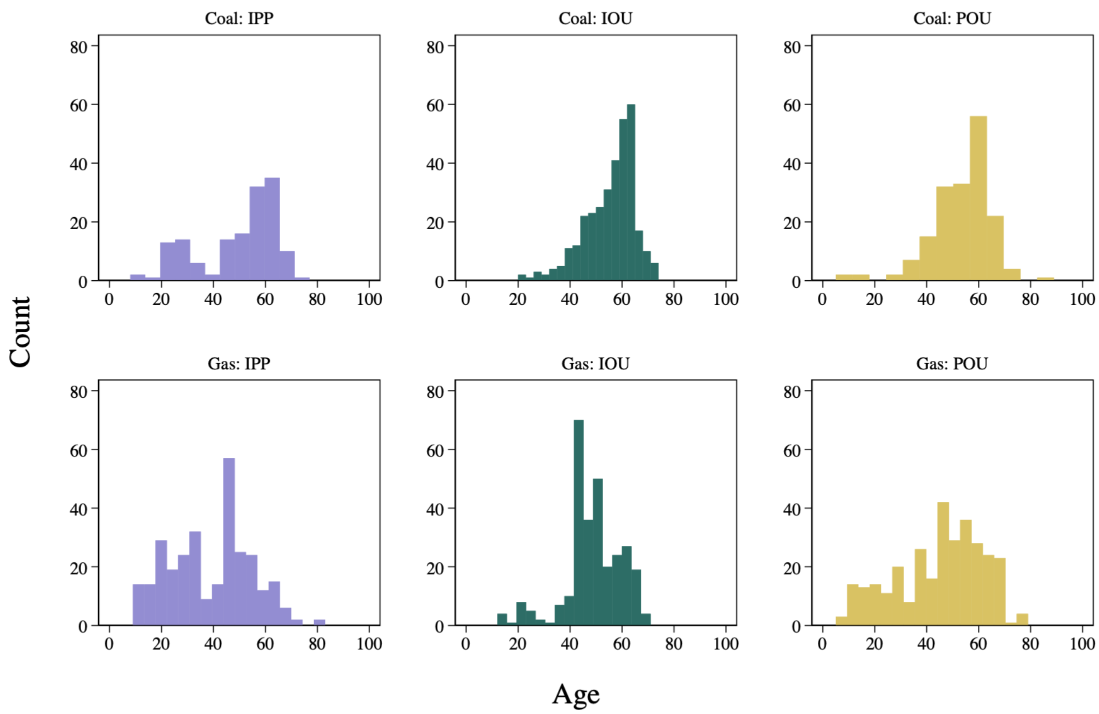

Notes: This figure shows histograms of unit age at retirement by fuel and regulatory status among the units online in 2010 that retire over the period 2011–2023. Across both fuel types, 9% of IOU units retire before age 40, compared to 40% and 26% for IPP and POU, respectively.

4 Empirical Analysis

4.1 Proportional Hazard Model

In this section, we use a proportional (Cox) hazard model to estimate the impact of regulation on the relative risk of unit exit in a given year. In the parlance of traditional survival analysis, a unit's life begins at its opening year, and the unit continues to "survive" until it exits. We use a hazard model rather than, for example, linear regression because this setting is intrinsically about duration. We want to understand how different factors shift the time until exit and to be able to say how regulation type shifts the expected lifetime for generating units.

The hazard model is specified:

$$\lambda(t; x_i) = \exp(x_i \beta) \lambda_0(t), \tag{1}$$

where t denotes the duration of unit i's operating life (in years), and $\lambda_0(t)$ is the baseline hazard function. We estimate hazard models separately for coal and natural gas, which allows the baseline hazard model $\lambda_0(t)$ to vary by fuel type.[9]

This is called a proportional hazard model because covariates x_i shift this baseline hazard function proportionally. Our covariate of interest is *Regulated*, an indicator equal to one for units owned by regulated companies. Here we include among reg-

[9] The covariate vector x_i scales the baseline hazard ratio proportionally by a factor of $\exp(x_i\beta)$, and the ratio of hazard rates between two units i and j can be expressed: $\frac{\lambda(t;x_i)}{\lambda(t;x_j)} = \frac{\lambda_0(t) exp(\beta x_i)}{\lambda_0(t) exp(\beta x_j)} = exp(\beta(x_i - x_j))$. Thus, the relative risk between two units i and j differs by $exp(\beta(x_i - x_j))$ regardless of time.

ulated companies both IOUs and POUs, but later we estimate more flexible models that distinguish between IOUs and POUs.

Among the covariates x_i we also include unit capacity (in MW) and NERC by 5-year fixed effects (2010–2014, 2015–2019, 2020–2023). The fixed effects are included to control for market conditions and other time-varying factors unique to the unit's NERC region. This includes, for example, regional electricity demand and entry of renewables. There are 10 NERC regions (ASCC, FRCC, HICC, MRO, NPCC, RFC, SERC, SPP, TRE, WECC) and 14 years in our study period, so this is a total of 30 fixed effects.[10]

In some specifications, we also include vintage by technology strata. This is introduced not as additional covariates, but by allowing the baseline hazard function to differ by strata, which means that we estimate separate baseline hazard functions for each vintage by technology group, which are then all shifted by the same proportional $\exp(x_i\beta)$ term. For vintage, we use the decade the unit opened. For technology, we use the unit's technology type – prime mover in EIA-860 data. This richer specification allows, for example, combustion gas turbine units built in the 1980s to have a different baseline hazard function than combined cycle units built in the 2000s.

[10] One might have envisioned also controlling explicitly for wholesale electricity prices but, of course, wholesale electricity prices are not available in parts of the United States without wholesale electricity markets. While alternative measures do exist, the NERC by 5-year fixed effects provide a transparent and flexible approach to controlling for wholesale prices and other market conditions. Our results are robust to a variety of sensitivity analyses including an alternative specification with NERC and 1-year fixed effects, not interacted. We would have liked to include even richer fixed effects, for example NERC by 1-year fixed effects, but there is not enough variation in exit within NERC by 1-year groups for this to be estimated. See the online appendix for details.

4.2 Main Results

Table 4 presents our main results. For each of six specifications, we report the hazard ratio and standard error corresponding to *Regulated*, an indicator equal to one for units that are regulated. Columns (1), (3), and (5) control for unit capacity only. Columns (2), (4), and (6) additionally incorporate region by 5-year fixed effects and vintage by technology strata.

Regulated units are much less likely to exit than unregulated units. Across the six specifications, we estimate hazard ratios ranging from 0.43 to 0.60. As usual, a hazard ratio of 1.0 would indicate equally likely, so regulated units exit at approximately half the rate of unregulated units, controlling for observable factors. In column (6), the estimated hazard ratio of 0.55 implies that regulated coal and gas units are 45% less likely to exit than unregulated units.

Standard errors are clustered at the power plant level in all specifications to allow for correlation over time and across units within a plant. Statistical significance is assessed relative to the null hypothesis that the hazard ratio equals 1.0, i.e. that regulated and unregulated units are equally likely to exit. This null is rejected in all six specifications, with p-values below 0.001 in all cases.

The estimated hazard ratios are similar for coal and gas, in both cases indicating lower exit rates for regulated units. The estimated hazard ratios are somewhat larger for coal – 0.59–0.60 – compared to for gas – 0.43–0.53, but the confidence intervals overlap. Thus, the main takeaway from Table 4 is that regulated coal and gas units were both significantly less likely to exit than unregulated units.

Figure 4 shows these results visually. Based on the specification in Table 4 column (5), the figure plots the cumulative retirement probabilities for regulated and unregulated units by age. For example, 60 years after opening 39% of regulated units have exited, compared to 64% of unregulated units. Converting the cumulative probability density functions shown in the figure into their implied probability density functions and computing mean durations across the two groups, we find that the average age at retirement is 8.4 years higher for regulated units compared to unregulated units.

Could this pattern reflect unregulated companies cutting corners in operations and maintenance (O&M)? Possibly, but neither the academic literature nor the popular press emphasizes equipment deterioration as a major driver of plant exit decisions.[11] Furthermore, it is not clear whether unregulated companies have less incentive than regulated companies to invest in O&M. On the one hand, an unregulated company cannot pass O&M costs along to ratepayers. However, on the other hand, unregulated companies have a strong incentive to keep their plants running.[12] Unregulated companies make profit from selling electricity and therefore do not want operational

[11]For the academic literature see Linn and McCormack (2019) and the other papers cited in the introduction, and for the popular press see: "Power Company Losses Some of Its Appetite for Coal," by Eric Lipton, *New York Times*, December 19, 2012. "Utilities Speed Up Closure of Coal-Fired Power Plants," by Katherine Blunt, *Wall Street Journal*, January 9, 2019. "New York's Last Coal-Fired Power Plant Is Closing," by Anne Barnard, *New York Times*, March 20, 2020. "$615,000 a Day: Order to Keep Coal Plant Open Ignites Debate in Michigan," by Joe Barrett and Jennifer Hiller, *Wall Street Journal*, November 3, 2025.

[12]In related work, Hausman (2014) finds that unregulated U.S. nuclear power plants are safer than regulated plants along multiple observable measures. The paper argues that unregulated companies work hard to improve safety because they have "strong incentives to avoid outages". For an unregulated company, "any outage leads to large losses in operating profits", while a regulated company can pass along the cost of an outage to ratepayers. This incentive to avoid outages is particularly strong for nuclear plants due to their low marginal cost of generation, but is also present for coal and gas plants.

problems that lead to outages.

Could this pattern reflect that regulated units were less affected by MATS? Fowlie (2010) finds that regulated companies (both IOUs and POUs) are more likely to invest in capital intensive emission equipment for coal plants, which the paper attributes to the Averch-Johnson effect. With more emission equipment already installed, regulated units might then have been less affected by MATS. This may well be part of the mechanism here, but it is worth emphasizing that MATS applies to coal but not gas, so this mechanism cannot explain the estimates in columns (3) and (4). Moreover, as we mentioned before, previous research finds that MATS played only a small role in explaining recent U.S. coal plant retirements (Linn and McCormack, 2019).

Table 4: Does Regulation Distort Exit Decisions?

	Coal		Gas		All	
	(1)	(2)	(3)	(4)	(5)	(6)
Regulated	0.598***	0.589***	0.432***	0.529***	0.505***	0.554***
	(0.091)	(0.094)	(0.068)	(0.086)	(0.056)	(0.064)
Size (MW)	Yes	Yes	Yes	Yes	Yes	Yes
NERC x 5-Year FE	No	Yes	No	Yes	No	Yes
Vintage x Technology Strata	No	Yes	No	Yes	No	Yes
AIC	7,683	6,085	12,213	8,046	21,943	14,142
Observations	14,073	14,073	65,743	65,743	79,816	79,816
Unique units	1,189	1,189	4,960	4,960	6,149	6,149
Number retirements	652	652	898	898	1550	1550

NOTES: This table reports estimated hazard ratios from a proportional hazard (cox) model described in equation (1). The model is estimated using maximum likelihood. The variable *Regulated* equals one for units owned by either an investor-owned utility or a publicly-owned utility. Columns (2), (4), and (6) include region by 5-year fixed effects for the 10 NERC regions. Columns with vintage by technology strata estimate separate baseline hazard functions for each combination of entry decade and prime mover group (e.g. natural gas turbine built in the 1980s vs natural gas combined cycle built in the 2000s). For column (6), the included strata are vintage by technology by fuel type and the included fixed effects are NERC by 5-year by fuel type. Standard errors are clustered at the power plant level. Statistical significance is assessed relative to the null hypothesis that the hazard ratio equals one.

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

Figure 4: Does Regulation Distort Exit Decisions?

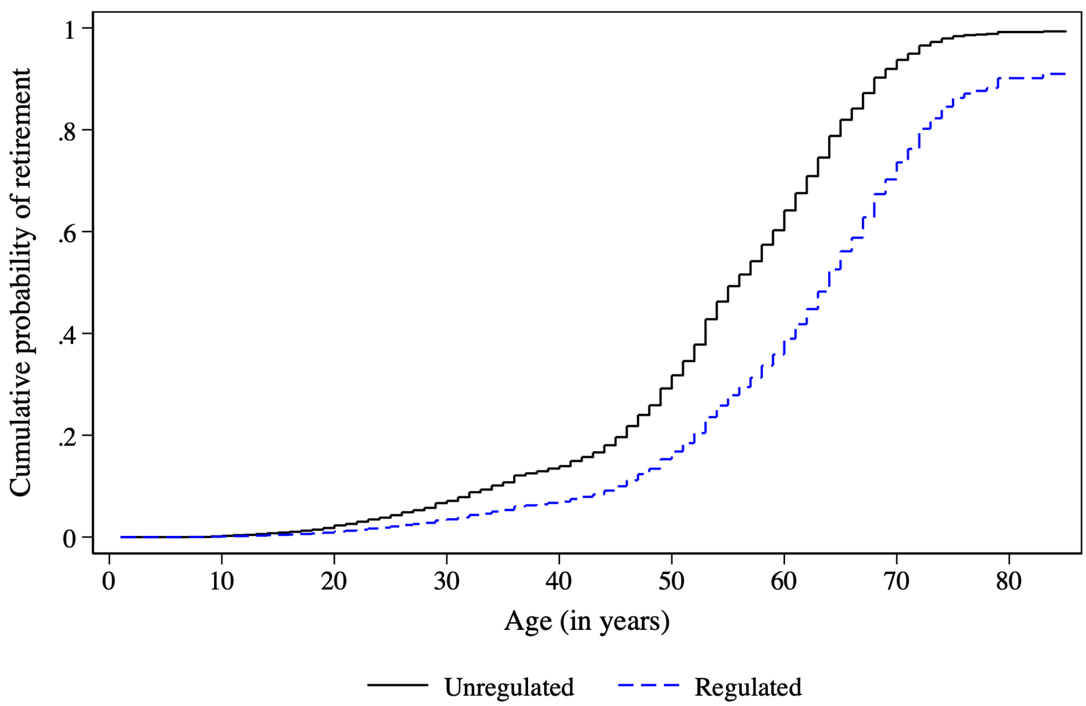

Notes: This figure plots cumulative retirement probabilities for regulated and unregulated units by age. For example, 60 years after opening 39% of regulated units have exited, compared to 64% of unregulated units. The estimates imply that the average age at retirement is 8.4 years higher for regulated units. Cumulative retirement probabilities are estimated using a Breslow estimator to recover the baseline hazard function after estimating the Cox model specification in Table 4 column (5).

4.3 IOU versus POU

We next estimate more flexible models that distinguish between investor-owned utilities (IOUs) and publicly-owned utilities (POUs). The results in the previous subsection combine IOUs and POUs into a single "regulated" category, but here we separate them into two categories.

Table 5 reports the estimated hazard ratios for IOUs and POUs. The structure of the table is identical to our main results in Table 4, except that we replace the variable *Regulated* in all specifications with two variables *Investor Owned Utility* and *Public Owned Utility*, which are indicators for IOU and POU units, respectively. The excluded category continues to be units owned by unregulated companies.

Both IOU and POU units are less likely to exit than unregulated units. In the full specification in column (6), the estimated hazard ratios are 0.62 and 0.44, indicating that IOU and POU units are 38% and 56% less likely to exit than unregulated units, respectively. The estimated hazard ratios are similar across columns, and twelve of the twelve estimated hazard ratios are statistically significant with p-values below 0.001. The table also reports in the last row p-values for tests of equality for the two estimated hazard ratios. The differences are statistically significant in some, but not all cases.

These results are consistent with IOUs and POUs being less exposed to economic conditions than IPPs. For IOUs, the lower exit rate also reflects that, in some cases, these companies have an incentive to continue operating capital-intensive equipment. With POUs, it was unclear ex ante whether these utilities would be more or less exposed to economic conditions than IOUs, but the point estimates suggest that, if anything, POUs are less exposed. We find this somewhat surprising, but could reflect POUs being more influenced by local priorities and preferences. As mentioned earlier, a POU might value, for example, retaining local employment or simply have strong preferences toward maintaining the status quo.

Table 5: Does Regulation Distort Exit Decisions? IOU versus POU

	Coal		Gas		All	
	(1)	(2)	(3)	(4)	(5)	(6)
Investor Owned Utility	0.608***	0.645**	0.584***	0.598***	0.592***	0.620***
	(0.096)	(0.111)	(0.100)	(0.115)	(0.069)	(0.080)
Public Owned Utility	0.550***	0.453***	0.319***	0.429***	0.405***	0.439***
	(0.114)	(0.098)	(0.059)	(0.085)	(0.055)	(0.063)
Size (MW)	Yes	Yes	Yes	Yes	Yes	Yes
NERC x 5-Year FE	No	Yes	No	Yes	No	Yes
Vintage x Technology Strata	No	Yes	No	Yes	No	Yes
AIC	7,636	6,005	12,137	7,961	21,839	13,953
Observations	14,073	14,073	65,743	65,743	79,816	79,816
Unique units	1,189	1,189	4,960	4,960	6,149	6,149
Number retirements	652	652	898	898	1550	1550
P-value: IOU vs POU	0.552	0.042	0.000	0.067	0.001	0.005

NOTES: This table is identical to our main results in Table 4 except that we replace in all specifications the variable *Regulated* with two variables *Investor Owned Utility* and *Public Owned Utility* which are indicators for IOU and POU units, respectively. The table has also been augmented with a new last row which reports the p-value for a Wald test of equality for the two estimated hazard ratios.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

4.4 Market Power

Market power provides a potential alternative explanation for why unregulated units would have exited at higher rates than regulated units. In particular, an unregulated company with a portfolio of power plants might choose to strategically retire certain units in order to increase electricity prices for the remaining units owned by the same company.[13]

In this subsection, we test for market power. We identify the potential to exercise market power using each unregulated company's total 2010 U.S. capacity. From EIA-860, we observe the operator name for each unit using the variable "Utility Name". Given that many parent companies have unique LLCs for different power plants (e.g. Calpine Bethlehem LLC and Calpine Bosque Energy Center LLC; AES Alamitos LLC and AES Beaver Valley), we extract the first word of the operator name and sum 2010 capacity based on that name. For example, we sum capacity for all units with an operator name that starts with the word Calpine, AES, Dynegy, and so on. This generates 481 unique operators for 1882 unregulated units at 574 power plants. See the online appendix for details.

Table 6 reports the results. This table is identical to our main results in Table 4 except that we replace in all specifications the variable *Regulated* with three variables which are indicators for units operated by independent power producers (i.e., unregulated) with different total portfolio sizes.

The estimated hazard ratios are all above 1.0, consistent with our previous results. Notice that for this table, we have flipped the excluded category. Whereas in the previous tables the excluded category was unregulated units, the excluded category now is regulated units. Thus, hazard ratios above 1.0 in this table are equivalent

[13] Myatt (2018) concludes based on a structural oligopoly model that there has been too much coal exit in two major U.S. electricity markets (MISO and PJM), relative to the optimal choices of a hypothetical least-cost planner, interpreting these exits as companies with portfolios of power plants exercising market power.

to hazard ratios less than one in the previous tables, in both cases indicating that regulated units are *less* likely to exit relative to unregulated units.

If market power was driving exit behavior, then we would expect to see larger hazard ratios for large portfolios. Estimates vary between specifications, but, in general, this is not what we find. Medium-sized portfolios tend to have the largest hazard ratios, with smaller hazard ratios for large-sized portfolios. In addition, most of the differences are not statistically significant. The three bottom rows in the table report p-values for tests of equality and, in most cases, the differences are not statistically significant.

It is also notable that the estimated hazard ratios are considerably higher than 1.0 even for smaller operators. This is not what we would expect with market power. The motivation for companies in this setting is that in retiring one unit, this pushes up prices for remaining units. But in this case, these operators are small enough that this strategy becomes unprofitable because there are not enough other units to benefit.

Thus, these results suggest that market power is not driving our main results. Estimated hazard ratios vary widely, but there is no consistent evidence of higher exit rates for larger operators, nor consistent evidence of lower exit rates for smaller operators. Market power may well be a factor in retirement decisions, but it does not seem to be the dominant factor, nor does it seem that market power can explain the patterns observed in Tables 4 and 5.

Table 6: Does Regulation Distort Exit Decisions? Market Power Test

	Coal		Gas		All	
	(1)	(2)	(3)	(4)	(5)	(6)
Unregulated, Operator < 3 GW	2.121**	1.820***	2.420***	1.616**	2.337***	1.662***
	(0.627)	(0.397)	(0.508)	(0.368)	(0.381)	(0.271)
Unregulated, Operator 3–11 GW	1.786***	2.344***	2.404***	2.383***	1.921***	2.342***
	(0.264)	(0.406)	(0.678)	(0.591)	(0.361)	(0.371)
Unregulated, Operator > 11 GW	1.069	1.022	2.081***	1.994***	1.570***	1.576**
	(0.305)	(0.377)	(0.439)	(0.434)	(0.274)	(0.311)
Size (MW)	Yes	Yes	Yes	Yes	Yes	Yes
NERC x 5-Year FE	No	Yes	No	Yes	No	Yes
Vintage x Technology Strata	No	Yes	No	Yes	No	Yes
AIC	7,676	6,089	12,216	8,045	21,935	14,123
Observations	14,073	14,073	65,743	65,743	79,816	79,816
Unique units	1189	1189	4960	4960	6149	6149
Number retirements	652	652	898	898	1550	1550
P-value: Small vs. Medium	0.589	0.337	0.983	0.215	0.406	0.110
P-value: Small vs. Large	0.085	0.167	0.564	0.433	0.076	0.818
P-value: Medium vs Large	0.090	0.031	0.645	0.560	0.398	0.098

NOTES: This table is identical to our main results in Table 4 except we replace in all specifications the variable *Regulated* with three variables which are indicators for units operated by independent power producers (i.e. unregulated) with different total portfolio sizes. The table has also been augmented with three new rows at the bottom which reports p-values for Wald tests of equality for the various estimated hazard ratios.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

4.5 Local Priorities

As discussed earlier, regulated firms are subject to the preferences and priorities of local regulators. In this subsection, we test for one particular form of regulatory capture. We hypothesize that in states with large amounts of coal and gas production, we would expect more political pressure to keep regulated coal and gas plants operating.[14] To test this hypothesis, we compile additional data and identify states

[14]This hypothesis is inspired by Cicala (2015), which compares procurement behavior by U.S. coal plants before and after deregulation. Using difference-in-differences, the paper finds that deregulation reduced the price paid for coal by 12 percent, primarily because unregulated plants became more willing to source coal from out of state. The paper attributes this pattern to regulatory capture in the spirit of Stigler (1971) and Peltzman (1976), with coal producers exerting political pressure on state politicians and regulators.

that are large producers of these inputs. We define "large" as the ten states with the highest cumulative production of coal and gas, respectively, over the period 2011–2023. These large producing states account of 88% and 90% of total cumulative production of coal and gas, respectively during that period.[15]

Table 7 reports the results of this test of regulatory capture. This table is similar to our main results in Table 4 except that we replace in all specifications the variable *Regulated* with two variables that indicate regulated units in states with or without large in-state production of coal and natural gas, respectively. Our hypothesis is about coal producers influencing coal plants and natural gas producers influencing gas plants, so we do not estimate a "stacked" specification for all units.

We find no evidence of regulatory capture along this dimension. Regulated units are again less likely to exit than unregulated units. The estimated hazard ratios are all well below one and are all highly statistically significant. However, we see no difference in exit behavior for units located in states with or without large in-state production.[16] Perhaps this is unsurprising for natural gas. Cicala (2015) finds little evidence that deregulation changed procurement behavior for natural gas, which the paper attributes to gas being a "homogeneous commodity traded in regional markets with transparent prices". If natural gas markets are sufficiently transparent

[15]These calculations were made using state-by-year data from U.S. Department of Energy, Energy Information Administration, State Energy Data System. `https://www.eia.gov/state/seds/`. The top ten states for coal are Wyoming, West Virginia, Kentucky, Pennsylvania, Illinois, Indiana, Montana, Texas, North Dakota, and Colorado. The top ten states for gas are Texas, Pennsylvania, Louisiana, Oklahoma, Colorado, West Virginia, New Mexico, Wyoming, Ohio, and North Dakota.

[16]We also conducted an alternative test based on county-level coal/gas employment rather than state-level coal/gas production, again finding no evidence of regulatory capture. See the online appendix for details.

and regional, then it may simply not be worth lobbying to keep in-state gas plants operating.

In contrast, coal is more heterogeneous. Local coal plants are often tuned to the specific characteristics of locally produced coal, and there is more asymmetric information with coal contracts (Joskow, 1987; Cicala, 2015). Given this, it is somewhat surprising that units in states with large coal production do not exhibit lower exit rates. However, there could be other counteracting differences between states with and without large coal production or that the magnitude of the effect is too small to detect. It is also worth emphasizing that there are many forms that local priorities and preferences could take, and these results should not be interpreted as providing evidence against all forms of local priorities and preferences.

Table 7: Does Regulation Distort Exit Decisions? Local Priorities

	Coal		Gas	
	(1)	(2)	(3)	(4)
Regulated x Large coal production	0.571***	0.538***		
	(0.116)	(0.107)		
Regulated x Small coal production	0.608***	0.616***		
	(0.095)	(0.105)		
Regulated x Large gas production			0.396***	0.527***
			(0.072)	(0.102)
Regulated x Small gas production			0.444***	0.530***
			(0.075)	(0.093)
Size (MW)	Yes	Yes	Yes	Yes
NERC x 5-Year FE	No	Yes	No	Yes
Vintage x Technology Strata	No	Yes	No	Yes
AIC	7,684	6,093	12,214	8,048
Observations	14,073	14,073	65,743	65,743
Unique units	1,189	1,189	4,960	4,960
Number retirements	652	652	898	898

NOTES: This table is identical to our main results in Table 4 except we replace the variable *Regulated* with two variables that indicate regulated units in states with or without large in-state production of coal and natural gas, respectively. Large producing states are the ten states with the highest cumulative production of coal and gas, respectively, over the period 2011–2023.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

5 Conclusion

U.S. power companies made the decision to retire more than 180 gigawatts of coal and natural gas between 2011 and 2023. In this paper, we focus on the role that regulation played in these decisions. Using generator-level data, we show that units owned by regulated companies were 45% less likely to exit than units owned by unregulated companies. This difference remains after controlling for unit characteristics and region-by-time fixed effects for the ten NERC regions, is large and statistically significant for both coal and natural gas plants, and holds for both investor-owned utilities and publicly-owned utilities.

We argue that this behavior reflects incentives. Unregulated companies were exposed to unfavorable market conditions during this period, so many decided to exit. Regulated companies, in contrast, were insulated from economic conditions and therefore less likely to exit. For investor-owned utilities, there is also a well-known regulatory distortion which in some cases creates a strong incentive to continue operating capital-intensive equipment.

Since 2024, the fortunes for U.S. power companies have changed dramatically. The increased demand for electricity from data centers shifted the market overnight from surplus to shortage and left companies scrambling to add new capacity, and, in a few cases, even to bring back capacity that had been recently closed. One of the implications of this dramatic reversal is that units that did not exit 2011–2023 are now more likely to remain open for the foreseeable future.

Moving forward, incentives will continue to matter. More exposed to the positive economic conditions, we would expect unregulated units to be *less* likely to exit than regulated units. There will also be pressure to keep regulated units operating, but it remains to be seen whether this pressure will be as strong as the profit motive for unregulated units. The broader point is that the effect of regulation on exit decisions can be either positive or negative, but that we should expect regulated companies to be largely insulated from whatever current economic conditions may be.

References

Averch, Harvey and Leland L Johnson, "Behavior of the Firm Under Regulatory Constraint," *American Economic Review*, 1962, *52* (5), 1052–1069.

Borenstein, Severin, "The Trouble with Electricity Markets: Understanding California's Restructuring Disaster," *Journal of Economic Perspectives*, 2002, *16* (1), 191–211.

_ **and James Bushnell**, "The U.S. Electricity Industry After 20 Years of Restructuring," *Annual Review of Economics*, 2015, *7* (1), 437–463.

_ **, James B Bushnell, and Frank A Wolak**, "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market," *American Economic Review*, 2002, *92* (5), 1376–1405.

Borrero, Miguel, Gautam Gowrisankaran, and Ashley Langer, "Ramping Costs and Coal Generator Exit," *Working Paper*, 2023.

Bushnell, James B, Erin T Mansur, and Celeste Saravia, "Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured U.S. Electricity Markets," *American Economic Review*, 2008, *98* (1), 237–266.

Chan, H Ron, Harrison Fell, Ian Lange, and Shanjun Li, "Efficiency and Environmental Impacts of Electricity Restructuring on Coal-Fired Power Plants," *Journal of Environmental Economics and Management*, 2017, *81*, 1–18.

Cicala, Steve, "When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation," *American Economic Review*, 2015, *105* (1), 411–444.

_ , "Imperfect Markets Versus Imperfect Regulation in US Electricity Generation," *American Economic Review*, 2022, *112* (2), 409–441.

_ , "Restructuring the Rate Base," *AEA Papers and Proceedings*, 2025, *115*, 363–368.

Coglianese, John, Todd D Gerarden, and James H Stock, "The Effects of Fuel Prices, Environmental Regulations, and Other Factors on U.S. Coal Production, 2008-2016," *Energy Journal*, 2020, *41* (1), 55–82.

Cullen, Joseph A and Erin T Mansur, "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution," *American Economic Journal: Economic Policy*, 2017, *9* (3), 106–133.

Davis, Lucas W and Catherine Wolfram, "Deregulation, Consolidation, and Efficiency: Evidence from US Nuclear Power," *American Economic Journal: Applied Economics*, 2012, *4* (4), 194–225.

Davis, Rebecca J, J Scott Holladay, and Charles Sims, "Coal-Fired Power Plant Retirements in the United States," *NBER Environmental and Energy Policy and the Economy*, 2022, *3* (1), 4–36.

Demirer, Mert and Ömer Karaduman, "Do Mergers and Acquisitions Improve Efficiency? Evidence from Power Plants," *forthcoming at Journal of Political Economy*, 2025.

Duff, Cameron, "Three Essays on the Economics of Electricity." Unpublished doctoral dissertation, University of Colorado Boulder 2020.

Dunkle Werner, Karl and Stephen Jarvis, "Rate of Return Regulation Revisited," *Energy Institute Working Paper, UC Berkeley*, 2025.

EIA, U.S. Energy Information Administration, "Electricity Data Browser," 2024.

EIA, U.S. Energy Information Administration U.S., "Hydraulically Fractured Horizontal Wells Account for Most New Oil and Natural Gas Wells," January 2018.

Fabrizio, Kira R, Nancy L Rose, and Catherine D Wolfram, "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency," *American Economic Review*, 2007, *97* (4), 1250–1277.

Fell, Harrison and Daniel T Kaffine, "The Fall of Coal: Joint Impacts of Fuel Prices and Renewables on Generation and Emissions," *American Economic Journal: Economic Policy*, 2018, *10* (2), 90–116.

Fowlie, Meredith, "Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement," *American Economic Review*, 2010, *100* (3), 837–869.

Gowrisankaran, Gautam, Ashley Langer, and Mar Reguant, "Energy Transitions in Regulated Markets," *American Economic Review*, forthcoming.

_ , _ , and Wendan Zhang, "Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards," *Journal of Political Economy*, 2025, *133* (6), 1757–1795.

Hausman, Catherine, "Corporate Incentives and Nuclear Safety," *American Economic Journal: Economic Policy*, 2014, *6* (3), 178–206.

_ and Ryan Kellogg, "Welfare and Distributional Implications of Shale Gas," *Brookings Papers on Economic Activity*, 2015, pp. 71–125.

Holland, Stephen P, Erin T Mansur, Nicholas Z Muller, and Andrew J Yates, "Decompositions and Policy Consequences of an Extraordinary Decline in Air Pollution from Electricity Generation," *American Economic Journal: Economic Policy*, 2020, *12* (4), 244–274.

Joskow, Paul L, "The Determination of the Allowed Rate of Return in a Formal Regulatory Hearing," *Bell Journal of Economics and Management Science*, 1972, pp. 632–644.

_ , "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets," *American Economic Review*, 1987, *77* (1), 168–185.

_ **and Richard Schmalensee**, "Markets for Power: An Analysis of Electrical Utility Deregulation," *MIT Press Books*, 1988.

Linn, Joshua and Kristen McCormack, "The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions," *RAND Journal of Economics*, 2019, *50* (4), 733–767.

Mansur, Erin T and Matthew White, "Market Organization and Efficiency in Electricity Markets," *Working Paper*, 2012.

Myatt, James, "Market Power and Long-Run Technology Choice in the U.S. Electricity Industry," *Northwestern University Department of Economics Working Paper*, 2018.

Peltzman, Sam, "Toward a More General Theory of Regulation," *Journal of Law and Economics*, 1976, *19* (2), 211–240.

Peng, Jing, Jesse Buchsbaum, Catherine Hausman, and Johanna L Mathieu, "Power System Decarbonization: A Comparison Between Carbon Taxes and Forcing Coal Power Plant Retirements," *IEEE Transactions on Sustainable Energy*, 2024.

Stigler, George J, "The Theory of Economic Regulation," *Bell Journal of Economics and Management Science*, 1971, *2* (1), 3–21.

Warner, Cody, Duncan Callaway, and Meredith Fowlie, "Risk-Cost Tradeoffs in Power Sector Wildfire Prevention," *Energy Institute Working Paper, UC Berkeley*, 2024.

White, Matthew W, "Power Struggles: Explaining Deregulatory Reforms in Electricity Markets," *Brookings Papers on Economic Activity*, 1996, *1996*, 201–267.

Wolfram, Catherine D, "Measuring Duopoly Power in the British Electricity Spot Market," *American Economic Review*, 1999, *89* (4), 805–826.

Online Appendix

Appendix Table 1: Balance of Covariates

	(1) Regulated Units	(2) Unregulated Units	(3) P-value (1) vs. (2)
A. Coal			
Size (MW)	279	278	0.961
	(273.8)	(266.6)	
Year Opened (MW)	1968	1972	0.000
	(14)	(16)	
B. Gas			
Size (MW)	81	112	0.000
	(97.8)	(111.7)	
Year Opened (MW)	1987	1993	0.000
	(19)	(14)	
Combined Cycle	0.24	0.47	0.000
Combustion Gas Turbine	0.43	0.36	0.000
Combustion Steam Turbine	0.12	0.07	0.000
Other	0.22	0.11	0.000

Notes: This table reports summary statistics for our dataset of electricity generating units by fuel type and regulatory status. Panel (A) describes all coal units in operation in 2010 and panel (B) describes all gas units. Mean capacity and opening year is reported, with standard deviations in parenthesis. Unit technology shares are reported for gas units; "Other" includes internal combustion, combined cycle units that operate with a single shaft, and units with energy storage. The final column reports p-values from a t-test of the null hypothesis that values in the two columns are equal for regulated and unregulated units.

Appendix Table 2: Does Regulation Distort Exit Decisions? Alternative Specification Excluding Units Smaller Than 50MW

	Coal		Gas		All	
	(1)	(2)	(3)	(4)	(5)	(6)
Regulated	0.644***	0.676**	0.520***	0.390***	0.625***	0.557***
	(0.108)	(0.111)	(0.118)	(0.105)	(0.084)	(0.079)
Size (MW)	Yes	Yes	Yes	Yes	Yes	Yes
NERC x 5-Year FE	No	Yes	No	Yes	No	Yes
Vintage x Technology Strata	No	Yes	No	Yes	No	Yes
AIC	5,844	4,632	2,827	2,047	9,620	6,687
Observations	11,524	11,524	36,797	36,797	48,321	48,321
Unique units	958	958	2,709	2,709	3,667	3,667
Number retirements	518	518	255	255	773	773

NOTES: This table reports estimates from an alternative specification which is identical to our main results in Table 4, except we exclude units smaller than 50MW. A nice feature of EIA-860 is that it includes all U.S. power plants over one megawatt of capacity. Thus, these data include many units that are quite small, just a fraction of the capacity of, for example, the average coal and gas unit in our sample. This table shows that the estimates are quite similar when we exclude smaller units. The estimated hazard ratios are modestly larger for coal and modestly smaller for gas with controls, but overall quite similar both in magnitude and statistical significance.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix Table 3: Does Regulation Distort Exit Decisions? IOU vs POU Alternative Specification Excluding Units Smaller Than 50MW

	Coal		Gas		All	
	(1)	(2)	(3)	(4)	(5)	(6)
Investor Owned Utility	0.645**	0.678**	0.523***	0.429***	0.623***	0.574***
	(0.110)	(0.121)	(0.124)	(0.123)	(0.086)	(0.088)
Public Owned Utility	0.599*	0.602**	0.496**	0.270***	0.606**	0.450***
	(0.166)	(0.150)	(0.160)	(0.101)	(0.131)	(0.093)
Size (MW)	Yes	Yes	Yes	Yes	Yes	Yes
NERC x 5-Year	No	Yes	No	Yes	No	Yes
Vintage x Technology	No	Yes	No	Yes	No	Yes
AIC	5,801	4,564	2,816	2,030	9,567	6,596
Observations	11,524	11,524	36,797	36,797	48,321	48,321
Unique units	958	958	2,709	2,709	3,667	3,667
Retirements	518	518	255	255	773	773
P-value: IOU vs POU	0.764	0.592	0.850	0.119	0.888	0.172

NOTES: This table reports estimates from an alternative specification which is identical to the specification used for Table 5, except we exclude units smaller than 50MW.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix Table 4: Does Regulation Distort Exit Decisions? Alternative Specification Using Richer Time Fixed Effects

	Coal		Gas		All	
	(1)	(2)	(3)	(4)	(5)	(6)
Regulated	0.598***	0.588***	0.432***	0.525***	0.505***	0.556***
	(0.091)	(0.098)	(0.068)	(0.086)	(0.056)	(0.066)
Size (MW)	Yes	Yes	Yes	Yes	Yes	Yes
NERC FE	No	Yes	No	Yes	No	Yes
Year FE	No	Yes	No	Yes	No	Yes
Vintage x Technology Strata	No	Yes	No	Yes	No	Yes
AIC	7,683	5,993	12,213	7,936	21,943	13,940
Observations	14,073	14,073	65,743	65,743	79,816	79,816
Unique units	1,189	1,189	4,960	4,960	6,149	6,149
Number retirements	652	652	898	898	1550	1550

NOTES: This table reports estimates from an alternative specification which is identical to our main results in Table 4, except that NERC by 5-year fixed effects for the ten NERC regions are replaced with separate (non-interacted) NERC fixed effects and annual (1-year) fixed effects. This alternative specification allows for more richness over time, for example 2021 and 2022 have separate fixed effects, but these two variables are not interacted, so this specification allows for less flexibility across NERC regions. The estimates in this alternative specification are very similar to the main results, indicating that the results are not unduly sensitive to this choice of fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix Table 5: Does Regulation Distort Exit Decisions? Sample Restricted to Units with Aligned Sector and Entity Type Designations

	Coal		Gas		All	
	(1)	(2)	(3)	(4)	(5)	(6)
Regulated	0.592***	0.545***	0.405***	0.487***	0.481***	0.509***
	(0.093)	(0.093)	(0.065)	(0.083)	(0.054)	(0.062)
Size (MW)	Yes	Yes	Yes	Yes	Yes	Yes
NERC x 5-Year FE	No	Yes	No	Yes	No	Yes
Vintage x Technology Strata	No	Yes	No	Yes	No	Yes
AIC	6,992	5,509	11,622	7,633	20,532	13,143
Observations	13,306	13,306	60,451	60,451	73,757	73,757
Unique units	1,117	1,117	4,575	4,575	5,692	5,692
Number retirements	600	600	862	862	1462	1462

NOTES: This table reports estimates from an alternative specification which is identical to our main results in Table 4, except we exclude units for which there is some disagreement in the EIA-860 about regulation type. In our main results, we measure regulation type using the response to the "entity type" question. This question provides the most complete information, including the distinction between IOU and POU. But there is also a "sector type" question which does not include as much detail but does distinguish between regulated and unregulated companies. In 7% of cases, there is disagreement between the two questions, and in this alternative specification, we exclude those observations. The results are very similar and, if anything, indicate somewhat stronger differences in exit behavior between regulated and unregulated companies.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix Table 6: Does Regulation Distort Exit Decisions? Cross Sectional Model

	Coal			Gas			All		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Regulated	-0.124**	-0.132***		-0.0260	-0.0619***		0.0175	-0.0733***	
	(0.0486)	(0.0474)		(0.0252)	(0.0214)		(0.0245)	(0.0195)	
IOU			-0.0856*			-0.0175			-0.0312
			(0.0500)			(0.0253)			(0.0227)
POU			-0.204***			-0.111***			-0.123***
			(0.0612)			(0.0247)			(0.0229)
Size (MW) Quadratic	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age Quadratic	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
NERC FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Vintage FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Technology FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Dependent var. mean	0.548	0.548	0.548	0.181	0.181	0.181	0.252	0.252	0.252
R-squared	0.112	0.243	0.252	0.049	0.274	0.282	0.005	0.341	0.347
Observations	1,189	1,189	1,189	4,960	4,960	4,960	6,149	6,149	6,149

NOTES: This table reports the estimates from a cross-sectional model of retirement estimated using least squares. For this regression, we collapse the data to a single observation per generating unit. The outcome variable equals one if the unit was retired at any point during 2011–2023 and zero otherwise. Controls are similar to Tables 4 and 5, with some modifications as time-varying fixed effects cannot be estimated in the cross section. It is worth emphasizing that this cross-sectional model is poorly-suited to our setting for several reasons. Perhaps most importantly, the cross-sectional model throws out a considerable amount of information used in the hazard model, such as the timing of retirements. For example, it treats a retirement in 2013 the same as a retirement in 2023. In addition, the cross-sectional model does not allow for time-varying covariates so we cannot include region-by-time fixed effects to capture changes over time in market conditions. As we explain in the paper, the hazard model is a better match because our research question is intrinsically about duration. We want to study the impact of regulation on time until exit and estimate the change in generating units' expected lifetime due to its regulation type. Despite these significant limitations, the results from the cross-sectional model are broadly similar. In column (2), regulated coal units are 13.2 percentage points less likely to exit, compared to a base of 54.8%, while in column (5) regulated gas units are 6.1 percentage points less likely to exit, compared to a base of 18.1%. Thus, in both cases, the probability of exit decreases by about one-third. Both effects are statistically significant at the 1% level. In columns (3) and (6), the point estimates for IOU are negative but smaller and less statistically significant, and the effects for POU are negative, large, and statistically significant.

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

Appendix Table 7: Large Operators

Operator	Capacity (GW)	Coal Units	Gas Units
GenOn	24.06	38	84
Luminant	17.52	11	40
Calpine	16.19	2	104
NRG	15.52	7	93
PSEG	9.47	4	71
Dynegy	9.06	12	32
Midwest	8.33	15	5
AES	7.20	11	18
Ameren	5.42	13	10
Tenaska	4.69		20
New	3.84		24
CAMS	3.66		23
Dominion	3.38	5	9
PPL	3.23	5	17
Louisiana	2.57	3	13
TC	2.55		21
Gila	2.48		12
TPF	2.47		25
Union	2.43		12
Constellation	2.40	6	1

NOTES: This table lists the top 20 operators by 2010 capacity of unregulated units. Unique operators are identified from the EIA-860 data using the first word of the identified operator name.

Appendix Table 8: Does Regulation Distort Exit Decisions? County Coal and Gas Employment Per Capita

	Coal				Gas			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Regulated x Large Coal employment (Top 20)	0.567**	0.856						
	(0.148)	(0.267)						
Regulated x Small Coal employment (Outside Top 20)	0.601***	0.572***						
	(0.093)	(0.092)						
Regulated x Large Coal employment (Top 50)			0.588**	0.796				
			(0.144)	(0.215)				
Regulated x Small Coal employment (Outside Top 50)			0.599***	0.567***				
			(0.093)	(0.091)				
Regulated x Large Gas employment (Top 20)					0.788	0.854		
					(0.447)	(0.306)		
Regulated x Small Gas employment (Outside Top 20)					0.428***	0.523***		
					(0.068)	(0.085)		
Regulated x Large Gas employment (Top 50)							0.416***	0.473**
							(0.104)	(0.146)
Regulated x Small Gas employment (Outside Top 50)							0.433***	0.532***
							(0.069)	(0.088)
Size (MW)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NERC x 5-Year FE	No	Yes	No	Yes	No	Yes	No	Yes
Vintage x Technology Strata	No	Yes	No	Yes	No	Yes	No	Yes
AIC	7,685	6,090	7,685	6,082	12,213	8,047	12,215	8,048
Observations	14,073	14,073	14,073	14,073	65,743	65,743	65,743	65,743
Unique units	1,189	1,189	1,189	1,189	4,960	4,960	4,960	4,960
Number retirements	652	652	652	652	898	898	898	898
P-value Large vs. Small	0.813	0.143	0.932	0.156	0.278	0.158	0.867	0.681

NOTES: Another mechanism generating local priorities and preferences could be local employment in the coal and gas sectors. Local regulators could be especially interested in keeping power plants online in counties with a high level of coal and gas employment per capita. This table is identical to our main results in Table 4 except we replace the variable *Regulated* with indicators for regulated units located in counties ranked in the top 20 or 50 coal or gas employment per capita, respectively, as of 2010. Overall, we find no evidence that county-level coal or gas employment matters for exit rates. The estimated hazard ratios are consistently well below 1.0, but not statistically different between counties with and without high levels of coal/gas employment. This is again a bit surprising because we might have expected local employment to generate political pressure to keep plants open. That said, it could be that there are other offsetting differences between counties or that the magnitude of the effect is too small to detect. We define coal employment as including both coal mining (NAICS 2121) and fossil fuel generation (NAICS 221112). We define gas employment as including both oil and gas extraction (NAICS 2111) and fossil fuel generation (NAICS 221112). We use publicly-available data files from the Quarterly Census of Employment and Wages (QCEW, https://www.bls.gov/cew/downloadable-data-files.htm), which suppresses data for some counties, for example, when there is a single dominant firm that accounts for more than 80% of employment. We use 2010 county population from the Census to calculate county-level coal and gas employment per capita.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.