Competitive markets and capital investment: Evidence from US power plants

Madeline Yozwiak*

September 4, 2024

Abstract: To mitigate climate change, the US will need to invest billions in new power plants. However, cost-of-service regulation could lead electric utilities to "over-" invest, increasing costs for consumers. I evaluate if competition can achieve better societal outcomes by studying a set of power plants exposed to both regulated and competitive market structures. I find competitive reforms led to a large decrease in investment among power plants that use fossil fuels. Deregulated power plants did not operate less efficiently, suggesting the change increased welfare. The results imply that competitive reforms could be used to lower the infrastructure costs of decarbonization.

JEL Codes: L51, L94, Q48

*Indiana University Bloomington; 1315 East 10th Street, Bloomington, IN 47405; myozwiak@iu.edu.

Mitigating climate change relies heavily upon the electricity system. Simulations suggest that to decarbonize the economy, the electricity sector needs to power many activities that today use fossil fuels directly, such as transportation or heating and cooking, and that this power needs to be supplied by renewable resources (Jacobson et al., 2017). To meet the higher level of demand for electricity with carbon-free generation will require a significant amount of new power plants to be built. For example, the National Renewable Energy Lab recently estimated that, to reach 100% clean energy by 2035, total power plant capacity in the US would need to more than double (Denholm et al., 2022). This could require on the order of $1.7 trillion in capital investment over the next decade (Phadke et al., 2020).

One factor that could affect the total cost to decarbonize is market structure. In many US states, the electricity sector remains fully regulated, and utility commissions determine the prices for vertically integrated utilities, who are granted a local monopoly over their given service territory. In other states, electricity is partially or fully deregulated, with competitive reforms introduced to part or all of the functions needed to generate and deliver power to customers. The heterogeneity creates different incentives for utility investment in infrastructure. Economists have long theorized that the way states regulate utilities can distort the companies' investment decisions, biasing production towards capital (Averch and Johnson, 1962). In the next decade, a natural concern is that regulation may inflate utilities' investment in a renewable transition beyond the level that may be socially optimal. Given the high baseline estimates of how much it might cost to decarbonize, any additional increase in costs risks undermining the already tenuous public support for change and reducing the net social benefit realized.

Reforms passed during the 1990s and early 2000s create an opportunity to study how electricity regulation affects investment and to evaluate if a more competitive market structure can achieve better societal outcomes. The most comprehensive competitive reforms occurred within electricity generation: federal policy enabled the creation of open wholesale power markets (Cicala, 2022), and many states removed generation as a regulated function, such that investment by utilities in power plants would not be "recoverable" through higher, regulated prices. Despite a formidable body of economic research on electricity restructuring, the effect of reform on capital investment has been largely unexplored (Borenstein and Bushnell, 2015; Bushnell, Mansur and Novan, 2017). The gap is especially notable because, prior to reform, scholars anticipated that competition would yield the largest benefit through reductions to generation investment.[1] A single paper recently found that the transition to competitive markets may have led to a disproportionate increase in entry among new power plants (Hill, 2021). However, the analysis focused on the short-term effect of changing market structures. Whether there have been long-term changes to investment in electricity generation, as a result of competitive reforms, has not been established.

[1] As Paul Joskow stated in 1997: "The most important opportunities for cost savings are associated with long-run investments in generation capacity" (Joskow (1997), at p. 125).

In this paper, I quantify how competition affects investment in electricity generation. To do so, I take advantage of a natural experiment that exposed a select set of power plants to both regulated and competitive market structures. In the late 1990s, investor-owned utilities sold or transferred hundreds of power plants to unregulated entities. These 'divestitures' were largely prompted by policymakers in reform states who sought to preempt a single company from exerting undue market power in newly-formed wholesale electricity markets. Once a plant was divested, competitive markets defined the economic incentives for new investment, rather than regulation. After an asset was sold or transferred to an unregulated company, the primary way for the power plant to earn revenue was to bid and sell its generation into wholesale exchanges. By studying how a power plant evolves after it was divested, it's thus possible to identify the effect of competitive reforms in electricity generation on capital investment, relative to a regulated counter-factual.

Using a difference-in-difference design, I find that competition led to large and robust decrease in the level of capital investment a power plant receives. To overcome that capital investment is unobserved after a plant is deregulated, I use the size or 'capacity' of a power plant as a proxy, which will become larger when capital is invested to either uprate existing generators or to add a new generator, and vice versa. Capacity acts as a conservative proxy, because capital can be spent on the plant in ways that is not reflected in a larger size (e.g., pollution controls). I focus the analysis on a sample of 800 power plants that burn fossil fuels–such as coal, oil, and natural gas–because they are more likely to be modular than other generating technologies. After divestment, I show that a fossil power plant is 9.1% smaller, on average across years, relative to comparable plants that remain regulated. This is equivalent to a decrease of 47 megawatts (MW) from the mean plant size of 522 MW. The reduction in capacity translates to about $16.4 billion in avoided investment, each year, across the cohort of 273 divested power plants, if we assume that additions would have been combined cycle generators.

I show the decrease in investment was likely welfare-enhancing by incorporating data on how power plants operated. A new generator can change the operating characteristics of a power plant, in addition to its size, with associated social benefit. Of particular interest in this setting is the use of natural gas; power plants that remain regulated become larger over time primarily by adding new natural gas generators, which could increase the plants' operating efficiencies or reduce their carbon dioxide emissions through coal-to-gas switching. However, I do not find any evidence of degradation in how divested plants operated as a result of their lower level of investment. While fossil power plants do generate less frequently, after divestment, there is no observable change in their rate of generation, the efficiency with which they use fuel, nor the carbon intensity of their fuel mix, relative to regulated power plants. This suggests the considerable savings for consumers in avoided capital investment is likely to be a net social benefit.

Finally, I provide evidence that the treatment effect is likely to be externally valid for investment in other power plant types, such as solar or wind. The decrease in investment among fossil power plants is identified, in part, on counter-factual additions among regulated plants. Using financial data for utilities which own these control plants, I show that the higher level of investment coincided with an increase in firms' production costs–the exact conditions under which theory predicts the Averch-Johnson effect (or, a bias towards capital investment) is most likely emerge (Joskow, 1974). This indicates that exposure to competitive markets reduced costs, in part, by removing a distortion to investment created by cost-of-service regulation. Because the investment incentive is not limited to one technology, it implies the treatment effect is not specific to fossil power plants. In other words, regulation can lead to "over-" investment in future renewable power plants, relative to a competitive market, as it did for past investment in fossil plants.

The paper makes two primary contributions to the literature. First, research on electricity restructuring in the US has documented that competitive reforms led to efficiency gains in how power plants operate (Bushnell and Wolfram, 2005; Fabrizio, Rose and Wolfram, 2007; Davis and Wolfram, 2012; Cicala, 2015, 2022). This analysis explores if there are additional effects of competitive reforms within capital investment decisions. I show that savings to consumers from avoided investment are likely an order of magnitude larger than operational gains, aligning with researchers' expectations prior to reform. This expands our understanding of the full benefits of competitive markets in electricity, both in the US (Borenstein and Bushnell, 2015; Bushnell, Mansur and Novan, 2017) and abroad (Pepermans, 2019).

Second, the literature on decarbonization has primarily explored the ability of incentive-based policy mechanisms, such as carbon taxes or clean energy standards, to encourage efficient investment in clean energy and achieve zero-carbon emissions in the power sector (Goulder, Hafstead and Williams, 2016; Kellogg, 2020; Stock and Stuart, 2021; Borenstein and Kellogg, 2023; Bistline et al., 2024). This paper demonstrates that existing regulation will also determine investment into power plants, even in the absence of additional environmental policy. The reduced-form approach I use also complements recent work by Gowrisankaran, Langer and Reguant (2024), who simulate the effect of regulation on utility investment in natural gas and coal using structural estimation. Overall, the results indicate that differences in electricity regulation across the US could create regional disparities in the cost-effectiveness of a renewable transition. Given this, competitive reform in electricity generation warrants further study among policy tools aimed at decarbonizing the economy.

The remainder of the paper proceeds as follows. Section 1 details the empirical strategy and provides the intuition for the difference-in-difference design.[2] Section 2 describes the data used. Section 3 provides the main empirical results and estimates the effect of

[2]Given the volume of prior work on restructuring, I do not provide a detailed policy background in the paper. Interested readers are referred to Borenstein and Bushnell (2015) and Bushnell, Mansur and Novan (2017) for two excellent, comprehensive, and recent introductions to the literature.

divestment on power plant capacity. Section 4 explores how divestment may have affected power plant operation and what this tells us about the possible welfare implications of a reduction in investment. Section 5 introduces the intuition and evidence for how regulation affected counter-factual investment in control plants, in order to inform if the treatment effects identified on fossil plants are likely to be externally valid for other generating technologies, like renewables. Section 6 summarizes and concludes.

1 Empirical Strategy

I take advantage of a natural experiment–the divestment of hundreds of power plants during the late 1990s–to identify the effect of competition on capital investment in electricity generation. "Divestment" refers to the sale or transfer of a power plant from an investor-owned utility, whose revenues are determined by cost-of-service regulation, to an unregulated entity, whose revenues are determined by the sale of generation within competitive wholesale markets.[3] Most divestments were prompted by broader market reform that occurred at the turn of the century; many states that chose to deregulate electricity generation also required or encouraged utilities to sell their power plants, in order to reduce the companies' ability to exert undue market power in nascent competitive exchanges (Borenstein and Bushnell, 2015; Bushnell, Mansur and Novan, 2017). Using a difference-in-difference design, I trace how the level of investment into divested power plants changed, once exposed to competitive markets.

The divestiture of US power plants has been widely used in the literature to study how competitive markets affect a range of outcomes, including market power (Mansur, 2007), investment crowd-out (Ishii and Yan, 2008), fuel procurement (Cicala, 2015), and the operating efficiency of power plants (Bushnell and Wolfram, 2005; Davis and Wolfram, 2012) and firms (Kwoka, Pollitt and Sergici, 2010). As an identification strategy, it has several strengths. First, divestment allows researchers to observe the same plant under both regulated and competitive market structures. This strengthens the argument for causality; researchers can control for unobserved, unit-specific factors that are time-invariant and would otherwise confound results, in a way that is difficult once markets are established. In addition, the sheer number of plants that were divested, as part of reform, has not been replicated since, and this increases the statistical power of regressions to detect changes. Third, selection effects among divested plants are less likely to be prominent because most divestitures were prompted by state policy, rather than voluntary sales among owners.[4] Finally, and specific to this setting, divestiture allows us to shorten the time-horizon on which medium- to long-run effects on investment could normally be observed. New power

[3] The owner can also sign a power purchase agreement with another party, for a fixed offtake of generation; this would enable a fixed revenue stream, separate from the market.

[4] The states which adopted restructuring legislation first tended to be those with higher retail electricity prices under regulation (Andrews, 2000; Kwoka, 2008; Borenstein and Bushnell, 2015). This affects the geographic distribution and operating characteristics of divested plants, relative to those that remain regulated, which I discuss directly throughout the paper.

plants typically require years of planning and development before they are operational (and appear in data). Expansions at existing plants can occur more quickly, because major bottlenecks, such as interconnection, are already overcome.

While capital investment is observed while power plants are regulated, due to reporting requirements to the Federal Energy Regulatory Commission (FERC), actual investment is unobserved once a plant is deregulated. As a result, I use the capacity of a power plant as a proxy, mirroring an approach in prior work (Davis and Wolfram, 2012). The "capacity" of a power plant is the technical name for its size–roughly, it can be thought of as how much energy the plant is "capable" of generating–and is measured in megawatts (MW). Because power plants are modular and can consist of multiple generators, a plant will become bigger when capital is spent to either add a generator or uprate existing units, and vice versa. Capacity acts as a conservative proxy for investment, because capital can be spent on the plant in other ways, such as pollution controls (Fowlie, 2010), that would not be reflected in a larger plant. Later in the paper, I convert the identified change in capacity (MW) to an estimated change in investment ($) using average cost factors ($/MW) for different generator types.

I focus the analysis on a sample of 800 fossil power plants that were all operating when systematic data collection on plant capacity began in 1990. I choose to study fossil power plants–those that burn coal, oil, and natural gas–because they are more likely to be modular than other generation technologies. (For example, the capacity of a hydro plant is likely to be determined by the size of the river and dam, and thus less likely to change over time.) In 1990, at the start of the panel, restructuring had yet to occur, and all power plants in the sample were regulated–that is, built, owned, and operated by investor-owned utilities whose prices were set according to cost-of-service regulation that was defined by states. Beginning in 1998, a total of 273 of the plants within the sample were eventually divested and removed from the rate base of regulated utilities. I limit the panel to 1990-2010, for a total of 21 years, as a way to avoid confounding from build driven by renewable portfolio standards and the reduction in natural gas prices that dominate the next decade. A central identifying assumption in this design is that the plants which remain regulated form an appropriate counterfactual for the level of investment divested plants would have received over time. Empirical evidence supporting this assumption can be provided by assessing the similarity in baseline characteristics between treated and control plants.

The main estimating equation is a two-way fixed effect model:

$$y_{pt} = \beta d_{pt} + \theta_p + \gamma_t + \epsilon_{pt}. \tag{1}$$

Here, p indexes a power plant, t indexes a year, and d is a binary variable equal to 1 once a plant has been divested. The outcome, y, measures the operable capacity of plant p in period t; if the power plant retires, y is 0. The plant-level fixed effect, θ_p, captures time-invariant characteristics of each power plant that can affect investment decisions, and the year fixed effect, γ_t, captures common shocks to new build, such as changes in resource availability, fuel prices or macro-economic shifts to demand. The coefficient of interest is β: the average change in capacity, after a plant is divested, in a given year. With capacity used as a proxy for capital, the design identifies the effect of deregulation on new investment in power plants.

Because plants are divested in different years (i.e., treatment timing is staggered), and the effect of divestment is unlikely to be homogenous across both plants and time, the estimate of β in Equation 1 can be biased due to problematic comparisons between adopting groups embedded in the simple two-way fixed effect model (Goodman-Bacon, 2021; Wing et al., 2024). To address the issue, I include additional estimators that are robust to staggered treatments and heterogeneous treatment effects (Callaway and Sant'Anna, 2021; Gardner, 2022; Wing, Freedman and Hollingsworth, 2024). However, because I find, in practice, that the simple two-way fixed effect estimates are very similar to those from stagger-robust estimators, I maintain the two-way fixed model as the main specification and provide the results from the alternative estimators as robustness checks.

Four additional points about the research design are worth highlighting. First, regional market conditions–such as, differential levels of new build, or changes in electricity demand–can create the possibility of group-specific time trends, based on any difference in the geographic distribution of divested and regulated plants. This can be addressed empirically by assessing the evidence for differential market conditions across the sample. Second, strict exogeneity can be violated if the original plant owners under-invested in anticipation of divestiture. This should manifest as a differential pre-trend which, assuming sufficient statistical power, can be detected within an event study. Third, SUTVA violations are a natural concern, because electricity markets are interconnected and the likelihood of a plant expansion can be affected by the capacity decisions of others within a region. This threat is a weakness of many empirical studies in electricity and is difficult to overcome with existing methods. Finally, divestment, as a treatment, can capture the combined effect of two underlying changes: the effect of a new owner and the effect of new economic incentives. The relative contribution of each can be explored by leveraging heterogeneity in the type of divestment that occurred (i.e., a transfer to an unregulated subsidiary of the original utility versus sale to an un-affiliated entity).

2 Data

The data I use for the analysis comes from three sources: (i) the Energy Information Administration (EIA), (ii) the Environmental Protection Agency (EPA), and (iii) the Federal Energy Regulatory Commission (FERC). What follows is a brief description of the source and construction of the primary variables used in the analysis.

2.1 Capacity and other characteristics of power plants

The measure of plant-level capacity, and my primary outcome variable, comes from EIA Form-860 (EIA, 1990-2010*b*). Beginning in 1990, the annual survey has collected detailed information about the technical characteristics of all US power plants above 1 MW. In each year of the panel, I observe the total operable capacity of each power plant, as well as the underlying characteristics of its generators, such as their age, primary fuel source, and turbine type. "Operable" generators include those that are actively in-service and those on standby and exclude those that are retired or under construction. I use the reported nameplate capacity of each generator, because summer de-rated capacity is not available consistently for all years of the survey. I limit the sample to plants that, in 1990, were located in the continental US and that were owned by utilities subject to cost-of-service regulation. I identify regulated utilities using the type of each plant's owner in EIA Form-861 (EIA, 1990-2010*c*), which is an annual census of the activities of US electric utilities, and limiting to investor-owned utilities that served end-load customers. Finally, I include only fossil plants that burn primarily coal, oil, or natural gas to generate power, removing hydroelectric, nuclear, and a small number of geothermal plants from the sample. I identify fossil plants based on their historic fuel use during the decade prior to the panel (1980-1990), obtained from EIA Form-759 (EIA, 1980-1990). The resulting sample has a total of 800 fossil plants, operating in 1990, that were built and initially owned by regulated utilities.

Table 1 shows the mean characteristics of power plants in the resulting sample, at the start of the panel. The Cohen's d statistic provides the standardized difference of means between the regulated and divested power plants, in (pooled) standard deviation units. Plants in both arms burn a mix of coal, natural gas, and oil, and are roughly 25 years old in 1990. Relative to control plants that remain regulated, the mean divested fossil plant is slightly larger, older, and less likely to have burned coal in the decade prior. A large share of the divested plants are also located in the Northeast, as shown in Figure 1, reflecting the states that adopted restructuring legislation and required or encouraged the divestment of power plants by utilities. Finally, it is worth noting the high average heat rate within the sample.[5] The mean value, near 12,000 Btu per kWh, reflects that the sample contains a mix of plants serving base, intermediate, and peaking load.

2.2 Divestment

I code a plant as divested during the first year in which it reports generation to the EIA as a non-utility (EIA, 1990-2010*d*), following the approach in Davis and Wolfram (2012) and Cicala (2015). Each plant that reported any non-utility generation was considered a "candidate" divestment, and I then cross-verified the fact and timing of each sale with

[5] 'Heat rate' is the inverse of the thermal efficiency of the plant. It captures the amount of fuel a plant uses as input, in British thermal units (Btu), to produce a single kWh of electricity output. The higher the heat rate, the less efficient the plant.

Table 1: Baseline Characteristics of Regulated and Divested Power Plants in the Sample

	Regulated	Divested	Difference	Cohen's d
Count of Plants	527	273		
Total Capacity	252,682	159,514		
Average Size (MW)	479.47	584.30	-104.83	0.16
Generators per Plant	3.50	4.18	-0.68	0.22
Year In-Service	1966.21	1965.61	0.60	0.06
Reported Heat Rate	12,021.90	12,068.93	-47.03	0.02
Fuel Use, Share for Average Plant:				
Coal	39.89	30.31	9.57	0.21
Natural Gas	29.95	33.33	-3.38	0.08
Oil	30.12	36.35	-6.23	0.14
Other (Wood/Waste)	0.04	0.01	0.04	0.06
Generator Types, Share of Average Plant that Uses:				
Combined Cycle	1.46	2.19	-0.73	0.06
Steam Turbine	57.72	62.04	-4.32	0.09
Gas Turbine	28.03	32.74	-4.71	0.11
Internal Combustion	12.79	3.03	9.76	0.34
Other Generator	0.00	0.00	0.00	0.05
Plant Location, Share in:				
Northeast census region	10.44	44.32	-33.89	0.89
Midwest census region	38.14	16.12	22.02	0.49
West census region	16.89	10.62	6.27	0.18
South census region	34.54	28.94	5.60	0.12

Notes: Fuel use is based on each plant's annual average consumption between 1980-1990, as reported in EIA Form-759. Heat rates are self-reported by plant operators in EIA Form-860 in 1990, and the average shown excludes 39 plants that did not report values.

a second data source, such as SEC filings of the original owner. Because I aim to use divestments to identify how competitive markets affect capital investment, I exclude any sales of power plants to a government-owned or cooperative utility, which are subject to unique forms of regulation. I also remove three power plants whose treatment status was non-absorbing (i.e., they were sold back to an investor-owned utility or re-regulated). This left a total of 273 plants that were divested within the sample.

The timing of divestments are shown in Figure 2. Divestments of plants in the sample are staggered between 1998 and 2008, but the majority (93%) occur between 1998 and 2002. Intuitively, this treatment timing is likely to limit the extent of "problematic comparisons" between early and late adopters that would be embedded in a simple two-way fixed estimate, reducing the possible magnitude of bias (Goodman-Bacon, 2021).

Figure 1: Spatial Distribution of Regulated and Divested Plants

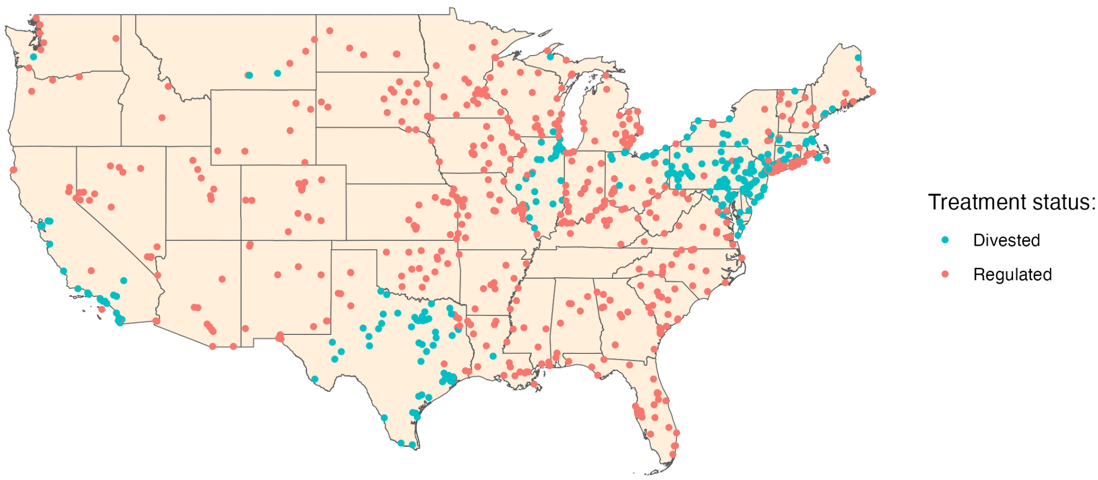

2.3 Fuel use, generation, and carbon emissions

I use data on power plants' fuel use, generation, and carbon emissions to help contextualize the possible welfare effects of a change in investment due to deregulation. I obtain data on plants' operation from the EPA Continuous Energy Monitoring System ('CEMS') (EPA, 1997-2010). The CEMS data began in 1995 to track compliance among plants regulated under the federal Acid Rain Program (which aimed to reduce sulfur dioxide emissions through a cap-and-trade program), though its scope has expanded over time in line with EPA air quality regulations. The two key strengths of the CEMS data are its granularity (variables are collected hourly, at the sub-plant level) and quality (it is based on physical monitors, rather than self-reported data). The trade-off for quality is scope: not all plants who reported capacity to the EIA were initially required to monitor their operations for the EPA.

The analysis on plant operations is thus based on a subsample of power plants for whom complete CEMS data is available. I limit the analysis to plants who reported for all 14

Figure 2: Timing of Divestments

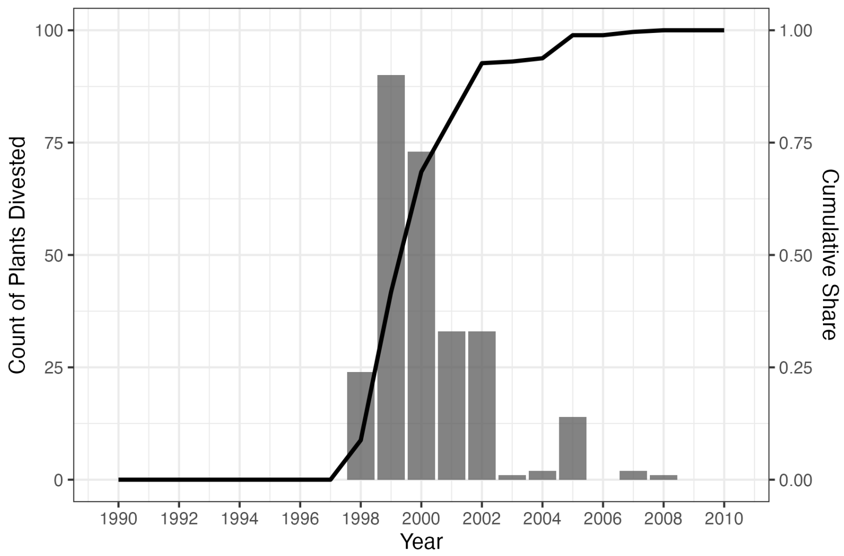

years of available data (1997-2010) and whose operational data did not contain obvious measurement error (e.g., reported generation without fuel use). This identified a total of 310 power plants (the 'CEMS Sub-Sample') that comprise 72% of the capacity in the full sample in 1997. Relative to the EIA capacity sample (Table 1), the plants in the CEMS sub-sample are larger, more efficient, and more likely to burn coal, which reflects the type of plant included in the initial cohorts of the Acid Rain Program. The mean plant was around 1 GW in 1990, obtained about 60-73% of its fuel from coal in the decade prior, and had an average heat rate near 10,500 Btu per kWh. There are also comparatively fewer control plants in the Northeast.[6]

2.4 Regulated utility financial data

I use data in the Form 1 filing from the Federal Energy Regulatory Commission (FERC) to analyze the financial context in which regulated utilities made their investment decisions (FERC, 1994-2010). Form 1 is an annual filing, available digitally from 1994 onwards, that is required of all major utilities (annual sales exceed 1 million MWh), and it contains detailed financial information for the companies' operations and investments. Among the utilities which own power plants in the control group in 1990, I limit the sample to the 74 who report to the FERC and who have a full panel of data (i.e., the reporting company definition is stable from 1994-2010).[7] These companies owned 94% of the control plant

[6]I provide a balance table of mean baseline characteristics, analogous to Table 1, for the CEMS sub-sample in the Online Appendix.

[7]Within the sample, it is possible for a utility to own both control and treated plants. This occurs for one of three reasons: (i) the utility is in a state that required divestment, but the control plant retires before restructuring occurs; (ii) the utility voluntarily divested one or two of its power plants; or (iii) the utility is in a state, such as California or New York, where divestment prompted by restructuring legislation was partial. For the FERC data, I exclude utilities in category (i) and include those in (ii) and (iii) for whom the majority of plants owned by the utility in 1990 remain regulated.

capacity at the start of the panel.

3 The Effect of Divestment on Capital Investment

This section provides the results of the main analysis: the effect of competition on capital investment into fossil power plants, estimated using a difference-in-difference model (Equation 1). The results proceed in three steps. Section 3.1 provides descriptive evidence for how regulated and divested plants evolved over time. Section 3.2 then estimates the difference-in-difference and converts the effect of divestment on capacity into an equivalent dollar amount of investment. Finally, Section 3.3 addresses the different validity threats that could undermine the causal interpretation of the coefficient. Figure 3 provides a visual summary of the coefficient estimates and 95% confidence intervals across specifications in this section.

Figure 3: Visual Summary of Treatment Effect Estimates and 95% Confidence Intervals

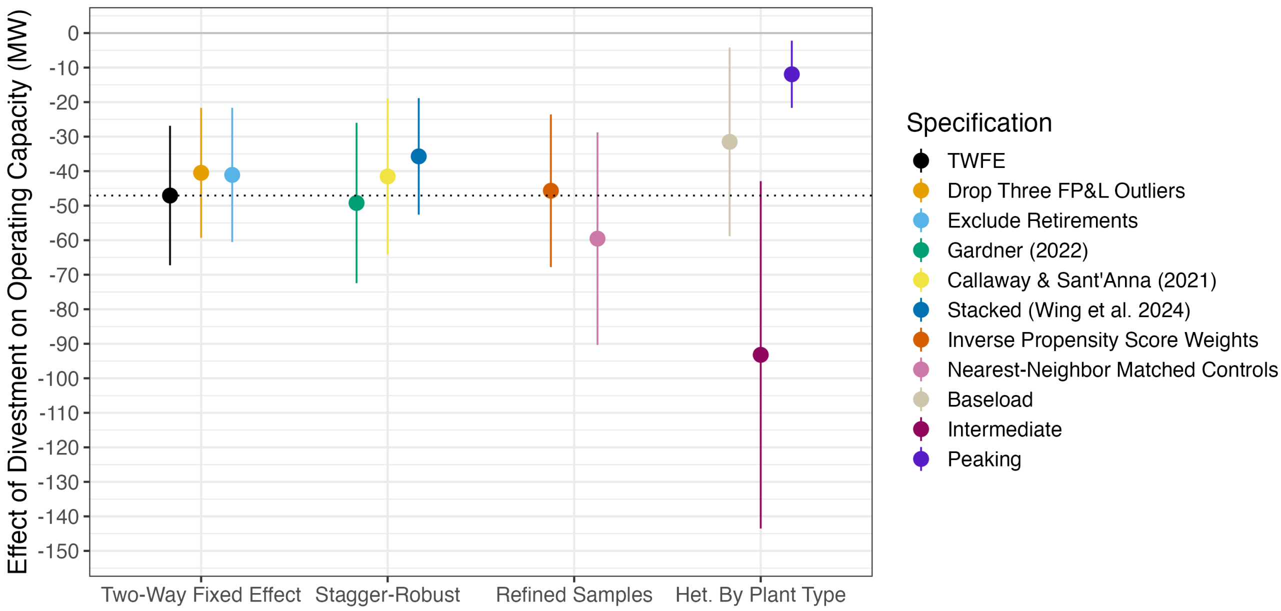

Notes: Standard errors are clustered at the plant level. The 'Two-Way Fixed Effect' and 'Stagger-Robust' groups correspond to Table 2; the 'Refined Samples' to Table 4; and the 'Heterogeneity by Plant Type' to Table 3. As shown in Table 3, plants that serve baseload, intermediate, and peak load have different mean sizes, which affects the interpretation of the relative reduction in size from the absolute effects charted here.

3.1 Descriptive Evidence

The mean evolution of power plant capacity among divested and regulated plants is shown in Figure 4. Two points are evident. First, divested and regulated plants appear to evolve similarly in the pre-period, with minimal change in plant capacity prior to 1998 when divestments begin, as demarcated by the gray dotted line. This supports the validity of the difference-in-difference design. Second, divested and regulated plants evolve differently in the post-period: the mean deregulated plant becomes smaller, while the mean regulated plant becomes larger.

Figure 4: Mean Plant Capacity by Treatment Arm and Year

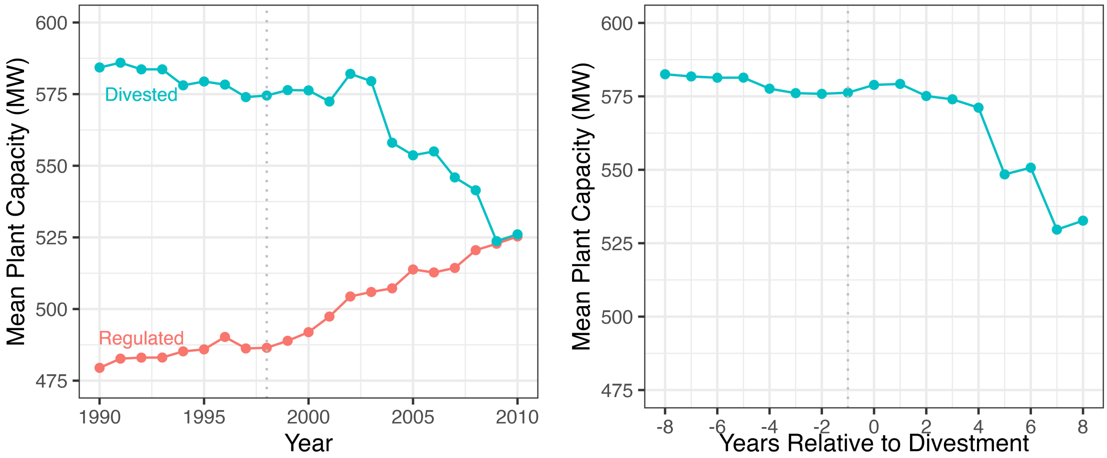

(a) Calendar Time

(b) Event Time

This difference is driven primarily by infra-marginal changes in plant size–that is, a plant adds a generator or uprates its capacity, and vice versa–rather than plant retirements. Divested and regulated plants retire at similar rates, with 19% and 17% of plants retired by 2010, respectively.[8] Instead, divested plants are more likely to become slightly smaller; in 2010, 34% of divested plants had operating capacities below their starting size, 6% more than the equivalent distribution among regulated plants. Regulated plants are more likely to expand; 20% of regulated plants were larger by 2010, 7% more than divested. I also note that there are three clear outlier additions in the sample, all of which happen to be regulated plants owned by one utility, Florida Power and Light: the Fort Meyers, Sanford and Martin plants each more than doubled in size by 2010, equivalent to adding 1.5-2.5 GW of new capacity per plant. Later in the results, I will show the robustness of the effect to excluding these three outliers.

The change in regulated plant capacity is driven by generators that use natural gas. In 2010, 30% of regulated plant capacity reported natural gas as its first energy source, up 11 percentage points over 1990. (Among divested plants, the distribution of capacity reporting natural gas, oil, or coal as its first energy source remains relatively stable over time.) Notably, plants in both arms see a similar rate of adoption of combined cycle generators in the post-period (Figure 5a, shown as a percent of operable capacity). Finally, despite the decline in the mean plant size, the mean age of divested plants still decreases in the post-period (average in-service year increases–Figure 5b). In other words, the mean divested plant is "newer," without adding new capacity. This implies that divested plants retire older generators, such that the composition of generators shifts towards newer capacity within the plant.

[8] Regulated plants that retire are, on average, smaller in size, relative to both the mean of all regulated plants and the mean size of divested plants which retire. Given this, I will show a sample specification that includes only plants which remain operable throughout the panel, excluding retirements in both arms.

Figure 5: Change Combined Cycle Generators and Plant Age Among Operable Capacity

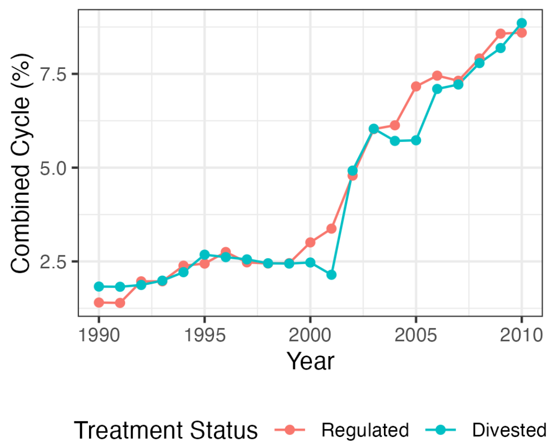

(a) Share of Combined Cycle

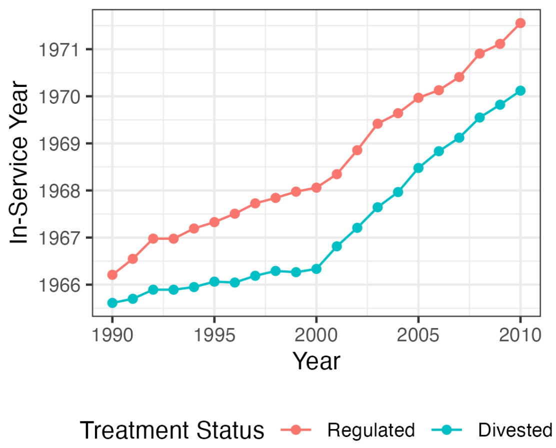

(b) Mean Plant Age

3.2 Difference-in-Difference Estimates

Table 2 shows the results of the difference-in-difference model (Equation 1). Divestment leads to a reduction in plant capacity of 47 MW per year, on average. This effect is equivalent to an 9.1% reduction in plant capacity, relative to the mean plant size among untreated observations (Column 2).[9] The treatment effect is precisely estimated and does not appear to be driven by plant retirements (excluded in Column 3) nor the three outlier expansions among the regulated plants, owned by Florida Power and Light (excluded in Column 4). As mentioned earlier, the simple two-way fixed effect estimate in Column (1) can be biased if the treatment effect is heterogeneous, due to the staggered nature of divestment timing. However, stagger-robust estimators provide a similar treatment effect estimate to the two-way fixed effect model, suggesting the likelihood of bias is low (Columns 5-7).

Fossil power plants are designed to meet different portions of electricity demand, based on the underlying generator technology used. For example, steam turbines that burn coal have a low average cost when operated consistently, but are slow to increase generation; plants that use them usually serve "base" load and will generate throughout the year. In contrast, internal combustion or gas turbines that burn oil or natural gas can ramp generation quickly but incur high marginal costs to do so; plants that use these generators are typically "peaking" capacity that will operate infrequently, during the hours with the highest demand. In Table 3, I examine if the effect of divestment is heterogeneous across plants that share similar operating fundamentals. I assign plants to one of three categories, based on the type of load a plant is likely to serve as a function of its baseline characteristics.[10] I find that the effect of divestment is consistently negative and significant

[9]Note that 9.1% is obtained by transforming the coefficient from the Poisson model using: $(1 - exp(\hat{\beta}_{pois})) * 100$.

[10]Fuel use and operation at the plant level can be quite complex and defy simple categorization, due to

Table 2: The Effect of Divestment on Operating Capacity

	Operating Capacity (MW)						
	Full Sample		Drop Retirements	Drop Outliers	Stagger-Robust		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	OLS	Poisson	OLS	OLS	Gardner	Stacked	C&S'A
Divested	-47.06***	-0.0875***	-41.09***	-40.49***	-49.22***	-35.72***	-41.53***
	(10.29)	(0.0186)	(9.902)	(9.576)	(11.85)	(8.604)	(11.524)
Squared Correlation	0.97123	0.97190	0.97737	0.97574	–	–	–
Observations	16,800	16,800	13,902	16,737	16,800	49,096	16,800
Dependent variable mean	521.94	521.94	610.20	515.43	–	–	–
Plant fixed effects	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓	✓	✓

Notes: Standard errors are clustered at the plant level. The sample in Column (3) includes only plants that remain operable at the end of the panel. The sample in Column (4) excludes the three largest absolute expansions–Fort Meyers, Sanford and Martin power plants, owned by Florida Power and Light–from the control group. Column (5) is estimated using the two-stage difference-in-difference estimator proposed by Gardner (2022). Column (6) is estimated using the stacked estimator as proposed by Wing, Freedman and Hollingsworth (2024). The stacked estimator in Column (6) uses 8 pre-periods and 8 post-periods and includes only plants that were divested in 1998-2002. Column (7) is estimated using Callaway and Sant'Anna (2021), and the aggregate coefficient shown is the average of the mean treatment effect for each timing group. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level.

across plant types, but that it has a relatively larger effect on the size of plants that initially served intermediate and peaking load versus those that served base load. This suggests that the magnitude of the effect of deregulation on capacity may be mediated by the specific way a plant's economics are affected by a competitive market structure.

An event study shows that divested plants do not receive statistically different investment in the pre-period, prior to sale (Figure 6). The figure includes both the simple two-way fixed effect estimate, as well as a stagger-robust estimator (Gardner, 2022). This supports the assumption that parallel trends hold in the pre-period. In addition, the event study is sufficiently powered to identify economically meaningful violations of this assumption (Roth, 2022). A pre-trend with slope 2.56 MW per year and 0.65 MW per year would be identified 99% of the time under the two way fixed effect and stagger-robust event studies, respectively–both of which are much smaller than the estimated treatment effect in the post-period (47 MW per year). This suggests the risk of confounding due to an un-identified pre-trend is minimal. Finally, in the post-period, Figure 6 shows that the effect of divestment on capacity is dynamic and grows larger the longer a plant has been deregulated. (The overall coefficient of 47 MW per year, from Table 2, should be interpreted as, in any given year, after divestment, a deregulated plant is on average 47 MW smaller, relative to its regulated peers.)

the many permutations with which generators of different types can be combined within a single plant. For example, it is not uncommon for a coal plant, which I assign to the "baseload" type, to have a single gas turbine generator that it uses in a peaking capacity. As a result, these categories should be viewed as a coarse but well-informed guess as to the operating profile of the plant overall, as a way to group plants that share similar fundamentals. In Section 3.3, I use a nearest neighbor matching methodology that is better able to accommodate idiosyncrasies in plant design. Finally, I caveat that plant type is endogenous over time and can change based on the addition or retirement of specific generators.

Table 3: Heterogeneous Effects by Plant Type

	Operating Capacity (MW)					
	Baseload		Intermediate		Peaking	
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	Poisson	OLS	Poisson	OLS	Poisson
Divested	-31.52**	-0.0359**	-93.21***	-0.1461***	-11.94**	-0.1301***
	(13.88)	(0.0159)	(25.53)	(0.0388)	(4.947)	(0.0501)
Squared Correlation	0.98330	0.98344	0.91856	0.92208	0.96123	0.96746
Observations	6,048	6,048	4,788	4,788	5,964	5,964
Dependent variable mean	878.46	878.46	608.93	608.93	90.557	90.557
Plant fixed effects	✓	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓	✓
Plant N	288		228		284	
Divested N	81		106		86	
Coal (%)	97.62		3.9		0.48	
Natural Gas (%)	1.34		65.88		33.41	
Oil (%)	0.96		30.23		66.1	
Steam Turbine (%)	96.15		85.84		0.34	
Gas Turbine (%)	3.48		8.1		73.45	
Internal Combustion (%)	0.15		0.46		26.12	

Notes: Standard errors are clustered at the plant level. Plant type is assigned based on a plant's historic fuel use in the decade prior to the panel (1980-1990), as reported in EIA Form-759, as well as the technology of its generators, as reported in EIA Form-860. Further details are provided in the Online Appendix. The coal, natural gas, and oil percentages report the average use of each fuel, at the plant-level, between 1980-1990, as reported in EIA Form-759. The steam turbine, gas turbine, and internal combustion turbine percentages report the average share of each generator type, at the plant-level, in 1990, as reported in EIA Form-860. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level.

The decrease in capacity translates into a meaningful magnitude of avoided investment that otherwise would have been recovered in electricity rates. For example, in 2010, the EIA estimated that capital costs for new combined cycle plants were around $1,273 per kW ($ 2023).[11] The overall treatment effect on capacity (47 MW per year) thus implies avoided investment of $59.9 million per plant and year ($34.2-$85.7 million for the 95% confidence interval), if the counter-factual addition would have been a combined cycle generator. Under a simplified assumption that *all* additions would have been combined cycle, this equates to a total of $16.4 billion ($9.3-$23.4 billion) in avoided investment, each year, across the cohort of 273 divested plants.[12] While this is only a back-of-the-envelope estimate, because actual investment is unobserved once a plant is deregulated, the exercise helps contextualize the magnitude of savings in capital investment, relative to operational savings that have been quantified previously in the literature. Most recently, Cicala (2022) estimated that competitive dispatch of power plants within wholesale power

[11] This value is obtained from EIA (2010*a*). The choice to use costs at the end of the panel is intentionally conservative; technology costs for power plants tend to decrease over time.

[12] The equivalent range for scrubbed coal ($2951 per kW) capacity is $21.6-$54.2 billion, and $6.4-$16.1 billion for combustion turbines ($876 per kW). All values are given in $ 2023.

Figure 6: Event Study of the Effect of Divestment on Plant Capacity

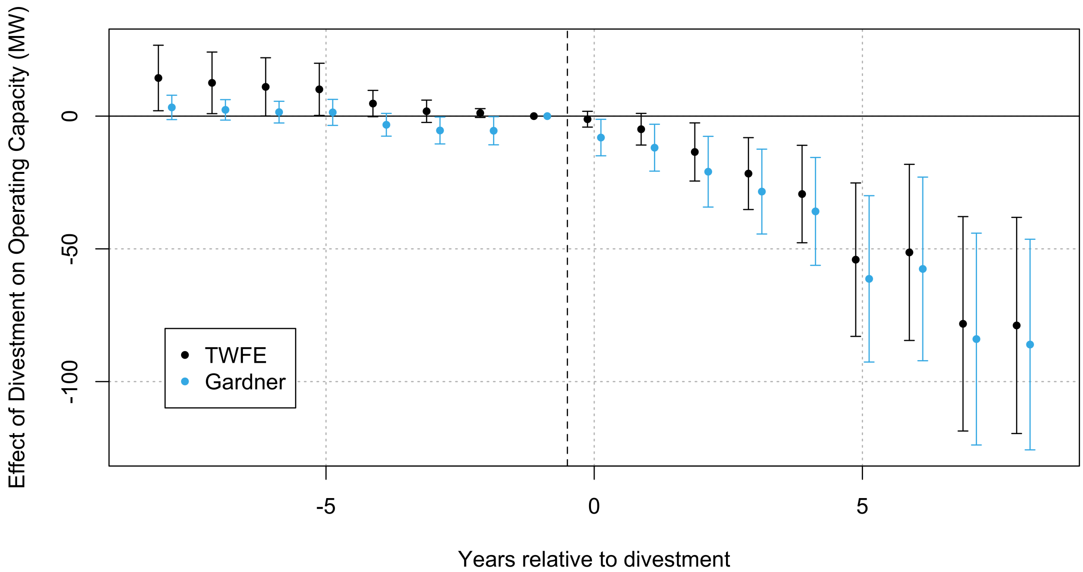

markets led to a reduction of production costs in generation of around $3-5 billion dollars per year. The findings here suggest that reductions in capital investment are likely an order of magnitude higher–greatly expanding our understanding of the total cost reductions achieved by competitive markets in electricity generation in the United States.

3.3 Threats to Validity

A central identifying assumption of the difference-in-difference design is that regulated power plants form an appropriate counter-factual for the level of capital investment divested plants would have received, but for the change in market structure. While the event study provides evidence that the parallel trend assumption likely holds in the pre-period, there are a few reasons why one might imagine that regulated plants evolved differently in the post-period, separate from the effect of regulation. First, regional differences in load growth or entry of new capacity could confound the observed differences in plant capacities, given the underlying spatial distribution of divested and regulated plants (Figure 1). That is, the growth in regulated plants could reflect differential changes in demand in their regions; similarly, the contraction in divested plants could result from differential rates of entry, crowding out the need for more capacity. However, I observe similar patterns of both load growth and entry among divested and regulated plants in the sample, as shown in Figure 7. On the whole, regulated plants tend to be located in states with less load growth and higher rates of entry in the post-period,[13] which would lead one to expect

[13]MacKay and Mercadal (2024) find a similar result regarding entry, when aggregating to the level of regulated and deregulated utility territories.

fewer expansions among these plants–the opposite of what is observed.[14]

Figure 7: Regional Trends for the Mean Plant by Treatment Arm

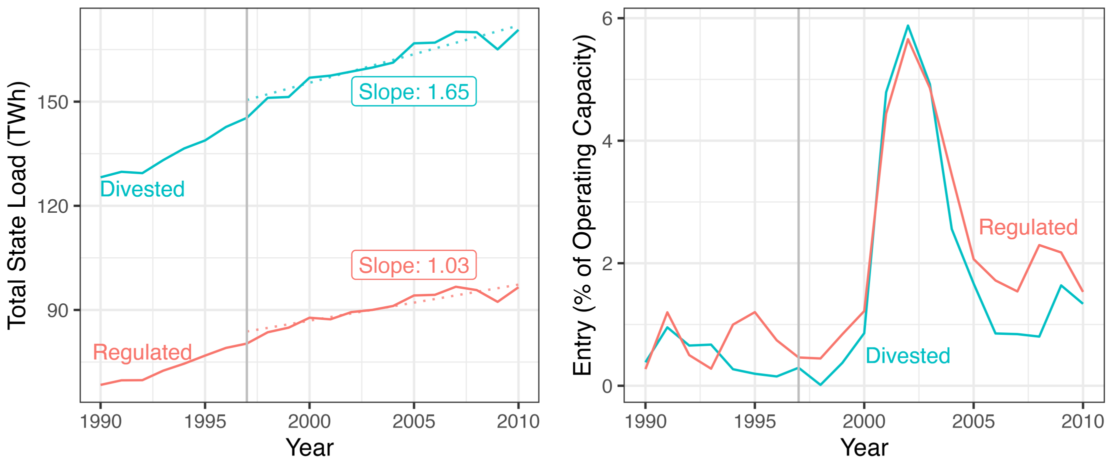

(a) Demand (b) Entry

Note: The figure shows the weighted-average outcomes, at the state-level, where weights are defined by the share (count of plants) in each state, by treatment arm. State is used as a simple way to define the market conditions in which each plant operates. The "Slope" term is the average change in TWh per year between 1998-2010 (the post-period). Entry is calculated as the share of capacity that comes online each year, as a percentage of all operable capacity within the region.

Separately, baseline differences in plant technologies could differentially affect plants' abilities to withstand fuel price shocks in the post-period, thereby impacting profitability and the owner's resulting investment decisions. Of particular concern is the experience and trajectory of natural gas prices during 2000-10 (Joskow, 2013). During this decade, prior to the rise of shale gas, prices at the Henry Hub both rose and became more volatile, spiking multiple times above \$10 per MMBtu, relative to a baseline near \$2.50 in 1999. Divested plants are more exposed to natural gas, at baseline, than their regulated counterparts (Table 1), suggesting they could be differentially vulnerable to gas price swings in the post-period. The average divested plant is also slightly older and less likely to use combustion turbines than their regulated peers. It is possible these differences could confound the estimated treatment effect, to the extent that they are predictive of differential, time-varying trends in investment in the post-period.[15]

To explore the role that plant characteristics may play, I use two strategies to refine the sample and curate control plants that are most similar to the divested plants, at base-

[14] One might be tempted to include load and entry as covariates in the two-way fixed effect model in Equation 1, to test the stability of the treatment effect estimate to these confounds. However, recent methodological literature has noted that an easily-interpretable treatment effect is not identified when time-varying covariates are included in a difference-in-difference setting, unless one is willing to make the strong assumption that the covariates are unaffected by the treatment (Caetano and Callaway, 2023). In this context, while it is unlikely that the divestment of an individual *plant* affects demand or entry in the state in which the plant operates, divestments are highly correlated with state-level restructuring *legislation*, for which the independence assumption is less tenable (Borenstein and Bushnell, 2015).

[15] Time-invariant differences will naturally be absorbed in the plant fixed effect.

line: re-weighting regulated plants according to the inverse propensity-score, and using only a fixed number of the most similar ("nearest neighbor") regulated plants as the counter-factual, for each divested plant. The former is similar to the double-robust estimator proposed by Callaway and Sant'Anna (2021) while the latter mirrors the matched difference-in-difference strategy employed by Cicala (2015); MacKay and Mercadal (2024); Deryugina, MacKay and Reif (2020) and defined formally by Imai, Kim and Wang (2023). The propensity score is estimated using a logit model of divestment on the age, capacity, heat rate, and fuel mix of each plant in 1990, and 5 nearest neighbor matches are identified using an equally-weighted distance metric of the log of capacity, fuel mix, and the plant's generation profile (share of generation by month).[16] As shown in Table 4, the magnitude and precision of the estimated treatment effect remains stable when refining the control plants to those that are most similar to divested plants. This suggests that baseline differences in the observed characteristics are not confounding the estimated effect of deregulation on plant capacity.

Table 4: Sensitivity of Results to Differences in Baseline Plant Characteristics

	Operating Capacity (MW)	
	IPS Weights	Nearest-Neighbor
	(1)	(2)
Divested	-45.68***	-59.55***
	(11.27)	(15.69)
Squared Correlation	0.97114	0.96950
Observations	16,716	13,860
Plant fixed effects	✓	✓
Year fixed effects	✓	✓

Notes: Standard errors are clustered at the plant level. In Column (1), "IPS" stands for inverse-propensity score. In Column (2), divested plants are matched to their 5 most-similar controls. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level.

A different validity threat concerns treatment assignment. If a utility had the ability to choose which plants to divest, we might be concerned that the company opted to only sell those power plants that were the equivalent of "lemons." Importantly, the fear here is that the firm would have private knowledge about plant characteristics that would be unobservable to buyers (or to researchers ex-post). This story would be compatible with the observed decline among divested plants; post-sale, the lower quality of the plant would be revealed to the new owner, who might become less willing to invest in the plant moving

[16]For the inverse propensity-score weights, all divested plants receive a weight of 1, while control plants are weighted according to $\hat{e}_p/(1-\hat{e}_p)$, where $\hat{e}$ is the estimated propensity score. For the nearest-neighbor control sample, the model is estimated using pooled OLS, where control plants are re-weighted according to the frequency with which they are chosen. The two strategies achieve similar improvements in balance within the sample; the key difference is that all regulated plants are included as controls within the inverse-propensity score model (but may have a weight of 0), while only selected plants are included in the matched approach (and may be included in the counter-factual for multiple divested plants).

forward. Intuitively, we would expect selection bias to be minimal when utilities were required to divest their generation assets (i.e., when there was no choice available to the firm regarding which specific plants to sell).

Table 5 leverages heterogeneity in state restructuring policy to split divested plants into three groups: those in states which passed policy requiring the deregulation of all fossil assets (Column 1); those in states where policy only required partial divestment, or where divestment was encouraged but not required (Column 2); and those in states without any legislation that sought to restructure generation (Column 3). Owners of plants in both Columns 2 and 3 had a degree of choice over whether and which plants to sell. The results show a larger effect of divestment on capacity among plants where divestment was partial or encouraged (Column 2), suggesting selection may play a role in these regions.[17] However, I also find that the treatment effect among plants where divestment was required (Column 1) is similar in magnitude and precision to the overall treatment effect in Table 2. This implies that, if selection effects are present, they do not dominate the overall treatment effect estimate presented in the main specification.

Table 5: Exploring Concerns About Selection Into Divestment

	Operating Capacity (MW)		
	Required	Partial or Voluntary	No Legislation
	(1)	(2)	(3)
Divested	-41.40***	-77.19***	29.45
	(10.24)	(24.03)	(38.16)
Squared Correlation	0.97575	0.96982	0.97589
Observations	14,868	12,768	11,298
Dependent variable mean	506.80	518.64	499.60
N Divested Plants	181	81	11
States	CT, MA, MD ME, NJ, OH PA, TX	CA, DE, IL IL, MI, NY	DC, IN, LA, MT, VA, VT, WA
Plant fixed effects	✓	✓	✓
Year fixed effects	✓	✓	✓

Notes: Standard errors are clustered at the plant level. State coding relied primarily upon Andrews (2000); FTC (2000) and individual state's regulation, as needed. Treated plants in each model are limited to those that fall in the states specified. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level.

Finally, the treatment effect for divested plants that were sold will capture the combined effect of two underlying changes: exposure to new economic incentives and to a new owner. One concern is that the observed effect of divestment may be driven entirely by the effect of ownership, rather than incentives, such that attributing the reduction solely to competitive markets would be inappropriate. To explore the relative contribution of

[17]Interestingly, I find no effect of divestment among plants in states without legislation (Column 3); however, the number of divested plants is very small in this subgroup.

each factor, I re-estimate the effect separately among divested plants that were transferred to an affiliate of the original owner and among plants that were sold to an unaffiliated entity (Table 6). Plants that were transferred remain within the same parent company as the original, regulated utility, such that we would expect the effect of new ownership to be minimal, whereas plants that were sold will reflect both changes. The results show that the effect of divestment is larger among plants that were sold, indicating that a change in ownership contributes to the observed reduction in investment. However, the treatment effect among plants that were transferred is similar in magnitude to the overall estimate (Table 2), which suggests that the change in economic incentives is responsible for the majority of the observed effect of divestment.

Table 6: Separating the Effects of Incentives and Ownership

	Operating Capacity (MW)	
	Transfer	Sale
	(1)	(2)
Divested	-41.95***	-57.77***
	(10.92)	(17.51)
Squared Correlation	0.97606	0.96994
Observations	14,112	13,755
Dependent variable mean	508.61	516.72
N Divested Plants	145	128
Plant fixed effects	✓	✓
Year fixed effects	✓	✓

Notes: Standard errors are clustered at the plant level. 'Transfer' includes divested plants that were transferred to a company that is an unregulated affiliate of the original utility. 'Sale' includes divested plants that were sold to a different entity that was not associated with the original owner. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level.

4 Plant Operation and Welfare Implications

The results in Section 3 show that divestment leads to a large and robust decrease in capital investment among fossil power plants. However, avoided investment does not necessarily represent a welfare gain. This is because the investment, in addition to changing the size of the plant, can also affect how it operates. For example, a new generator is likely to be more fuel efficient than older technologies. This might allow a power plant, all else equal, to burn less fuel while generating the same amount of electricity. A new generator may also allow a plant to burn a different fuel, such as natural gas instead of coal. This is particularly relevant in our setting, because most new generators at regulated plants were natural gas-fired (Section 3.1). Because natural gas is roughly half as carbon intensive as coal, the investment might significantly reduce carbon dioxide emissions from the plant. Thus, before concluding that the decrease in capital investment among divested plants represents a social benefit, we need to assess if the change led to any corresponding *loss* in operating performance.

In Table 7, I use data from the EPA for a sub-sample of 310 plants (72% of capacity), from 1997 onwards, to assess the effect of divestment on how plants operated. (The 'CEMS sub-sample' is explained further in Section 2.3.) Each coefficient is obtained from a two-way fixed effect regression, following the main specification in Equation 1, and Figure 8 graphs the corresponding mean of each outcome. (In the Online Appendix, I show the results are not sensitive to the use of stagger-robust estimators.) Before moving to the new operational outcomes from EPA data, I first re-estimate the effect of divestment on plant capacity, to assess the stability of the coefficient between the full and sub-samples (Column 1). I find the effect is nearly identical in magnitude, sign and precision, suggesting plants in the sub-sample are, on average, representative of changes in investment within the full sample.

Table 7: Changes in Plant Operation due to Divestment in CEMS Sub-Sample

	Capacity (1)	Capacity Factor (2)	Operating Hours (3)	Heat Rate (4)	CO2 Intensity (5)	CO2 Emissions (6)
			Panel A: Outcomes in Levels			
Divested	-45.22***	0.3535	-581.2***	-16.68	-0.1487	-101,393.4
	(14.53)	(1.506)	(171.0)	(145.1)	(1.209)	(110,018.6)
Mean Outcome	994.70	53.666	7,518.2	10,619.4	91.422	4,596,126.6
Squared Correlation	0.97499	0.87510	0.78904	0.67788	0.89390	0.97985
Observations	4,340	4,340	4,340	4,340	4,340	4,340
			Panel B: Logged Outcomes			
Divested	-0.0445**	-0.0149	-0.1392***	-0.0027	-0.0039	-0.2046***
	(0.0192)	(0.0387)	(0.0485)	(0.0124)	(0.0134)	(0.0726)
Squared Correlation	0.98397	0.85552	0.74064	0.69597	0.84727	0.92626
Observations	4,337	4,337	4,340	4,340	4,340	4,340
Plant fixed effects	✓	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓	✓

Notes: Standard errors are clustered at the plant level. Sample is limited to the 310 plants in the CEMS sub-sample. Capacity is obtained from the EIA, while all other outcomes are from the EPA. The units for capacity are MW; for capacity factor are percentage points; for heat rate are MBtu of heat input per MWh of gross generation; for CO2 intensity are short tons per MBtu; and for CO2 emissions are short tons. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level.

Moving from left to right in Table 7, I first assess how divestment affected the total generation of a plant (Columns 1-3), before estimating the impact on how the plant used fuel (Columns 4-5), and ending with the net effect on the total carbon dioxide emissions (Column 6). These six variables, including capacity, have the convenient relationship that:

$$CO_{2,t} = Generation_t \cdot HeatRate_t \cdot CO_2Intensity_t \tag{2}$$

$$= \{Capacity_t \cdot (CapacityFactor_t|Op.) \cdot OpTime_t\} \cdot HeatRate_t \cdot CO_2Intensity_t \tag{3}$$

$$= \{MW_t \cdot \frac{MWh_t}{MW_t * Hours_t} \cdot Hours_t\} \cdot \frac{MBtu_t}{MWh_t} \cdot \frac{st_t}{MBtu_t} \tag{4}$$

In other words, the total carbon emissions from a plant (Column 6) can be obtained from the product of its capacity, capacity factor, operating time, heat rate and the carbon intensity of its fuel mix (Columns 1-5). It follows, by taking the log of both sides of the equation, that the percentage change in carbon emissions is approximately equal to the sum of the percentage change in the other factors. This is shown in Panel B of Table 7, which reports coefficients from log-linear, two-way fixed effect specifications.

Starting with how much plants generate, I find that divestment reduces the number of hours a plant operates: operating time declines by approximately 13.0%, after deregulation (around 581 hours each year) (Column 3 and Figure 8c). This result suggests that divested plants may have been dispatched less frequently within competitive market structures. However, I do not find that divestment affects the *rate* at which the plant produces power in the hours when it does operate. This is captured by the capacity factor (Column 2 and Figure 8b), which normalizes how much generation each unit of capacity produces, on average, within a given hour.[18] The null result indicates that divestment does not affect the average productivity of plants, once operating. The combined effect of the changes in capacity and operating time is an overall decrease in annual generation among divested plants, relative to regulated, of approximately 18.0%, on average.

Turning to fuel use, I find no effect of divestment on the annual average heat rate of plants nor the carbon intensity of their fuel mix (Columns 4-5 and Figures 8d-8e). Both results are somewhat striking. Heat rate quantifies the amount of fuel a plant needs to burn as input, to produce each unit of output. It thus measures the thermal efficiency of the plant–the heat rate declines (less fuel, per unit output) when fuel efficiency improves–and, as described earlier, is expected to be lower among newer technology. As shown in Figure 8d, both divested and regulated plants demonstrate markedly similar gains in average heat rates over time. This suggests that the expansions at regulated plants did not lead to a disproportionate gain in the overall thermal efficiency of the plants, and that divested plants were able to achieve similar efficiency improvements, with less apparent capital input. Similarly, the carbon intensity of the plant's fuel mix would decline if a plant switched from coal towards natural gas. The fact that there is no detectable change in the carbon intensity of heat input among divested and regulated plants indicates that the natural gas additions at regulated plants did not substantially displace their use of coal. The end effect is an overall decrease in carbon emissions, after divestment, of about 18.5%, driven nearly entirely by the relative reductions in capacity and operating time.[19]

[18]For example, a 50 MW plant with a 10% capacity factor would generate 5MWh of electricity in one hour. Here, I have defined the capacity factor to be conditional on operation, meaning that the capacity in the denominator is multiplied by the number of hours the plant operates within a year, rather than 8,760 (the total number of hours in a non-leap year). This is shown in Equation 4.

[19]Why is there no detectable effect on the *level* of carbon dioxide emissions, in Panel A? As a variable, the mass of carbon dioxide emissions has a large range across plants in the sample: a power plant at the 75th percentile of emissions would have released over five times more carbon than a plant at the 25th in 1997. As shown in Figure 8f, the mean level of carbon emissions–which will be affected by high-emitting plants–is relatively stable over time and similar between groups, providing the intuition for the insignificant coefficient in Panel A. The decrease in carbon dioxide is instead concentrated among lower-emitting, infra-

Figure 8: Mean Operational Outcomes Among CEMS Sub-Sample

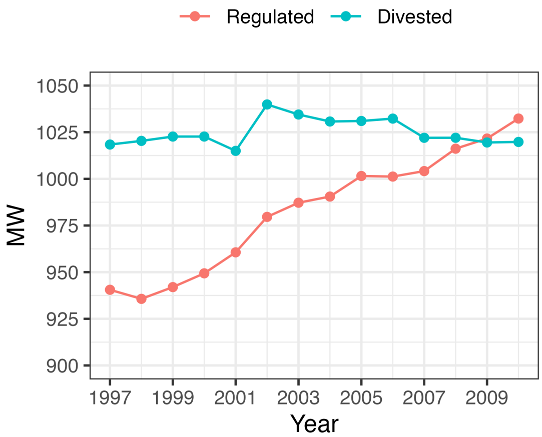

(a) Capacity

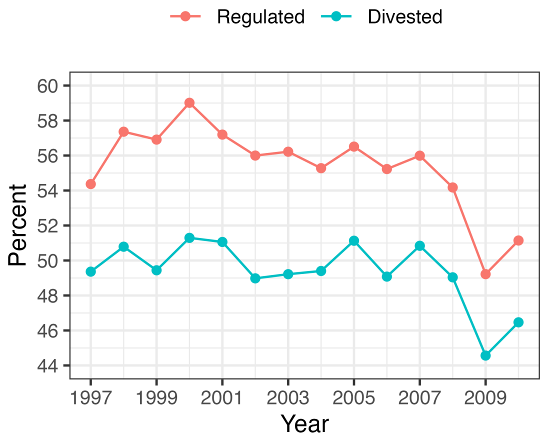

(b) Capacity Factor

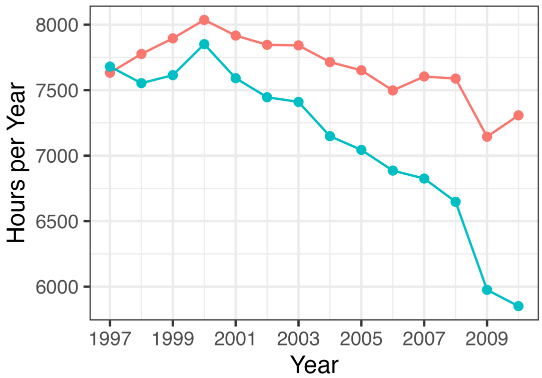

(c) Operating Time

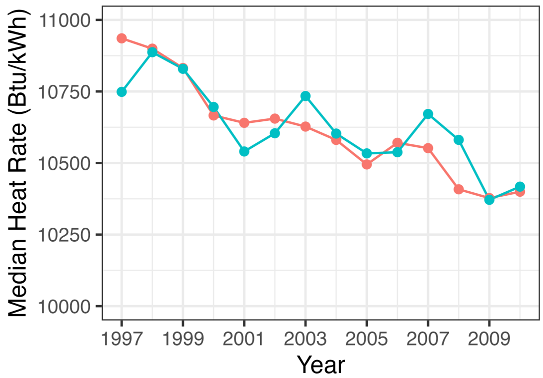

(d) Heat Rate

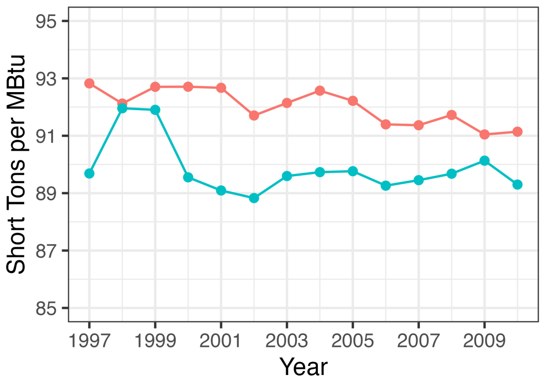

(e) Carbon Intensity of Fuel

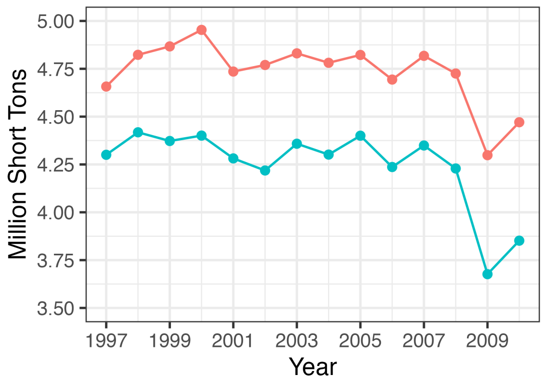

(f) Carbon Dioxide Emissions

One concern about the above results is that the impact of new natural gas generators among regulated plants may have been muted, due to the specific (high) realization of natural gas prices in the post-period. As a result, natural gas generators may have been less economic to run during these years, in a way that would not reflect their long-term

marginal plants, whose generation declines as they operate less.

economics after prices fell due to fracking. Luckily, we are able to learn something about plant behavior during "low" natural gas prices by limiting our attention to the final two years of the panel, 2009-10, when the annual average Henry Hub price fell to $3.94 and $4.37 (compared to $8.86 in 2008; all values in nominal $ per MMBtu). Even when allowing for a differential treatment effect during low (2009-10) and high (2008 and earlier) natural gas prices, I still do not find a statistical difference in average heat rates or carbon intensity among regulated and divested plants (Table 8). While the share of generation from natural gas was 2.97 percentage points higher among regulated plants under low prices (Column 5), this marginal increase was not large enough to detectably shift the overall carbon intensity nor thermal efficiency at the plant-level.[20]

Table 8: Sensitivity of Operational Outcomes to Natural Gas Prices

	Heat Rate		CO2 Intensity		% Natural Gas
	Level	Log	Level	Log	Level
	(1)	(2)	(3)	(4)	(5)
Divested (Low NG Prices)	-15.12	-0.0006	0.6824	0.0078	-2.970*
	(192.4)	(0.0168)	(1.500)	(0.0158)	(1.779)
Divested (High NG Prices)	-17.03	-0.0032	-0.3382	-0.0065	-0.8114
	(141.4)	(0.0120)	(1.158)	(0.0131)	(1.181)
Squared Correlation	0.67788	0.69598	0.89397	0.84735	0.96377
Observations	4,340	4,340	4,340	4,340	4,340
Mean Outcome	10,619.4	9.2618	91.422	4.4894	24.173
Plant fixed effects	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓

Notes: Standard errors are clustered at the plant level. Sample is limited to the 310 plants in the CEMS sub-sample. 'Low NG Prices' limits the post period years to 2009, and 'High NG Prices' includes all post years in 2008 and earlier. '% Natural Gas' is the share of generation using natural gas as the primary fuel. The units of heat rate are MBtu of heat input per MWh of gross generation; of CO2 intensity are short tons per MBtu; and '% Natural Gas' are percentage points. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level.

It's worth noting that the CEMS data does not help us identify the effect of divestment on the operation of peaking plants, specifically, which comprise an estimated 5.72% of capacity in the full sample in 1990. The sub-sample studied here consists primarily of plants that would serve base and intermediate load, based on the frequency of their operation at baseline. (The mean plant operated for 7,646 hours–87% of the year–in 1997, and only 9% of plants operated for half the year or less.) It is difficult to anticipate if divestment would affect peaking plants differently than those serving base or intermediate load. On the one hand, peaking plants already operate infrequently during the year, such that they may be

[20] Cullen and Mansur (2017) find evidence of meaningful coal-to-natural gas switching among US power plants, in response to natural gas price fluctuations. I note that this paper asks a slightly different analytical question–not, "Did plants switch?", but, "Did divested plants switch by a greater amount than regulated plants?"–and studies a more limited sample of power plants (only those owned by regulated utilities and operating in 1990).

less likely to experience the changes in operating time observed here. On the other hand, because heat rates are nonlinear and increasing at low levels of generation, it's possible any small shift in operating time among peaking plants could differentially affect their average thermal efficiencies. Given this, how deregulation may affect peaking plants, in particular, remains an open question for future work.

Overall, these findings suggest that the decrease in capacity among fossil power plants, due to divestment, can be interpreted as a likely social benefit. The considerable savings for consumers in avoided investment–on the order of $9.3-$23.4 billion per year–do not appear to be offset by any evident degradation in plant operation among divested plants. While divestment does lead fossil plants to operate less frequently,[21] there is no observable change in plants' rate of generation nor their thermal efficiency, nor is there a relative increase in the carbon intensity of their fuel mix. An equivalent way to frame these results is that I do not observe a meaningful *gain* in the performance of regulated plants, the social value of which might have offset the higher cost of capacity additions.

5 Averch-Johnson and External Validity

Economists have long noted that cost-of-service regulation can lead a firm to over-allocate production inputs towards capital. This is referred to formally as the "Averch-Johnson effect," after the eponymous authors of the original model (Averch and Johnson, 1962), and colloquially in electricity as "gold-plating" the grid. In Section 3, the effect of divestment on capacity is identified, in part, on counter-factual additions observed among regulated power plants. In this section, I demonstrate that this higher level of investment by regulated utilities coincided with the conditions under which theory predicts the Averch-Johnson effect is most likely to emerge. This affects how we interpret the estimated effect of divestment and the likelihood it would be externally valid for other generating technologies, such as renewables.

The intuition for these conditions rests on a difference between how Averch and Johnson modeled cost-of-service regulation and how it is implemented by US states to govern electric utilities. The Averch-Johnson model assumes that the regulator fixes the return on investment a firm earns. In practice, states regulate electric utilities' price, rather than their return (Lazar, 2016). Policymakers do use a "regulated" return as an input to determine prices within a rate case, but once the price is fixed, regulators do not control nor oversee the actual returns utilities achieve. As Paul Joskow showed, under this set-up, the regulated return is most likely to act as a binding constraint during periods when the utility's average production costs are increasing (Joskow, 1974). It follows that Averch-Johnson style effects–or, an increase in capital investment by utilities–are most likely to be observed when utilities' average costs are rising.

[21] I note that this result is specific to the type of power plant (fossil) studied here. Davis and Wolfram (2012), for example, find that divested nuclear plants operate more frequently.

In Figure 9 and Table 9, I use financial information reported by regulated utilities to the FERC to assess how their operating costs evolved over the course of the panel. The sample consists of 74 regulated utilities who own control plants in the sample and who report consistently to the FERC (Section 2.4). The analysis shows that average production costs increased significantly during the post-period. The utilities' operating expenses began to rise around 2000 (Figure 9a), leading average production costs to be 33.6% higher, across the sample, from 1998-10 (Figure 9b and Column 4 of Table 9).[22] In line with Joskow's prediction (Joskow, 1974), we observe a coincident and significant increase in overall capital investment: the utilities' invested 31.1% more on average, each year, during the post-period, when production costs were higher (Figure 9c and Column 6 of Table 9). These changes led to a decrease in estimated returns on investment earned by the companies, from a high of 8.9% in 1995-96 to 6.4% in 2010 (Figure 9d). Taken together, the data and theory indicate that the Averch-Johnson effect likely biased the investment decisions of regulated utilities during the post-period.

Table 9: Change in Regulated Utility Operating Costs and Investment

	Operating Expenses ($ mil)		Average Cost ($/kWh)		Net Plant Investment ($ mil)	
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	Poisson	OLS	Poisson	OLS	Poisson
Post-Period (1998-2010)	610.2***	0.4270***	0.0238***	0.2899***	810.5***	0.2707***
	(73.52)	(0.0314)	(0.0025)	(0.0288)	(112.4)	(0.0244)
Squared Correlation	0.85844	0.87785	0.62497	0.63776	0.88982	0.89939
Dependent variable mean	1,612.2	1,612.2	0.08916	0.08916	3,226.8	3,226.8
Observations	1,258	1,258	1,255	1,255	1,258	1,258
Utility fixed effects	✓	✓	✓	✓	✓	✓

Notes: Standard errors are clustered at the utility level. *** indicates significance at the 1% level, ** indicates significance at the 5% level, and * indicates significance at the 10% level.

Given this evidence, one way to view the treatment effect estimated in Section 3 is as a measure of the magnitude of the Averch-Johnson effect within electricity generation–that is, it captures what happens to investment in fossil power plants when the distortion is removed, relative to a competitive baseline. This interpretation can inform how we view the likely external validity of the treatment effect to other generating technologies. The Averch-Johnson distortion is not limited to investment in any one specific type of power plant, and as a result, it suggests that the effect of regulation identified here–that regulated fossil plants are on the order of 9.1% larger, on average (Table 2)–is more likely to hold for other generating technologies, such as solar and wind, given similar cost conditions for utilities.[23] In other words, the results suggest that there is a risk that regulation may inflate the level of investment in new renewable power plants, just as it did among existing fossil plants.

[22]The spike in 2000 is driven by an increase in purchased power costs, observed across utilities (i.e., it does not appear to be driven by a single company).

[23]In a recent paper, I showed that the increase in operating costs for regulated utilities has continued beyond 2010 and affected companies nationally, beyond those in this particular sample (Yozwiak, 2023).

Figure 9: Assessing if the Averch-Johnson Effect Distorted Regulated Utility Investment

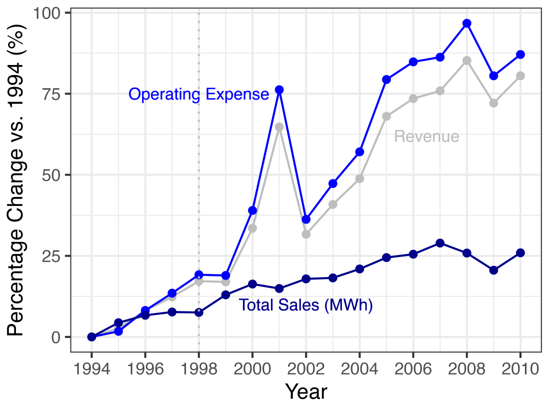

(a) Inputs to Average Production Cost

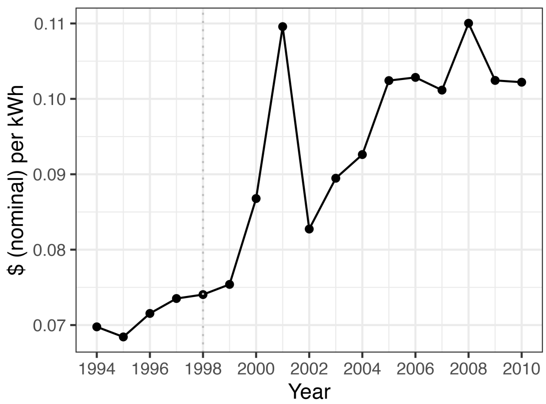

(b) Average Production Cost

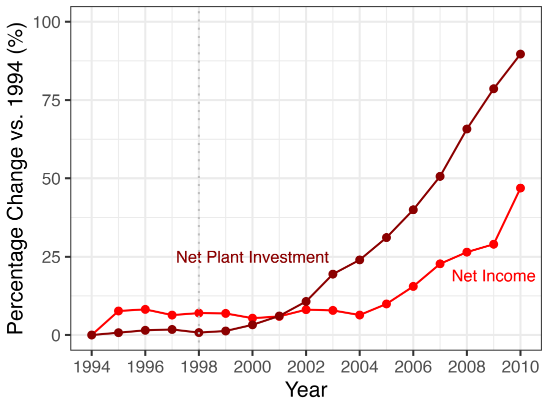

(c) Inputs to Return on Investment

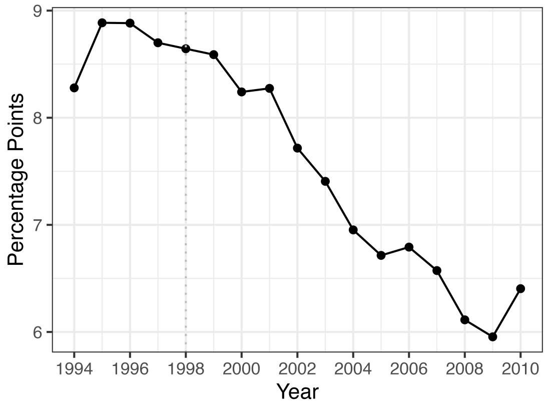

(d) Return on Investment

Notes: All data is obtained from FERC Form 1. The sample consists of 74 utilities who own control plants in the sample and who report consistently to the FERC. Operating expense is total of operation and maintenance costs, amortization, depreciation, and taxes, net of deferrals and adjustments. It captures the amount of operating costs a utility would recover within regulated rates, prior to any adjustments by regulators. Average cost is operating expense divided by total sales. Net plant investment is the total un-depreciated value of capital investment made by the utility across functions. It represents the investment on which a utility would earn a return within regulated rates, prior to any adjustments by regulators. Net income is the difference between revenue and operating expense. The return on investment is net income divided by net plant investment. The values shown are the mean outcome across the sample. The average production cost is given in nominal dollars (Joskow, 1974), because regulated prices do not typically contain inflation adjustments. An equivalent chart in real dollars is provided in the Online Appendix.

6 Conclusion

How do competitive markets affect capital investment? In this paper, I document that competitive reforms led to a large and robust decrease in capital investment among fossil power plants in the United States. I identify the effect of competition by using the large-scale divestiture of power plants that occurred in the late 1990s, in which regulated utilities sold or transferred hundreds of power plants to unregulated entities, as part of broader reform. By comparing fossil power plants before and after divestment, I find that power plants are 47 MW (9.1%) smaller, on average, after exposure to competitive markets, relative to similar units that remain regulated. The reduction in size corresponds to roughly \$9.3-\$23.4 billion in avoided investment, per year, across the cohort of 273 divested plants. Divested power plants do not operate less efficiently, as a result of the decreased expenditure, suggesting that the cost savings are welfare-enhancing. Finally, I show that competitive reforms reduced investment in fossil plants, in part, by removing a distortion to utility investment within cost-of-service regulation. This distortion is not specific to one technology, suggesting regulation may similarly inflate utilities' future investment in renewable power plants as it did for past investment in fossil.

I highlight two questions that emerge from the analysis which could motivate additional research. First, regulated utilities, to date, have not led new investment in renewables, raising concerns that they may actually under-invest in the clean energy transition (Andonov and Rauh, 2024; Fogler and Ver Beek, 2023). Future work may seek to explore the decision from a dynamic perspective, incorporating how regulated utilities' pre-existing stock of natural gas capacity—itself affected by regulation (Gowrisankaran, Langer and Reguant, 2024)–can impact firms' investment decisions going forward. Second, I find a connection between the operating characteristics of a power plant and how it is affected by competitive markets. Because renewables operate differently than traditional, fossil power plants (Joskow, 2011), one open question is how their characteristics may mediate the magnitude of the effect of competitive reforms (e.g., larger or smaller than that observed for fossil plants?). As one illustration, solar and wind can be curtailed, if the supply of renewable generation exceeds demand (Novan and Wang, 2024). We might imagine, as a result, that the effect of competitive reform could vary based on the level of renewable capacity that is already operating within a region, creating interesting patterns of spatial and time heterogeneity.

The United States stands on the precipice of massive investment in electricity generation. To decarbonize the electricity system, as a means to mitigate climate change, will require more than doubling the existing fleet of power plants. The analysis in this paper highlights the crucial role that market structure can play in determining the total cost of this transition to consumers. A new renewable grid, built to be around 9.1% too large, could cost consumers hundreds of billions in unnecessary investment. As a result, policymakers and academics may wish to consider competitive reforms as a policy tool to help reduce the long-term infrastructure costs of decarbonization.

References

Andonov, Aleksandar, and Joshua Rauh. 2024. "The Shifting Finance of Electricity Generation." National Bureau of Economic Research Working Paper 32733.

Andrews, Clinton J. 2000. "Diffusion Pathways for Electricity Deregulation." *Publius: The Journal of Federalism*, 30(3): 17–34.

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

Bistline, John, Kimberly A Clausing, Neil Mehrotra, James H Stock, and Catherine Wolfram. 2024. "Climate Policy Reform Options in 2025." National Bureau of Economic Research Working Paper 32168.

Borenstein, Severin, and James Bushnell. 2015. "The US Electricity Industry After 20 Years of Restructuring." *Annual Review of Economics*, 7(Volume 7): 437–463.

Borenstein, Severin, and Ryan Kellogg. 2023. "Carbon Pricing, Clean Electricity Standards, and Clean Electricity Subsidies on the Path to Zero Emissions." *Environmental and Energy Policy and the Economy*, 4: 125–176.

Bushnell, James B., and Catherine Wolfram. 2005. "Ownership Change, Incentives and Plant Efficiency: The Divestiture of U.S. Electric Generation Plants." *Working paper.*

Bushnell, James B., Erin T. Mansur, and Kevin Novan. 2017. "A Review of the Economics Literature on Electricity Restructuring." *Working paper.*

Caetano, Carolina, and Brantly Callaway. 2023. "Difference-in-Differences with Time-Varying Covariates in the Parallel Trends Assumption."

Callaway, Brantly, and Pedro H.C. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230. Themed Issue: Treatment Effect 1.

Cicala, Steve. 2015. "When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation." *American Economic Review*, 105(1): 411–44.

Cicala, Steve. 2022. "Imperfect Markets versus Imperfect Regulation in US Electricity Generation." *American Economic Review*, 112(2): 409–41.

Cullen, Joseph A., and Erin T. Mansur. 2017. "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution." *American Economic Journal: Economic Policy*, 9(3): 106–33.

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

Denholm, Paul, Patrick Brown, Wesley Cole, and et al. 2022. "Examining Supply-Side Options to Achieve 100% Clean Electricity by 2035." National Renewable Energy Laboratory NREL/TP-6A40-81644, Golden, CO.

Deryugina, Tatyana, Alexander MacKay, and Julian Reif. 2020. "The Long-Run Dynamics of Electricity Demand: Evidence from Municipal Aggregation." *American Economic Journal: Applied Economics*, 12(1): 86–114.

EIA. 1980-1990. "Historic Form EIA-906 Detailed Data with previous form data (EIA-759)."

EIA. 1990-2010*b*. "Form EIA-860."

EIA. 1990-2010*c*. "Form EIA-861."

EIA. 1990-2010*d*. "Form EIA-923 detailed data with previous form data (EIA-906/920)."

EIA. 2010*a*. "Assumptions to the Annual Energy Outlook 2010." DOE/EIA-0554(2010).

EPA. 1997-2010. "Clean Air Markets Program Data."

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

FERC. 1994-2010. "Form 1, 1-F, & 3-Q (Electric) Historical VFP Data."

Fogler, Cara, and Noah Ver Beek. 2023. "The Dirty Truth About Utility Climate Pledges." Sierra Club White Paper.

Fowlie, Meredith. 2010. "Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement." *American Economic Review*, 100(3): 837–69.

FTC. 2000. "Staff Report: Competition and Consumer Protection Perspectives on Electric Power Regulatory Reform."

Gardner, John. 2022. "Two-stage differences in differences." *Working paper.*

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, 225(2): 254–277.

Goulder, Lawrence H., Marc A. C. Hafstead, and III Williams, Roberton C. 2016. "General Equilibrium Impacts of a Federal Clean Energy Standard." *American Economic Journal: Economic Policy*, 8(2): 186–218.

Gowrisankaran, Gautam, Ashley Langer, and Mar Reguant. 2024. "Energy Transitions in Regulated Markets." National Bureau of Economic Research Working Paper 32088.

Hill, Alexander. 2021. "Excessive entry and investment in deregulated markets: Evidence from the electricity sector." *Journal of Environmental Economics and Management*, 110: 102543.

Imai, Kosuke, In Song Kim, and Erik H. Wang. 2023. "Matching Methods for Causal Inference with Time-Series Cross-Sectional Data." *American Journal of Political Science*, 67(3): 587–605.

Ishii, Jun, and Jingming Yan. 2008. "Does divestiture crowd out new investment? The "make or buy" decision in the U.S. electricity generation industry." *The RAND Journal of Economics*, 38(1): 185–213.

Jacobson, Mark Z., Mark A. Delucchi, Zack A.F. Bauer, Savannah C. Goodman, William E. Chapman, Mary A. Cameron, Cedric Bozonnat, Liat Chobadi, Hailey A. Clonts, Peter Enevoldsen, Jenny R. Erwin, Simone N. Fobi, Owen K. Goldstrom, Eleanor M. Hennessy, Jingyi Liu, Jonathan Lo, Clayton B. Meyer, Sean B. Morris, Kevin R. Moy, Patrick L. O'Neill, Ivalin Petkov, Stephanie Redfern, Robin Schucker, Michael A. Sontag, Jingfan Wang, Eric Weiner, and Alexander S. Yachanin. 2017. "100% Clean and Renewable Wind, Water, and Sunlight All-Sector Energy Roadmaps for 139 Countries of the World." *Joule*, 1(1): 108–121.

Joskow, Paul L. 1974. "Inflation and Environmental Concern: Structural Change in the Process of Public Utility Price Regulation." *Journal of Law & Economics*, 17(2): 291–328.

Joskow, Paul L. 1997. "Restructuring, Competition and Regulatory Reform in the U.S. Electricity Sector." *Journal of Economic Perspectives*, 11(3): 119–138.

Joskow, Paul L. 2011. "Comparing the Costs of Intermittent and Dispatchable Electricity Generating Technologies." *American Economic Review*, 101(3): 238–41.

Joskow, Paul L. 2013. "Natural Gas: From Shortages to Abundance in the United States." *The American Economic Review*, 103(3): 338–343.

Kellogg, Ryan. 2020. "Output and attribute-based carbon regulation under uncertainty." *Journal of Public Economics*, 190: 104246.

Kwoka, J., M. Pollitt, and S. Sergici. 2010. "Divestiture policy and operating efficiency in U.S. electric power distribution." *Journal of Regulatory Economics*, 38: 86–109.

Kwoka, John. 2008. "Restructuring the U.S. Electric Power Sector: A Review of Recent Studies." *Review of Industrial Organization*, 32(3): 165–196.

Lazar, Jim. 2016. "Electricity Regulation in the US: A Guide (Second Edition)." The Regulatory Assistance Project, Montpelier, VT.

MacKay, Alexander, and Ignacia Mercadal. 2024. "Do Markets Reduce Prices? Evidence from the U.S. Electricity Sector."

Mansur, Erin T. 2007. "Upstream Competition and Vertical Integration in Electricity Markets." *The Journal of Law and Economics*, 50(1): 125–156.

Novan, Kevin, and Yingzi Wang. 2024. "Estimates of the marginal curtailment rates for solar and wind generation." *Journal of Environmental Economics and Management*, 124: 102930.

Pepermans, Guido. 2019. "European energy market liberalization: experiences and challenges." *International Journal of Economic Policy Studies*, 13: 3–26.

Phadke, Amol, David Wooley, Ric O'Connell, and et al. 2020. "2035: The Report." Goldman School of Public Policy, University of California, Berkeley.

Roth, Jonathan. 2022. "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends." *American Economic Review: Insights*, 4(3): 305–22.

Stock, James H, and Daniel N Stuart. 2021. "Robust Decarbonization of the US Power Sector: Policy Options." National Bureau of Economic Research Working Paper 28677.

Wing, Coady, Madeline Yozwiak, Alex Hollingsworth, Seth Freedman, and Kosali Simon. 2024. "Designing Difference-in-Difference Studies with Staggered Treatment Adoption: Key Concepts and Practical Guidelines." *Annual Review of Public Health*, 45.

Wing, Coady, Seth Freedman, and Alex Hollingsworth. 2024. "Stacked Difference-in-Differences." *NBER Working Paper*, 32054.

Yozwiak, Madeline. 2023. "Calculating the realized investment returns of U.S. electric utilities." *Utilities Policy*, 85: 101684.