Do Mergers and Acquisitions Improve Efficiency: Evidence from Power Plants

Mert Demirer*

Ömer Karaduman†‡

October 27, 2022

Abstract

Using rich data on hourly physical productivity and five thousand ownership changes from US power plants, we study the effects of mergers and acquisitions on efficiency and provide evidence on the mechanisms. We find that acquired plants experience an average of 4 percent efficiency increase five to eight months after acquisition. Three-quarters of this efficiency gain is explained by increased productive efficiency; the rest comes from improved capacity management at the plant level and allocative efficiency at the portfolio level. Our findings suggest that acquisitions reallocate assets to more productive uses: we find that *high-productivity* firms buy *under-performing* assets from *low-productivity* firms and make the acquired asset almost *as productive as* their existing assets after acquisition. Finally, investigating the mechanism, the evidence suggests that acquired plants achieve higher efficiency through low-cost operational improvements rather than high-cost capital investments.

*MIT Sloan E-mail address: mdemirer@mit.edu

†Stanford GSB E-mail address: omerkara@stanford.edu

‡We thank Vivek Bhattacharya, Natalia Fabra, Bob Gibbons, Paul Joskow, Chris Knittel, Charlie Kolstad, Jacob LaRiviere, Ariel Pakes, Nancy Rose, Mike Whinston, Frank Wolak, and seminar participants at ASSA 2022 meeting and Princeton for excellent comments and suggestions. We thank Matt Barmack, Charlie Gates, Robert Kasle, Peter Rider, Jon Sepich, and Sang Ha Gang for insightful discussion on plant operations and mergers and acquisitions in the power generation industry.

1 Introduction

A fundamental issue in antitrust policy is the trade-off between the efficiency and market power effects of mergers. The increase in market power raises prices for consumers; however, potential efficiency gains can counteract this effect, making the net effect on welfare ambiguous. While there is an extensive literature on the price effects of mergers, we have limited evidence on how mergers affect efficiency. With little guidance from empirical evidence, studies analyzing competitive effects of prospective mergers often rely on hypothetical efficiency gains (Farrell and Shapiro (2010), Nocke and Whinston (2022)).[1]

A major challenge in analyzing the efficiency effects of mergers is distinguishing true efficiency gains from other factors, such as changes in market power, buyer power, and product quality. Due to the limitations of production datasets, most research has studied revenue-based productivity (TFPR), which is estimated from revenues and input expenditures, rather than quantity-based measures (Foster et al. (2008), Atalay (2014)). Using TFPR is particularly problematic in merger retrospectives because an increase in market power, buyer power, or quality decline raises TFPR without any efficiency gains. This makes it very difficult to identify true efficiency gains of mergers.

In this paper, we provide a detailed and large-scale analysis on the efficiency effects of mergers while tackling these issues. In particular, we ask: (i) Do mergers and acquisitions improve efficiency? (ii) What are the mechanisms? (iii) How do mergers reallocate assets between firms?

We focus on the US electricity generation industry between 2000 and 2020. Four distinct features of this industry and available rich data allow us to overcome the challenges of estimating the efficiency effects of mergers. First, we observe, at the hourly frequency, the physical quantity of output and the physical quantity of the largest single input, the consumption of fuel. With this unusually rich and high-frequency production data, we construct efficiency measures at the hourly level and analyze how they change around the time of acquisition. Second, electricity is a homogeneous product, ruling out potential quality changes that could confound our analysis. Third, the efficiency measure relies on accurate input and output sensor measurements rather than survey responses, as in many industries. Finally, and most importantly, the power generation industry experienced a significant number of mergers and acquisitions during our sample period. We identified

[1] As an example, consider these quotes from Nocke and Whinston (2022): "there is a clear need for much better evidence on the efficiency effects"; "we observe that the literature on efficiency effects of horizontal mergers is extremely limited"; "our reading of the current (meager) evidence in the literature"; "there is remarkably little solid empirical evidence on this point."

around 600 transactions in which about 5,000 generator units, corresponding to 95 percent of cumulative industry capacity, changed ownership between 2000 and 2020. These ownership changes exhibit significant heterogeneity based on transaction type, transaction size, firm type, and plant and market characteristics, which we use to study which merger characteristics predict efficiency effects and to test potential mechanisms.

Our analysis starts by employing a difference-in-differences (DID) estimator to compare the efficiency of acquired plants to those not directly involved in merger activity. Our first finding is that the efficiency of the acquired power plants increased by 4 percent on average after acquisition. The efficiency increase starts five months after acquisition and it reaches the new steady-state level after eight months, suggesting that it takes time for the new owner to implement changes required for efficiency improvements. Our calculations suggest that these efficiency improvements correspond to a total cost saving of 6 billion dollars, and the total cumulative decline in CO_2 emissions between 2000 and 2020 is roughly 50 million tons.

This clear evidence of efficiency gains from mergers is important. However, to inform merger policy and generalize the lessons from this industry to other industries, it is crucial to understand the underlying mechanisms that generate efficiency gains in a power plant. Motivated by this, we investigate which plant, firm, and transaction characteristics are correlated with efficiency gains and what potential mechanisms generate them.

Our first question is how mergers allocate assets between firms. There are two main theories on the efficiency effects of asset allocations with mergers. The first theory suggests a "high-buys-low" pattern (Jovanovic and Rousseau (2002)) in which acquisitions transfer assets from low-productivity firms to high-productivity firms. The second theory assumes a "like-buys-like" pattern (Rhodes-Kropf and Robinson (2008)), in which firms have no systemic productivity differences, but there are complementaries between assets and firms. According to this theory, assets are allocated to firms with a higher ability to utilize those assets. Quantifying the role of these theories in efficiency gains is important to understand whether mergers put assets to more productive use and contribute to aggregate productivity growth.

Our findings suggest that acquisitions reallocate assets to more productive uses: we find that *high-productivity* firms buy *under-performing* assets from *low-productivity* firms and make the acquired asset almost *as productive as* their existing assets after acquisition. Acquirers are, on average, one percent more productive than the target firms, and the target firms are selling their underperforming assets relative to their other assets. After acquisition, the productivity of the acquired asset goes up by 4 percent and becomes almost as productive as the existing assets in the acquiring firm's portfolio. This finding suggests that assets are allocated to firms with both relative and absolute advantages in utilizing those assets.

We then move to understand what mechanisms generate efficiency gains. A firm can improve the overall efficiency of a power plant with three distinct mechanisms: (i) increasing its productive efficiency, (ii) better allocating production dynamically, and (iii) allocating production more efficiently across plants (portfolio efficiency). We develop predictions for each of these mechanisms and test them empirically. The test for productive efficiency involves estimating firm-specific cost curves separately for pre- and post-acquisition periods and then comparing them. For dynamic efficiency, we study whether the production profile of acquired plants change post-acquisition. And finally, portfolio effects would suggest that existing plants of the acquirer in the same markets see efficiency improvements. After testing these predictions, we find that productive efficiency explains 75 percent of the total efficiency gain. Testing for dynamic and portfolio efficiency effects suggests evidence for these mechanisms; however, they play a minor role in explaining the total efficiency gain.

After establishing the role of productive efficiency, the next question is what firms do to improve productive efficiency. There are two alternatives: (i) low-cost process improvements, which involve adopting best practices and hiring more skilled personnel, and (ii) high-cost capital investments, which involve equipment upgrades. Process improvements indicate information transfers after acquisition; capital upgrades indicate liquidity constraints of the former owner. To distinguish between these two mechanisms, we augment our production data with data on plant managers, manager characteristics, non-fuel costs, and capital expenditures. This data allows us to study how important determinants of power plant efficiency change after acquisition. Starting with the manager data, we find that 55 percent of acquired power plants change managers within three months of acquisition. These managers are 5 percentage points more likely to have a master's degree and 4 percentage points more likely to have a bachelor's degree compared to non-merger manager changes. In contrast, we find no evidence of an increase in capital expenditures, non-fuel costs and number of employees after the acquisition. These findings suggest that the new owner of the power plant improves efficiency through operational improvements rather than high-cost capital investment.

As in all retrospective merger analyses, an important concern in our paper is the endogeneity of mergers. We include three additional analyses to address these concerns. First, we run placebo tests by looking at the efficiency effects of minority acquisitions and company name changes, finding no efficiency effects. Second, we run a battery of robustness tests and show that our results are robust to different specifications. Finally, we look at whether other important changes in the plant in the absence of mergers generate similar efficiency effects. For example, we look at how management changes in the absence of mergers affect efficiency and find that management change leads to a 0.8 percent efficiency gain, in contrast to four percent caused by mergers.

Although we believe our empirical setting is ideal for studying the efficiency effects of mergers, there are important caveats worth mentioning. First, the production process in electricity generation might be different from production in other industries in terms of variable cost structure and the role of labor. While we focus our analysis on a single industry to take advantage of the available data and numerous acquisitions, we provide detailed evidence for mechanisms to draw broader lessons from this study. Second, our efficiency measure is fuel efficiency rather than total factor productivity (TFP), which is most commonly used in the productivity literature. While it is possible to estimate TFP for power plants at the annual frequency, analyzing fuel efficiency provides a more detailed and complete picture due to the availability of high-frequency efficiency measures.

We conclude the introduction by highlighting that our results do not give a conclusive answer to the overall impact of mergers on consumer harm, as we identified only one factor going into the welfare analysis. Although more research is needed to understand the net effects of mergers, our paper provides a detailed analysis of the efficiency effects of mergers.

1.1 Literature

This article contributes to several bodies of literature. The first is the literature studying the effects of mergers and acquisitions on productivity. Since many merger retrospectives focus on price effects, there are only a few papers studying productivity effects of mergers (Braguinsky et al. (2015), Blonigen and Pierce (2016), Kulick (2017)). Blonigen and Pierce (2016) use the methods of De Loecker and Warzynski (2012) to separately identify markup power and productivity for manufacturing plants in the US and study how mergers affect them. Their findings suggest significant effects of mergers on market power but no evidence for a productivity effect. Kulick (2017) studies mergers in the ready-mix concrete industry. He finds evidence for price increase due to a rise in market power post-merger despite a 6 percent productivity increase in acquired plants.[2] Our paper is most closely

[2]Evidence from other industries include meat product industries (Nguyen and Ollinger (2006)), railroads (Bitzan and Wilson (2007)), healthcare (Schmitt (2017), Dranove and Lindrooth (2003), Harrison (2011)).

related to Braguinsky et al. (2015), who study the Japanese cotton spinning industry at the turn of the 20th century, which experienced a wave of acquisitions over 30 years. They find that acquirers were not more productive, but they were more profitable due to better inventory management and lower capacity utilization. After the acquisitions, the acquirer improves capacity utilization in the acquired plant, raising the productivity level by almost 13 percent.

This article also contributes to the literature studying efficiency in the power generation industry. This literature has primarily focused on how restructuring that started in the 1990s affected efficiency (Knittel (2002), Bushnell and Wolfram (2005), Davis and Wolfram (2012), Fabrizio et al. (2007)). These papers compared the performance of plants in states that pursued restructuring against plants in states that did not. Overall, the results point to a positive influence of restructuring on the operations of plants.

We contribute to the literature studying the allocative efficiency effects of mergers (Jovanovic and Rousseau (2008), Jovanovic and Rousseau (2002), Rhodes-Kropf and Robinson (2008), McGuckin and Nguyen (1995), Schoar (2002)). These papers study mergers and acquisitions across a range of industries. They investigate the characteristics of buyers and sellers, how acquisitions transfer assets between firms, and the effects of this on reallocations of resources in the economy. We contribute to this literature by providing detailed evidence from a single industry on how mergers allocate resources in the economy.

Finally, our paper is related to a recent wave of papers that use retrospective merger analyses to understand how mergers affect firm behavior. The insights from this growing literature advance the understanding of cross-market mergers (Lewis and Pflum (2017), Dafny et al. (2019)), monopsony power (Prager and Schmitt (2021)), buyer power (Craig et al. (2021)), quality (Eliason et al. (2020)), product availability (Atalay et al. (2020)), and the price effects of mergers (Luco and Marshall (2020), Bhattacharya et al. (2022)). We complement this literature by studying how mergers affect firm efficiency and providing evidence on the mechanisms.

2 Institutional Background and Plant Productivity

This section starts by providing an institutional background of the power generation sector and an overview of mergers and acquisitions in the industry. We then explain how to measure efficiency in a power plant.

2.1 Power Sector

The US electricity sector represents roughly 5 percent of the US GDP, with over 11,000 utility-scale power plants providing 2 to 3 million jobs across the US (Bradley Associates (2017)). Until the early 1990s, US electricity generation was overwhelmingly supplied by regulated and vertically integrated investor-owned utilities (IOUs) or government-owned utilities (municipal and state-owned). Typically, these entities served a specific territory and owned all parts of the power sector: transmission, distribution, and retailing. The returns of these utility services were regulated through rate-of-return on capital investments and cost-of-service regulation. This highly regulated market structure left little incentive for efficiency improvements, generating significant inefficiencies (Fabrizio et al. (2007), Cicala (2015)).

After the 1990s, the industry went through significant deregulation. Electricity generation was decoupled from transmission and distribution, and most generators began to earn profit through the market pricing system. This deregulation was accompanied by the creation of independent system operators (ISOs). ISOs organize the wholesale electricity market and meet electricity demand by running high-frequency auctions where power plants bid their willingness to produce. In 2020, roughly 70 percent of US electricity demand was provided through seven ISOs.[3] The deregulation also changed the electricity generation technology mix with a significant amount of plant entry and exit. In the early 1980s, coal was the primary fuel source for electricity generation. As the price of natural gas fell significantly with the expansion of fracking in the early 2010s, gas-fired generation became competitive with coal-fired plants, each providing roughly one-third of the market supply in 2015. In 2020, gas-fired generation reached roughly twice the size of coal-fired generation.

2.2 Mergers and Acquisitions

Large utility companies are usually organized into several subsidiaries under a big parent company, serving in different locations and segments of the power sector. The structure of the subsidiaries tends to follow the boundaries of the vertically integrated utilities before the deregulation. Parent companies typically own assets in generation, transmission, and distribution in the same region, with some parent companies having subsidiaries serving different parts of the country. After the deregulation, significant merger and acquisition

[3] The sevon ISOs are California ISO (CAISO), New York ISO (NYISO), Electric Reliability Council of Texas (ERCOT), Midcontinent ISO (MISO), ISO New England (ISO-NE), Southwest Power Pool (SPP), and Pennsylvania-New Jersey-Maryland Interconnection (PJM).

activity has occurred between these entities within and across regions. Moreover, financial firms, predominantly private equities and bank funds, started to invest in power generation.

Mergers and acquisitions in the power sector can be divided into three main groups: (i) asset sales, (ii) subsidiary acquisitions (divestitures), and (iii) mergers. In asset sales, a firm sells part of its power plant portfolio while maintaining its corporate structure. In this case, the acquired assets fall under the ownership of a subsidiary of the acquirer. In subsidiary acquisitions, a parent company acquires a subsidiary of another firm with all its assets. The plant's owner (subsidiary) remains the same in these cases, but the parent owner changes. The third type is a merger between two firms, where two companies merge and form a new company. Appendix Figure 25 provides a visual representation of these merger types.

All proposed plant acquisitions in the US electricity sector must be reviewed by the Federal Energy Regulatory Commission (FERC), Department of Justice (DOJ), and state Public Utility Commissions (PUC). The FERC reviews mergers under Section 203 of the Federal Power Act, relying on the 1996 revision of the Horizon Merger Guidelines (HMG) and putting more emphasis on market concentration levels. The DOJ's review is more extensive than FERC's, relying on the 2010 HMG and investigating potential anticompetitive effects of mergers. If either the DOJ, FERC, or state PUC concludes that the merger harms competition, they either block it or require assets divestiture on generation or distribution.[4] Despite extensive reviews by three government agencies, most proposed mergers have been approved in the US electricity sector, with only a few challenged mergers in the last two decades.

There are several merger motives in the electricity industry, and efficiency improvement in power plants is one of them.[5] Firms often argue that mergers will generate synergies between assets, citing enhanced financial flexibility, increased cash flow benefits, and complementarities between different sectors such as distribution and generation. As fuel represents roughly 80 percent of operational costs, many merging firms argue for potential operational efficiency improvements and synergies after the merger.[6] As an example,

[4] To give some examples, in 2005, the Exelon-PSEG merger was not completed after failing to get approval from state PUCs. In 2012, following the DOJ's request, Exelon Corporation and Constellation divested three generating plants in Maryland. In both cases, the FERC concluded that the merger would not harm competition.

[5] For most mergers in our sample, we have access to investor presentations and conference calls; therefore, we can identify the motives precisely.

[6] Some other examples: (i) AES-DPL merger, which argues that their earlier merger IPL led to improvements in management; (ii) NRG-GenOn merger, which cites measurable and actionable cost synergies of

Appendix Figure 24 shows a slide from the investor presentation of the 2018 Dynergy and Vista Energy merger, in which merging parties argue that heat rate improvements will lead to 125 million dollars in cost savings.

2.3 Electricity Production and Construction of the Efficiency Measure

A major challenge when studying the efficiency effects of mergers is the lack of suitable data because most industries do not have reliable measures of cost and physical productivity. The power generation industry is unusual in this respect because rich and high-frequency efficiency and cost data are publicly available. This section describes the efficiency measures used in this study and explains production in power plants.

A power plant is an industrial facility that generates electricity. In 2020, there were 11,070 utility-scale electric power plants in the US. A typical power plant include multiple generators, which transform a form of energy into electricity using different production technologies. Our research focuses on fossil fuel power plants, a type of thermal plants, because their efficiency is easier to measure. Fossil fuel power plants use energy of heat, obtained from burning gas or coal to make electricity.[7] In this process, the total input is measured as heat content of the fuel used in electricity generation. This leads to a natural efficiency measure, called *heat rate*, which specifies how much heat input is used to produce a given amount of electricity. Our main measure of efficiency is the inverse of this measure, defined as the ratio of energy output and input:

$$\text{Inverse heat rate} = \frac{\text{Energy Output (MWh)}}{\text{Energy Input (MMBtu)}}. \tag{2.1}$$

The heat rate characterizes the productivity of a generator, representing how efficiently fuel is converted into electricity. It is expressed as the ratio of the fuel's heat content, British thermal unit (Btu), and the plant's electricity output, megawatt-hour (MWh). Heat rate is a standard efficiency measure in the industry, widely used by regulatory agencies and firms.

A lower heat rate means more efficient production, as the plant generates the same MWh of electricity with less fuel. Thus, improving heat rates lowers fuel costs, which is

$175 million per year; (iii) Mirant-RRI Energy $150 million in annual cost savings; (iv) Vistra-Dynergy mentions geographic, fuel, market and earnings diversification benefits. Other cited reasons are increasing the consumer base, diversifying the portfolio across technologies and regions, and accelerating efforts to meet potential future environmental regulations.

[7] In a thermal power plant, water is heated in a boiler to generate steam, which is then moved through a turbine that attached to a shaft. As the steam moves, it causes the shaft to spin. This spinning shaft is connected to a generator, which produces electricity.

Figure 1: Heat Rate Curve

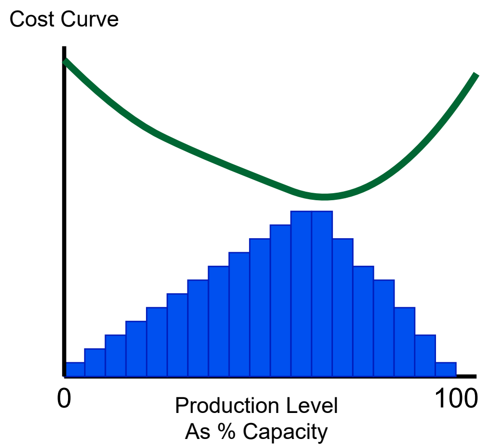

Note: The green line represents the marginal cost of producing electricity at a given capacity without potential ramping costs. The blue bars represent the distribution of actual production levels.

the major input representing roughly 80 percent of operating costs (Fabrizio et al. (2007)). Most importantly for our study, the fuel efficiency of a power plant provides an ideal setting to study the effects of mergers and acquisitions on productivity. First, our measure is a quantity-based efficiency measure obtained from quantity input and output, not affected by buyer or market power changes. Second, electricity is a homogeneous product, thus not affected by potential quality changes after acquisition. Finally, the efficiency measure relies on accurate sensor measurements of input and output rather than survey responses, as in many industries.

Several factors affect the heat rate in a power plant. Figure 1 shows an example of a heat rate curve, where blue bars show the production distribution as the percentage of capacity and the green line shows efficiency (inverse heat rate). First, the efficiency of a power plant varies with its production level. Production at high or low capacity leads to low efficiency. Second, power plants must rapidly adjust their production to respond to highly volatile demand, which requires technical expertise and can affect the power plant's overall efficiency. The associated cost with this adjustment is called ramp-up and ramp-down cost. Therefore, power plants whose production varies a lot tend to produce electricity less efficiently. These determinants of power plant efficiency depends on the skills and expertise of power plant personnel who monitor and control production (Bushnell and Wolfram (2009)). Finally, fuel type plays an important role. Coal-fired generators often have 10–12

MMBtu/MWh, whereas natural gas-fired generators have 7–9 MMBtu/MWh.[8]

Although the electricity generation process seems relatively mechanical, there is considerable heterogeneity in power plant productivity in the US. Figure 2 shows the distribution of yearly residual log productivity of power plants in the US after controlling for plant age, fuel type, technology, capacity, and other characteristics.[9] The difference between the 10th percentile and the 90th percentile is 0.38, indicating that plants in the top part of the productivity distribution are more than twice as productive as plants in the bottom part of the productivity distribution.[10] This heterogeneity in productivity has been observed by others in the literature (Sargent & Lundy (2009), Staudt and Macedonia (2014)).[11] The large dispersion in productivity conditional on a very rich set of observables highlights the role of unobserved heterogeneity in efficiency, indicating potential room for improvement in many power plants.

Improving heat rate is a complex process that can be done in two main ways: (i) low-cost operational improvements and (ii) costly capital upgrades. Low-cost practices, such as process optimization, personnel training, efficient maintenance, and avoiding ramp-up costs, can significantly improve the heat rate. Every year, power plant managers gather at the Heat Rate Improvement Conference to discuss these practices (EPRI (2022)).[12] An important determinant of operational practices is labor. As documented in detail in Bushnell and Wolfram (2009), individual skills of key personnel could make a significant difference in the performance of generating plants. Another way to improve plant efficiency is by upgrading key equipment, such as boilers, fuel feeders, and cooling systems, as the old equipment degrades and new technology becomes available.

Improving the efficiency of a plant is also important for environmental considerations. The more efficiently a plant operates, the less fuel it requires, emitting lower local pol-

[8] Even within a fuel type, there could be small differences in efficiency depending on the processing of the fuel. For example, using different coal types can lead to different efficiency levels due to factors like moisture content and pulverization procedures. Since we observe fuel and quality (for coal), we can control for changes in fuel types.

[9] We explain this estimation in Appendix B.1.

[10] This dispersion is slightly smaller than the dispersion identified in other manufacturing industries (Syverson (2011)

[11] For example, Sargent and Lundy's report commissioned by the EPA found that coal-fired power plants in the US require on average 10,400 British thermal units (Btu) to produce one kilowatt-hour (kWh) of electrical energy, but range from heat rates of 5 MMBtu/MWh to 32.7 MMBtu/MWh (Sargent & Lundy 2009). Staudt and Macedonia (2014) examine factors that contributed to heat rate using data from US EPA's National Electric Energy Database (NEEDS), including facility size, capacity factor, emission controls, steam cycle (supercritical versus subcritical), and coal type. They determined that each factor played an essential role in the generator's heat rate. However, they also determined that there was much unexplained variability in the data.

[12] Appendix Figure 25 highlights a few case studies of low-cost improvements.

Figure 2: Distribution of Residual Log Productivity

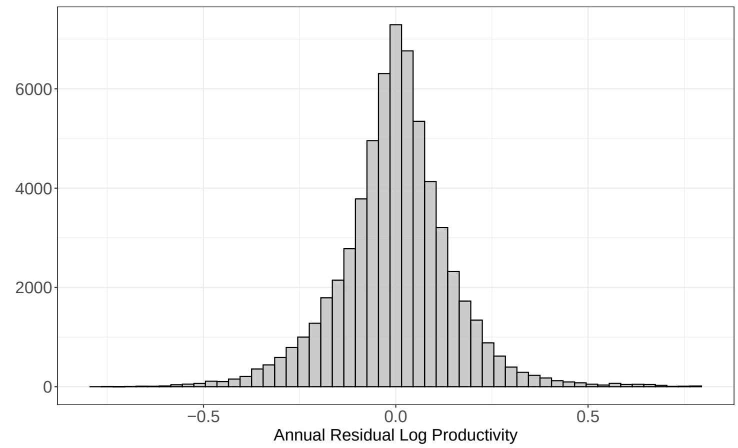

Note: This figure shows the distribution of residual yearly log productivity of fossil fuel plants in the US between 2000 and 2020 after controlling for plant age, fuel type, technology, capacity, and other characteristics.

lutants and GHG emissions. As a result, improving plant efficiency can be an effective tool to decrease local and global pollution. The potential role of increasing power plant efficiency has been recognized by policymakers in the US. In the 2016 Clean Power Plan Act introduced by the Obama administration, improving the heat rate of existing power plants was proposed as the first building block to reduce the carbon intensity of electricity generation (EPA (2018)).[13]

3 Data

Our primary goal is to construct an hourly measure of generator efficiency and the universe of ownership changes to examine the impacts of M&A on efficiency in electricity generation industry. An attractive feature of the power generation industry is that it has richer data on production and ownership than most industries. We take advantage of this and create a unique dataset on ownership and production.

We combine several datasets from the FERC, EPA, Department of Energy's Energy Information Administration (EIA), S&P Global, Velocity Suite, and S&P CapitalIQ Pro at the firm, plant, and generator level for all coal- and gas-fired power plants in the US between

[13]An analysis conducted by National Energy Technology Laboratory (NETL) supported this conclusion. Under a scenario where generation from coal is constant at the 2008 level, increasing average efficiency from 32.5 to 36 percent reduces US GHG by 175 MMmt/year, or 2.5 percent of total US GHG emissions in 2008. Moreover, NETL notes that "if each plant achieved their maximum efficiency each year, 5 percent reduction in CO_2 could result" (Campbell (2013)).

2000 and 2020. This results in close to a billion data points with the hourly generation at the generator level. This section briefly describes our process using several data sources. We provide a more detailed description of the data sources, construction of variables, and descriptive statistics in Appendix A.

Plant-Level Characteristics We use Velocity Suite, S&P Global, and EIA Forms 860 and 923 to construct detailed data on generator-level and plant-level characteristics for all fossil fuel power plants in the US. We assembled information on fuel type, capacity, regulation status, boiler model, and boiler manufacturer for generator units. For plants, we constructed data on plant age, location, ISO, and FERC region. For roughly half of the plants in the sample, we also have annual information on the number of employees, non-fuel costs, and capital expenditures between 2008 and 2020. Data provider Velocity Suite sourced this information from FERC Form 1, which is available only for investor-owned utilities.

Production and Efficiency Data We utilize the EPA's (EPA) Continuous Emissions Monitoring Systems (CEMS) for hourly generation and input data. The CEMS program was developed to systematically monitor power plant emissions for implementing environmental controls. It provides hourly power output, power input, emission, and heat rate of almost every fossil fuel power plant in the US.[14] We merged this dataset with the unit characteristics data from other datasets using generator names. In some cases, generator names in the EIA dataset and CEMS do not match.[15] We used EPA's Power Sector Data Crosswalk in those cases. Finally, we manually matched the retired and unmatched power plants using the principles described in the Appendix. We restrict our sample to all US fossil-fuel generators that comply with the CEMS program, except those in Alaska and Hawaii.

Mergers and Acquisitions Data We construct the ownership panel data for the universe of fossil fuel power plants from two separate ownership and transaction datasets obtained from S&P Global. The ownership data includes the ownership structure of all power plants at the subsidiary and parent company levels. We observe not only the majority owners but any firm that owns shares of the power plant. Transaction data provides detailed information about the transferred assets and transactions, such as acquired power plants, deal size, buyer, seller, announcement and close dates, conference call transcripts, and deal description. Since regulatory authorities must review all transactions, this data is available

[14]Every power plant in the US with more than 25 MW capacity that burns fossil fuel must comply with the EPA CEMS program. This sample represents approximately 95 percent of the US fossil-fuel generating capacity.

[15]The EPA uses boiler names as a unit, whereas EIA uses generator names.

for the universe of transactions during the sample period. Ownership and merger datasets often suffer from falsely identified ownership changes because firm name changes and restructuring of the parent company sometimes appear as ownership changes. We identify false ownership changes by cross-matching the plants in the transaction data with the ones in the ownership data. We also use corporate structure data to identify changes that are merely a restructuring within the same parent company. We use those false ownership changes for placebo tests. We also identified ownership changes that are forced divestitures due to deregulation in the early 2000's and excluded them from our sample.

Personnel Data Since plant personnel is an important determinant of efficiency, we assembled panel data on personnel information for each plant in our sample from the EPA. The EPA has this information because each power plant that complies with a EPA program must submit a plant representative to EPA. This data includes the representative's name, start and end date of their tenure, and contact information. To obtain more information about the personnel, we matched roughly 70 percent of them to their LinkedIn profiles, obtaining title, education, and employment history. From LinkedIn data, we confirmed that 80 percent of the reported personnel are plant managers, and the rest are mostly environmental compliance personnel. Therefore, we treat the personnel in our data as plant managers for the rest of the study.

Other Datasets To control for renewable generation, we gathered hourly data for solar and wind power generation from FERC and S&P Global for each Balancing Authority Area (BAA) and ISO Zone.[16] We also obtained information about firm characteristics, such as asset size, market cap, and industry information, from S&P CapitalIQ Pro.

4 Descriptive Statistics on M&A in the US Power Industry

This section presents descriptive statistics about mergers and acquisitions of fossil fuel power plants in the US electricity generation industry. Our goal is to demonstrate that the industry experienced a significant number of acquisitions with rich heterogeneity in terms of transaction size, acquirer and target firm types, and location. These facts allow us to study several aspects of how acquisitions affect efficiency, and they will be essential to keep in mind when we conduct our empirical analysis.

95 percent of Industry Capacity Changed Ownership between 2000 and 2020. There has been a large number of mergers and acquisitions in the US fossil fuel power generation

[16]Since demand-side variations lead to expected or unexpected changes in power plant production, controlling for demand shocks is critical to understanding plant efficiency.

Figure 3: Share of Capacity with Majority Owner Change

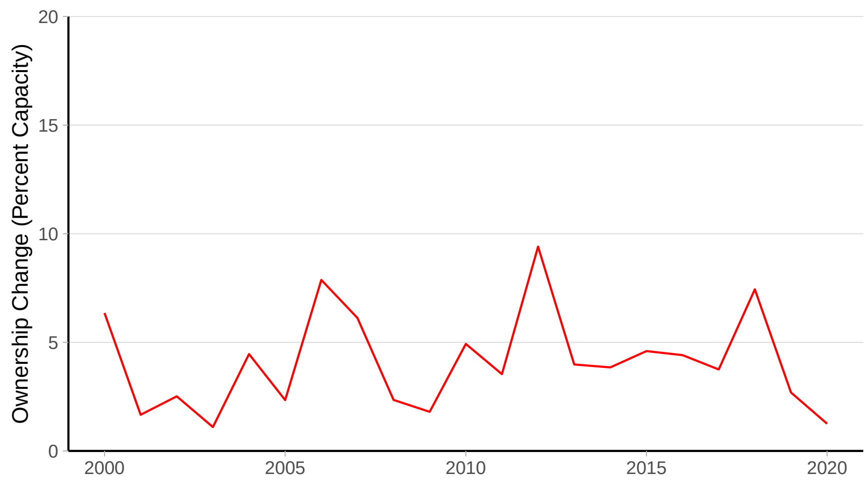

Note: This figure shows the change in percentage cumulative ownership of fossil fuel plants in the US between 2000 and 2020.

industry between 2000 and 2020. Figure 3 shows the percentage of fossil fuel electricity generation capacity that changed ownership from 2000 to 2020. We see that an average of 5 percent capacity changes ownership annually, with some fluctuations year-to-year.[17] Cumulatively, this corresponds to 95 percent of industry capacity during our sample period.[18] This is also reflected in the large number of power plants that changed ownership. Table 1 presents some summary statistics on plants, firms, and deal characteristics from the universe of acquisitions between 2000 and 2020. Our data includes 690 transactions involving 4,834 generation units and 1,567 plants. About 80 percent of these transactions involves a generator whose majority owner changes, giving us 4,030 generators with an ownership change. These generators will constitute our primary sample. Finally, looking at the firm characteristics, we see 267 unique acquirer firms and 266 target firms in the data.

Despite many acquisitions in our study period, we do not observe a meaningful change in market concentration. Appendix Figure 26 reports the national market shares of the largest 5, 10, 20, and 30 firms in terms of capacity owned. The concentration fluctuates

[17] We define acquisition as an ownership change if a different firm owns the majority of the plant's shares after the acquisition. For a small number of plants, no firm owns more than 50 percent of shares. For those plants, an acquisition is defined as the change of the largest shareholder.

[18] This is cumulative capacity, so it double-counts the capacity of generators that have changed ownership multiple times. We observe that 2,200 generators change ownership at least once, corresponding to 50 percent of the industry's total capacity.

Figure 4: Market Share by Firm Type

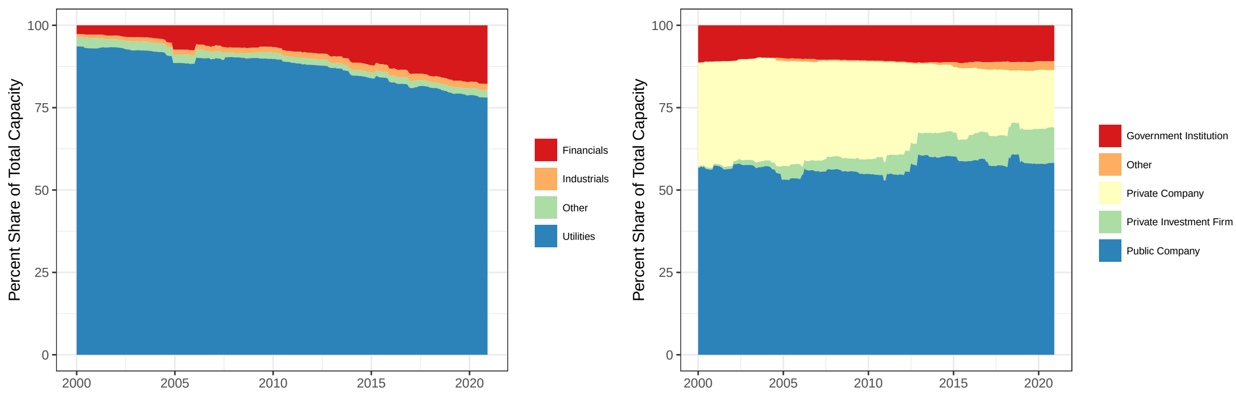

(a) Industry Type

(b) Owner Type

Note: These figures show change in market share of industry and owner types of fossil fuel plants in the US between 2000 and 2020

over time; however, it is broadly stable in the sample period.[19] This is because there is a considerable firm turnover in the industry, as suggested by the large number of acquirers and targets in Table 1. Some examples can be seen in Appendix Figures 17 and 18, where we report firms with the largest capacity increase and decrease between 2010 and 2020.

To show the composition of firms in the industry, Figure 4a displays the evolution of ownership by the primary activity of the parent company (utilities, industries, financials), and Figure 4b displays the evolution of ownership by whether the ultimate owner is a public firm. Looking at Panel 4a, we see an increasing presence of financial firms between 2000 and 2020 in the industry. The share of total capacity owned by financial firms goes from 3 percent in 2000 to 20 percent in 2020, suggesting substantial asset allocations from utilities to financial firms. Figure 4b highlights that public firms own most industry capacity, and the share of public firms remains stable over time. Finally, government institutions own 12 percent of the industry capacity. Except for the federally run Tennessee Valley Authority, these are local governments in rural areas that operate power plants that supply electricity to the public.

Most Acquisitions Reallocate Assets between Incumbent Firms. Most acquisitions are partial asset sales between two incumbent firms, but we also see some transactions where the target exits or the acquirer enters the industry. Columns 2 and 3 of Table 1 report the

[19]Note that these concentration ratios are not informative about the changes in market power due to the local nature of wholesale electricity markets. We report these changes at the national level to see whether large firms increase their dominance in this industry through acquisitions.

Table 1: M&A Summary Statistics

	(1)	(2)	(3)	(4)
	All	Acquirer Firm Enters Market	Target Firm Exits Market	Change in Majority Owner
	Unit Characteristics			
# of Units	4834	701	1692	4030
# of Plants	1567	268	482	1264
# of Unique Units	2365	585	1355	2198
# of Unique Plants	735	222	393	674
# of Acquirer Firms	267	126	107	234
# of Target Firms	266	99	148	229
% Gas	0.81	0.89	0.74	0.81
% Coal	0.13	0.09	0.18	0.13
% Oil	0.06	0.02	0.08	0.06
% Becomes Unregulated	0.04	0.06	0.03	0.04
% in Regulated State	0.24	0.25	0.39	0.24
% in ISO	0.77	0.76	0.71	0.76
	Firm Characteristics			
# of Units Target Owns	35.03	0.00	25.56	32.50
	(49.28)	(0.00)	(39.54)	(46.15)
# of Units Acquirer Owns	44.60	35.22	11.92	44.05
	(52.49)	(47.16)	(22.16)	(52.02)
Acquirer Firm Capacity	5459	0	4189	5055
	(8765)	(0)	(6883)	(8049)
Target Firm Capacity	7025	5312	1738	6912
	(9655)	(8314)	(3588)	(9453)
	Transaction Characteristics			
# of Deals	689	132	147	532
Deal Size in # of Units	7.0	5.3	11.5	7.6
	(14.7)	(8.9)	(21.8)	(16.0)
Deal Size in Capacity	1233	986	1922	1298
	(2640)	(1981)	(3798)	(2813)

Note: This table includes M&A activities that include fossil fuel-generating units in the US between 2000 and 2020. Each column reports the counts and characteristics in the data at varying levels of sample restrictions. Column (1) reports data from all acquisitions. Column (2) and Column (3) present data from transactions where the acquirer firm enters the market and the target firm exits the market, respectively. Column (4) reports data from majority acquisitions only.

Figure 5: Heterogeneity in Merger Size

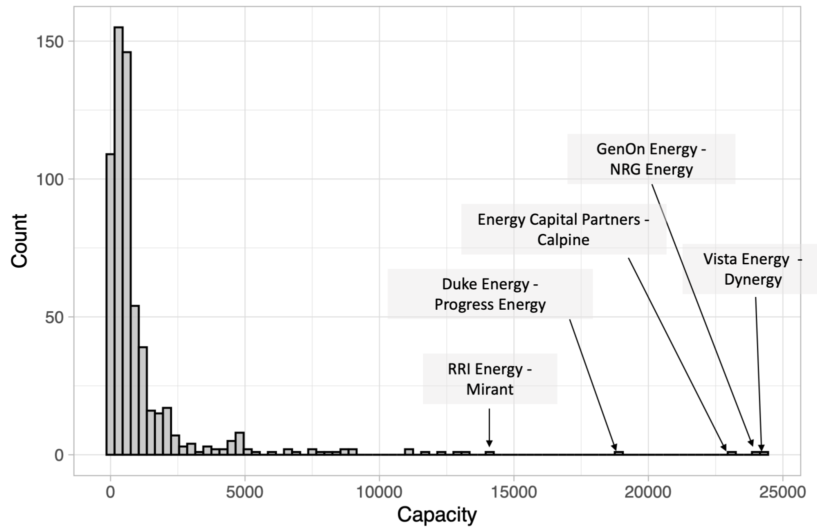

Note: This figure shows the differences in the size of deals that include fossil fuel plants in the US between 2000 and 2020.

summary statistics by transactions where the acquirer firm enters and the target firm exits the industry. The rest of the transactions occur between incumbent firms, which have a presence in the industry pre-and post-acquisition. In 20 percent of the transactions, the acquirer firm enters the market, and in about 20 percent of transactions, the target firm exits the market. Overall, these summary statistics suggest a significant reallocation of assets between firms in the industry, which allows us to test some important hypotheses about the allocative effects of acquisitions.

Heterogeneity in Transaction Size. Our sample includes mega-mergers that involve hundreds of generators and small transactions that involve only changes in minority owners. In Figure 5, we report the distribution of capacity that changed ownership across 689 transactions. While most small transactions involve one or a few plants with small capacity, some moderate-size transactions include ownership change of 5,000–10,000 MWh capacity. Finally, our sample includes mega-mergers that involve more than 10,000 MWh capacity. Observing this rich heterogeneity is useful because (i) our evidence does not come from a small number of large mergers, and (ii) we can test the heterogeneity of the effect by transaction size.

Ownership Changes at Different Level of Corporate Structure. Ownership changes can occur at two levels in a corporate structure: (i) the owner level (subsidiary) and (ii) the parent company level. Typically, a subsidiary of a holding company is the legal entity that

Figure 6: Ownership Change Types

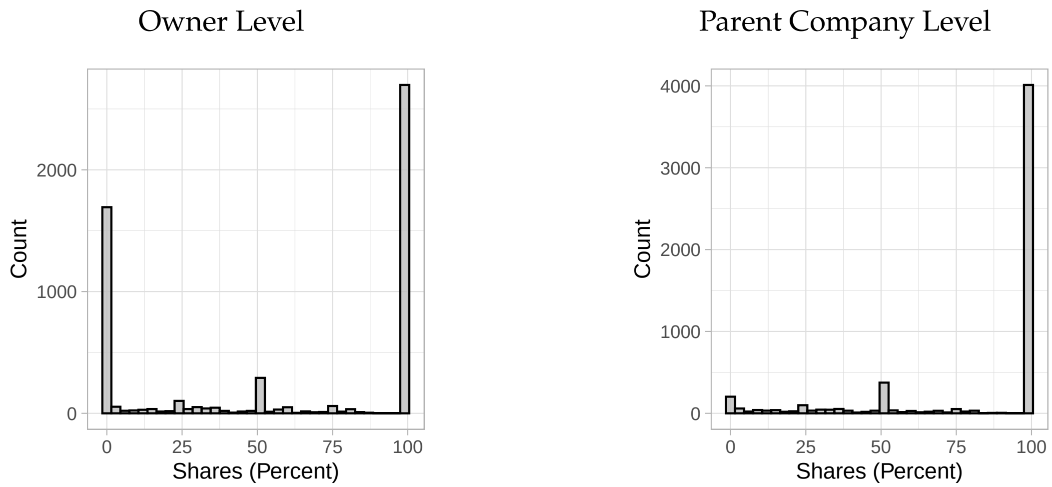

Note: This figure shows the differences in the ownership level of deals that includes fossil fuel plants in the US between 2000 and 2020.

owns the power plant, and a parent company owns that subsidiary. When a parent company acquires a subsidiary of another parent company, the entity that owns the power plant remains the same, but the parent company changes. On the other hand, if a subsidiary acquires a power plant in a partial asset sale, both the owner and parent company change. Figure 6a shows ownership changes at the owner and parent company levels. The figure suggests that in around 1,700 acquisitions, the owner remains the same and the parent company changes. We will use this variation to study the efficiency effects of acquisition at different levels of corporate structures. Moreover, in our data, we observe not only the majority of acquisitions, which involve more than 50 percent of the plant's shares, but also minority acquisitions, where the buyer acquires less than 50 percent of a power plant. The distribution of shares that change ownership is reported in Figure 6b. Perhaps surprisingly, there is a large number of minority acquisitions. We use these minority acquisitions in placebo tests, as one should expect no change in efficiency after a minority acquisition.

As explained in the previous section, some states regulate power plants' returns from power generation. One might be concerned about the role of regulations, as they might change merger motives and the incentive to improve productivity after an acquisition. To investigate this, Table 1 shows the fraction of acquisitions that occur in regulated vs. deregulated markets. We see that the majority of the ownership changes, 76 percent, occur in deregulated markets.[20] This is also reflected in the geographic variation of acquisitions (Appendix Figure 23). Another concern could be that ownership changes coincide with forced divestitures due to deregulation. Even though most state restructuring took place in the 1990s, some state restructuring overlaps with our sample period. For this reason, we look at how many ownership changes coincide with forced divestitures. Table 1 reports that only a handful of deregulation happened after 2000, which we exclude from our acquisition sample.

5 Empirical Results

Our empirical strategy aims to identify the causal effect of acquisitions on power plant productivity and study potential mechanisms. To do this, we compare productivity trends at acquired generators to those not directly involved in merger activity; we refer to these generators as "control generators" or "controls." In most estimations, each observation is a unique combination of generator and week, with variables containing productivity, ownership, and several units and plant characteristics.

The main advantage of our empirical setting is that we observe a high-frequency measure of the productivity of power plants, allowing us to track the productivity immediately before and after acquisition. As a result, we can identify the effects of acquisition within a short time window, not years before or after acquisition. This is in contrast to data availability in most manufacturing industries, where production data is typically at the annual level. This unique data feature provides an important advantage for identification because we can consider acquisitions as discrete events.

Before estimating our model, we make several additional sample restrictions. First, we eliminate generators that are inactive more than 90 percent of the time during their lifetime.[21] Second, we drop acquisition events that correspond to forced divestitures due to deregulation. Third, we require that target generators have at least one year of data before and after acquisition so that all in-sample acquired plants contribute to identifying variation in both the pre-merger and post-merger effect coefficients. Fourth, we focus only on the first acquisitions if a unit is acquired multiple times (half of all treated units). We remove the observations of units after the post-treatment period if they are acquired multiple times so that other acquisitions are not included in the sample. Fifth, we remove all treated units acquired again within a year of the first acquisition. We provide several robustness checks for these sample restrictions in Section 8.

[20] This is expected because, in most regulated markets, public ownership serves the local population.

[21] Lifetime is defined as the period between the first and last time of production observed in the data.

We find that acquisitions increase the productivity of power plants 4 percent, but only when ownership changes both at the parent company and owner level. In contrast, ownership changes at the parent company level only do not lead to a significant productivity increase. The productivity increase starts five months after the acquisition and reaches the new steady-state level after eight months. After documenting the evidence on efficiency, we close the section by studying the heterogeneity of the effects. To facilitate the exposition, we defer the detailed examination of mechanisms to Section 7.

5.1 Mergers and Efficiency

This section presents our main difference-in-differences results from estimating the effects of mergers on efficiency. To do this, we follow Braguinsky et al. (2015) and estimate a regression of the following form:

$$\log(y_{it}) = \theta_1 \text{l\_pre}_{it} + \theta_2 \text{e\_post}_{it} + \theta_3 \text{l\_post}_{it} + X_{it} + \mu_t + \alpha_i + \eta_{it}, \tag{5.1}$$

where y_{it} is the efficiency of generator i at week t (measured as inverse heat rate given in Equation (2.1)); the controls, X_{it}, include state-month fixed effects, time-varying generator characteristics such as age and fuel type (for coal), capacity and indicators for whether the unit is connected to the grid and whether it is an internal generator. α_i is generator fixed effect and μ_t is week fixed effect. By controlling for state-month fixed effect, we flexibly account for changes in demand and the supply of non-fossil fuel electricity generators (mainly entry of renewables) at the state level. Although it happens rarely, generators can change their capacity and fuel type; we include fuel type and capacity to control for these cases. Including generator and week fixed effects implies that merger effects are identified within generator changes following a merger event.

The regression includes three variables of interest: (i) l_pre, an indicator variable for 1 to 5 months pre-treatment, (ii) e_post, an indicator variable for 1 to 5 months post-treatment, and (iii) l_post, an indicator variable for 6 to 10 months post-treatment. By including early and late post-acquisition treatment indicators, we aim to capture the dynamic effects of mergers and identify when efficiency changes happen. We include l_pre to see whether there are any productivity effects of the acquisition before the acquisition. This could happen due to anticipation effects or disruption in the production process, as most acquisitions are announced months in advance. Finally, we cluster all standard errors at the plant level.

It is important to highlight that our unit of analysis is a generator rather than a plant.

Table 2: Balance Table

	Acquired	Never Acquired
Coal Unit	0.18	0.26
Gas Unit	0.74	0.67
Oil	0.08	0.07
Capacity	179	176
Install Year	1987	1985
Unit in ISO	0.75	0.58

Note: This figure shows the differences acquired and never acquired fossil fuel plants in the US between 2000 and 2020.

While the same firms typically own all generators in a plant, they might have different production profiles and maintenance schedules, which would affect efficiency estimates if inputs and production are aggregated at the plant level. Moreover, in a few ownership changes, we see that only some of the units in the plant change ownership. Therefore, we think that generator is the right level of analysis and maintain this throughout the paper.

Before moving to the results, in Table 2 we report the average characteristics of acquired power plants and power plants that have never been acquired. Treated and control groups look similar in terms of capacity and install year. There are slight differences in fuel type: acquired plants are more likely to be gas-powered than the control group. This difference is expected because there is substantial policy uncertainty about the future of coal power plants, which deters potential buyers. Finally, regarding geographical distribution, we see that treated and control groups are similar.

Table 6 shows the results from estimating Equation (5.1) for the outcome variable of log productivity. It does so for the entire sample of acquisitions (all M&A) as well as different acquisition events: (i) both the parent company and owner firm change; (ii) the parent company changes and the owner remains the same, (iii) minority acquisitions; and (iv) owner name changes. Estimates from (iii) and (iv) provide placebo tests, as discussed later.

The results in the first numerical column of 6 indicate the efficiency rises after the acquisition. The efficiency change in the early acquisition period is only 0.7 percent and marginally statistically significant, underscoring that there is no considerable change in efficiency immediately after acquisition. However, in the late post-acquisition period, 6 to 10 months after acquisition, the efficiency increases by 2 percent above the pre-acquisition level. Thus, acquired generators' efficiency levels improve following acquisition, though

Table 3: Regression Results

	All M&A	Owner/Parent Company Changes	Only Parent Company Changes	Minority Owner Changes (Placebo)	Name Changes (Placebo)
	(i)	(ii)	(iii)	(iv)	(v)
Late pre-acquisition	0.002	-0.003	-0.003	-0.004	-0.007
	(0.005)	(0.007)	(0.007)	(0.008)	(0.006)
Early post-acquisition	0	0.005	-0.002	-0.008	-0.004
	(0.005)	(0.007)	(0.007)	(0.020)	(0.043)
Late post-acquisition	0.014	0.039	-0.006	0.001	0.007
	(0.006)	(0.012)	(0.007)	(0.01)	(0.01)
Adj. R^2	0.622	0.635	0.622	0.652	0.635
# of Obs.	1.79M	1.38M	1.4M	1.12M	1.22M
# of Acq.	1760	897	921	405	456
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: This table presents the coefficient estimates from estimating Equation (5.1). Standard errors are clustered at the plant level.

it takes time for this to fully manifest.

Columns (ii) and (iii) of 6 test whether different types of ownership affect efficiency differently. Column ii estimates a difference-in-difference specification, where treatment is defined as ownership change at the parent company and owner (subsidiary) levels. In contrast, in Column (iii), a plant is treated when the parent company changes, but the owner firm remains the same after acquisition. One might expect the efficiency effects of an acquisition to be different in these two cases because the owner firm typically has direct control over the operation and personnel of the power plant. In contrast, the parent company controls the power plant indirectly.[22] Comparing these two columns reveals significant heterogeneity in the treatment effect. There is no effect when the owner company remains the same, whereas it is 4 percent if both the owner and parent company change. These results confirm our intuition that, for operational changes in the power plant, the direct owner plays a more important role than the ultimate parent. A change in indirect control does not affect power plant productivity.

The last two columns of the table serve as placebo tests. For the first placebo test,

[22]Moreover, ownership changes at the parent level tend to be financial acquisitions, in which the motive is to achieve diversification or safe returns.

we use the 512 minority acquisitions in the data. These are the transactions where the majority owner remains the same after an acquisition. For the second placebo test, we identify 612 generators whose owner changes their name, but the generator is not involved in an acquisition. These are "false" ownership changes that would have been classified as "acquisitions" from the ownership data, but cross-checking with the transaction data reveals no ownership change. We use these two types of events as placebo tests because we should expect no impact on power plant efficiency. The results confirm our expectation that we do not see any significant change in power plant efficiency after these events. These placebo tests give us confidence that confounders do not drive our main results.

To interpret the results from the specification in Equation (5.1) as causal, we rely on the assumption that an acquisition creates a discontinuous change in power plant behavior. In contrast, any efficiency trends that might lead to selection would be either common to the control plants or gradual enough to be distinguished from the more discrete direct effect. This assumption is likely to hold in our setting because we observe production at short intervals and include a rich set of control variables that account for industry-specific efficiency changes.[23] Still, there could be unobservable factors that would change efficiency in the absence of mergers, and those are observed by the acquirers affecting their acquisition decision. For example, the acquirer might observe that the target plant's manager will retire and decide to buy the plant, anticipating that the new manager will improve efficiency. To test whether the effect comes from manager changes, we estimate the effects of manager changes on efficiency in the absence of mergers and find that the efficiency increase is only 0.6 percent (Appendix Figure 19). Finally, we do a battery of robustness checks that are presented in Section 8, including matching estimators, Goodman-Bacon (2021) estimator, and estimation with daily and hourly data, to show that our results are robust to several specification choices.

After showing the significant impact of acquisitions on efficiency, we turn to the dynamic effects to identify more precisely when efficiency change occurs and whether there are any differences in pre-trends. For this purpose, we plot power plant efficiency during the period around acquisition using the following regression specification:

$$\log(y_{it}) = \sum_{s\in(-16,16)} \delta_s D_{t-s} + X_{it} + \mu_t + \alpha_i + \eta_{ijt}, \tag{5.2}$$

where D^{t-s} is a dummy variable for generator i being acquired at month t and controls are

[23] The average power plant efficiency gain in the industry is 0.3 percent. Therefore, in power plant generation, we do not see a significant fluctuation in efficiency.

Figure 7: Impact of Merger on Productivity

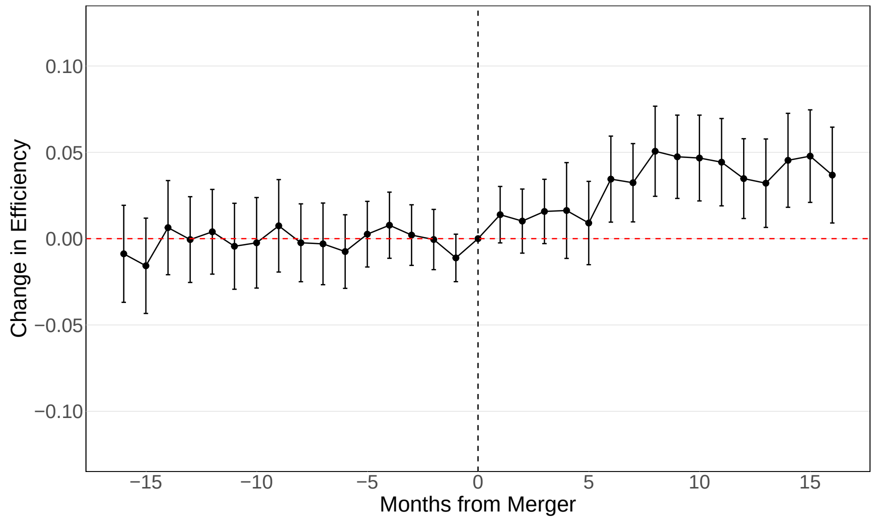

Note: The dynamic effects of acquisitions estimated from Equation (5.2). Standard errors are clustered at the plant level.

the same as the regressions estimated above, including generator fixed effects. Because we identify the significant efficiency impact of ownership changes at the owner and parent company level, we study only those acquisitions hereafter.

The results from the dynamic effect regression in Figure 7 suggest no significant evidence of pre-treatment trends. For the treatment group, the coefficients on $t \in (-16, 0)$ are small and insignificant, suggesting that acquired plants do not have a different efficiency trend before acquisition than those not acquired. The difference remains close to zero until five months post-acquisition. And after that, the efficiency of acquired plants and control plants diverges. The cumulative efficiency increase is 4 percent. The fact that plants experience efficiency gain five months after the acquisition suggests that it takes time for the new owner to change power plant operations necessary for efficiency improvements.[24]

5.2 Discussion

How large is the average 4 percent efficiency gain after acquisition, and what are the corresponding cost savings? To understand how large our finding is, we compare it to average within-plant productivity growth in the industry. In power generation, the contribution

[24]This finding is suggestive of how the efficiency gain occurs. Our interviews with power plant managers indicated that five months is not long enough to make costly capital investments and upgrades. This suggests that efficiency improvements occur primarily due to operational changes and adopting best practices rather than costly capital investments. We will return to this question later for a more formal analysis.

of within-plant productivity growth to aggregate productivity is small, as most of the productivity growth comes from plant entry and exit. We show this in Appendix Figure 22, which plots the average year-to-year within-plant productivity growth for the plants that were not involved in an acquisition. The within-plant productivity growth fluctuates around 0, with an average 0.3 percent annual increase over the sample period. This small average productivity growth makes the efficiency gain after acquisitions more striking because acquired plants achieve the level of productivity increase that an average plant achieves in twelve years without mergers.

Next, we quantify the cost saving due to the efficiency increase after an acquisition. We calculate this number under three simplifying assumptions; (i) acquired plants will keep the efficiency level after the acquisition; (ii) acquired plants will maintain the same production level that we observed pre-acquisition; and (iii) there is no redistribution of production across plants. Under these simplifying assumptions, we find that the total cost saving is six billion dollars. However, whether these cost savings would be passed to consumers is more complicated. It depends on whether the acquired plants are marginal or infra-marginal and how much market power they have. Although we think this is an important point, it is outside the scope of this paper.

Finally, we also aim to estimate the social gains due to positive externalities resulting from a decline in fuel usage in electricity generation. For this estimation, we assume that (i) CO_2 emissions are linearly increasing with heat rate; (ii) CO_2 emissions a roughly 0.4 tons per MWh for gas power plants, and 1 ton per MWh for coal power plants. Under these assumptions, the total cumulative decline in CO_2 emission between 2000 and 2020 is roughly 50 million tons. This corresponds to emission reduction from replacing 125 TWh gas power plant with renewables.

5.3 What Predicts Efficiency Gains: Heterogeneity Analysis

Estimating the average effects of past mergers is important to understand the overall impacts of mergers. However, to draw broader lessons from this industry and guide merger policy, it is crucial to learn the underlying mechanisms that lead to post-merger changes and understand what merger and firm characteristics predict them. Thus, our next set of analyses examines which characteristics of mergers predict merger effects. Results from this section are important because they can help anti-trust authorities predict which mergers will lead to more efficiency gains.

Our rich data will allow us to study the relationship between efficiency gains and several plant, firm, and transaction characteristics. For this purpose, we estimate the following regressions:

$$log(y_{it}) = \theta_1 D_{it} + \theta_2 D_{it} \times Z_{it} + X_{it} + \mu_t + \alpha_i + \eta_{it}, \tag{5.3}$$

where D_{it} is an indicator variable for treatment and Z_{it} is a plant, firm, or transaction characteristic. We estimate this equation for many plant, firm, or transaction characteristic separately and report the estimates of θ_2. The details of the estimation procedure for the heterogeneity analysis and how sub-samples are constructed are provided in Appendix C.

We begin by considering five plant characteristics: fuel type (natural gas or coal), age, regulation status, capacity, and whether the plant is infra-marginal or marginal.[25] Looking at the results, we do not detect heterogeneity by fuel type, mostly because a small number of coal power plants have been acquired during our sample period. However, there is significant heterogeneity based on other power plant characteristics. The efficiency increase is higher in older power plants. This is reasonable because there is degradation in performance over time, and therefore there is more room for efficiency improvements in older plants. The efficiency improvements are also higher if the plant is unregulated, larger, or infra-marginal. In all of these cases, the owner has more incentive to improve power plant efficiency.[26]

Next, we look at which firm characteristics predict efficiency gains. We focus on transactions where the target firm exits the market, the acquirer enters the market, the acquirer is a financial firm, the acquirer is large, and it is a serial acquirer. In this specification, we do not see any heterogeneity by target firm exiting, acquirer firm entering, and the acquirer being a financial firm. However, efficiency improvement is 0.02 percent higher when the acquirer is large (in total capacity) and is 0.05 percent larger when the acquirer is a serial acquirer. These results are consistent with the interpretation that a firm's experience in plant operations and acquisitions is an important predictor of efficiency increase after acquisition.

Finally, we study five transaction characteristics: deal size by deal value (which includes non-power plant assets), whether the transaction occurs after 2010, the acquirers' existing capacity in the market, transaction size by acquired fossil fuel capacity, and whether the transaction is a bankruptcy sale. First, we see that larger transactions by deal

[25]Details about the heterogeneity variables are provided in Appendix B.2. The coefficient estimates can be found in Appendix Tables 7-9.

[26]This is because in unregulated plants, any cost-savings will be retained as profit; for infra-marginal and larger plants, production is higher, so any efficiency improvement would lead to a higher return.

Figure 8: Heterogeneity of the Merger Effects on Productivity

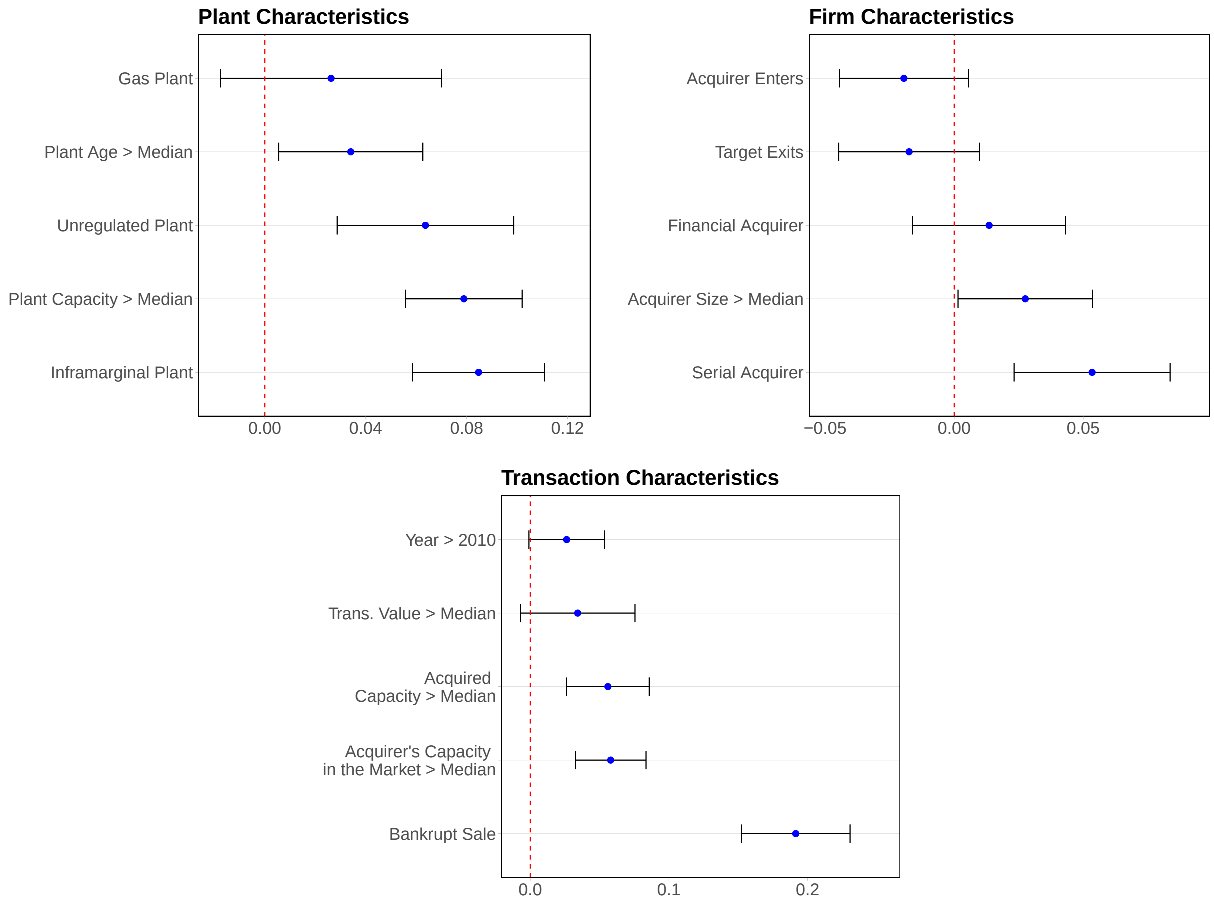

Note: Estimates of θ_2 from Equation (5.3) along with the 95 percent of confidence intervals. Standard errors are clustered at the plant level. Details about the heterogeneity variables are provided in Appendix B.2. The coefficient estimates can be found in Appendix Tables 7-9.

size do not lead to higher efficiency increase. Transactions after 2010 lead to slightly higher efficiency, suggesting that the results are not specific to deals that occur in a particular time frame. The acquirer's overall capacity and the capacity in the market are correlated with higher efficiency improvements. Finally, we see that efficiency improvements are significantly larger, close to 20 percent, when the transaction is due to bankruptcy sales.

The analysis in this section suggests that heterogeneity in efficiency effects can be explained by many plant, firm, and deal characteristics. The direction of these effects is consistent with firms having more incentives to improve the efficiency and experience of the firm. We would also like to note that these findings should not be interpreted as causal as, in many cases, the effect can be explained by other factors.

We believe that the results from this section are valuable for merger policy. Predicting the efficiency effects of mergers through counterfactual simulations is particularly difficult as most merger simulations focus on predicting price effects. Therefore, evidence on the efficiency effects of mergers conditional on a merger's attributes provides valuable information to assess which mergers would lead to efficiency gains ex-ante.

6 Do Mergers Allocate Resources Efficiently?

Mergers and acquisitions represent a significant source of reallocation in the economy and account for vast flows of resources between firms. This reallocation can generate allocative efficiency gain in the economy if the assets are allocated from less productive firms to more productive firms or acquirers utilize the acquired assets more productively.[27] Our empirical setting provides an ideal opportunity to study the allocative efficiency effects of mergers because we observe hundreds of asset reallocations between incumbent firms in our sample. Motivated by this, this section investigates whether (i) acquirers are more productive than target firms and (ii) acquirers have a comparative advantage in utilizing the acquired assets over the target firm.

There are two main theories on how acquisitions raise aggregate productivity through resource reallocation. The first, the Q theory of mergers (Jovanovic and Rousseau (2002)), posits that there are inherent productivity differences between firms, and acquisitions transfer resources from low- to high-productivity firms. This implies a "high-buys-low" pattern in the merger market. According to the second theory, proposed by Rhodes-Kropf and Robinson (2008), there are no systemic productivity differences between firms, but assets and firms could be complementary. Therefore, firms could have different levels of ability to operate different assets. This implies a "like-buys-like" pattern, as we expect to see acquisitions of complementary assets. A body of literature has tested these theories of merger gains (Maksimovic and Phillips (2001), Jovanovic and Rousseau (2008)) without conclusive empirical evidence.

The answer to these questions is an important input to quantify efficiency gains in merger analysis. In merger simulations of two firms with different marginal costs, the question of what should be the marginal cost post-merger is an empirical one. In other words, we need to know whether efficiency is transferable. It is well known that organizational challenges for integrating merged firms and dis-economies of scale at the senior

[27] Another view of the literature suggests that most acquisitions are undertaken for other motives, such as empire-building and managerial hubris.

Figure 9: Efficiency of Acquirer and Target's Plants

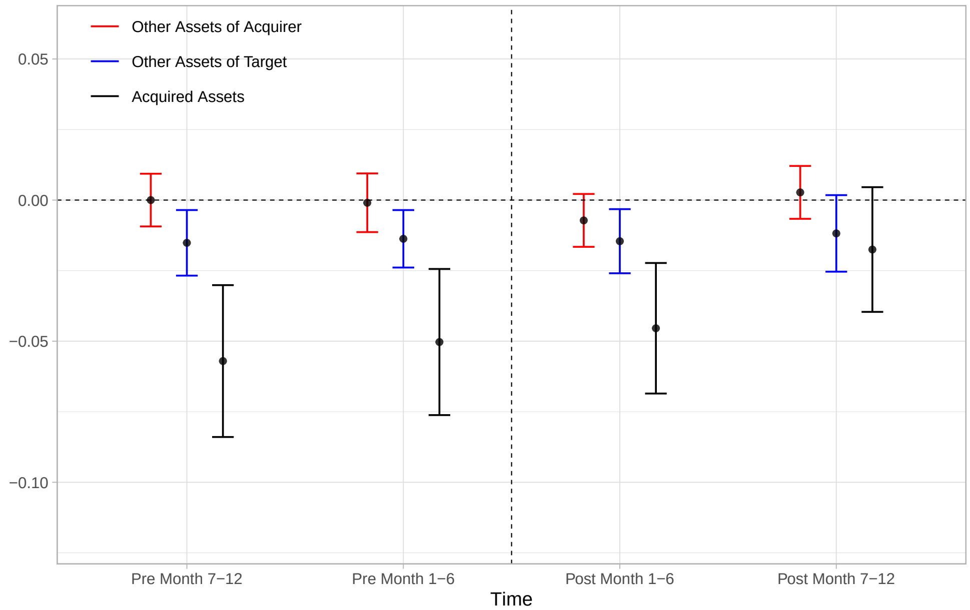

Note: Regression estimates from Equation (6.1). Error bars indicate 95 percent confidence intervals. Standard errors are clustered at the plant level.

management level could prevent firms from transferring best practices that improve productivity. In our empirical analysis, we also provide evidence for this important question.

To understand how mergers allocate resources in the economy, we estimate a difference-in-differences specification where we study the efficiency of three different types of assets: (i) acquired plants, (ii) plants of the acquirer not subject to the transaction, and (iii) plants of the target not subject to the transaction. In particular, we estimate the following specification:

$$log(y_{it}) = \sum_j \left(\theta_1^j \text{e\_pre}_{it}^j + \theta_2^j \text{l\_pre}_{it}^j + \theta_3^j \text{e\_post}_{it}^j + \theta_4^j \text{l\_post}_{it}^j\right) + X_{it} + \mu_t + \eta_{it}. \tag{6.1}$$

This regression aims to estimate both the level and change of efficiency, separately for the target's assets, acquirer's assets, and acquired assets around the time of acquisition. Equipped with these estimates, we can compare the efficiency of the target and acquirer's existing assets and identify how they perform relative to acquired plants. This regression does not include generator fixed effects because we are interested in estimating level differences, not only changes. We deal with potential endogeneity concerns due to the lack of generator fixed effects by including a very rich set of controls, including age, boiler type, fuel type, boiler model, boiler manufacturer, and plant capacity. We restrict the sample to acquisitions where both the acquirer and target firms own plants pre- and post-acquisition that are not subject to the transaction. We also normalize the efficiency level of acquirers' assets to zero in the early pre-acquisition period, so all other coefficients are estimated relative to θ_1^1.

Figure 9 reports the estimates of three sets of coefficients. The red, blue, and black colors, respectively, indicate the change in the efficiency of the existing assets of the acquirer, the existing assets of the target, and the acquired assets. First, comparing the efficiency levels of acquirer and target's assets reveals several interesting findings. First, we see that the productivity levels of both the target and the acquirer's existing plants are roughly constant around the time of the acquisition.[28] Second, comparing the efficiency level of the acquirer and target, we see that the acquirer is 1 percent more efficient than the target firm. These estimates suggest that assets are allocated from high-productive to low-productivity firms; however, the productivity differences are small.

We next compare the efficiency level of acquired assets with the other asset types. The first observation is that the assets sold by the target firm are underperforming relative to other assets in the target's portfolio: the acquired asset's efficiency is 4 percent lower than the average efficiency of other assets in the target's portfolio. This suggests some heterogeneity in within-firm plant productivity. More interestingly, it also shows that target firms sell their underperforming plants.

What happens to these underperforming assets after the acquisition? The efficiency of these assets improves after acquisition with an increase of 3 percent. This effect is similar to what we found in the previous section; however, it is estimated less precisely due to the decline in sample size. When we compare the post-merger efficiency of the acquired assets with the acquirer's portfolio, we see that the acquired assets become almost as efficient as the acquired firm's other assets.

Overall, the empirical findings in this section suggest that *high-productivity* firms buy underperforming assets of *low-productivity* firms and make the acquired asset almost as productive as its existing assets after acquisition. These results provide clear evidence for the two merger efficiency gain hypotheses discussed above. In particular, we find evidence for the high-buys-low pattern, as the acquirers are more efficient than the targets. We also find evidence for complementary asset theory in that acquirer firms have the ability to improve the productivity of the underperforming assets of the target firm. These results

[28] This result provides evidence against the target- or acquirer-level unobservables that lead to the selection.

overall are strong evidence for the allocative efficiency effects of acquisitions. Because the acquired assets are underutilized under the ownership of the target firm and improve performance under the ownership of the acquirer, acquisitions potentially contribute to aggregate productivity in the power generation sector.

7 Mechanisms

Our results so far uncovered large improvements in the efficiency of acquired plants after acquisition. However, if we want to draw broader lessons from this industry, we need to understand the mechanisms that generate efficiency gains. This section investigates the potential mechanisms of efficiency gains in power plants and provides two major findings. First, we find that most efficiency gains come from improvement in productive efficiency within a generator. Second, firms achieve these efficiency gains with operational improvements rather than costly capital investments.

7.1 Mechanisms of Efficiency Improvements

We propose three mechanisms that could generate the efficiency increase identified in the previous section: (i) increase in productive efficiency, (ii) dynamic efficiency, and (iii) portfolio (allocative) efficiency. We first explain these mechanisms and then develop an empirical prediction for each mechanism that can be tested in the data. We find evidence consistent with each of the three mechanisms, but most of the efficiency gain comes from an increase in productive efficiency.

Increase in Productive Efficiency. The first mechanism that could generates efficiency change is productive efficiency. Productive efficiency arises when the plant's new owner adopts operational processes that lower production cost or invests in new equipment. This mechanism does not rely on synergies with other plants in the same market; it arises because the new owner knows how to operate the plant more efficiently. An implication of productive efficiency is the decline in the cost curve at every production level, illustrated in Figure 10. Based on this implication, a distinct prediction of this mechanism is:

Prediction 1: If firms increase overall efficiency by improving productive efficiency, the cost curve of production will shift down.

Dynamic Efficiency. The second mechanism that could improve efficiency is dynamic efficiency, which arises due to better allocation of production over time. An important feature of power generation is that efficiency depends on the level of production as well as the

Figure 10: Productive Efficiency

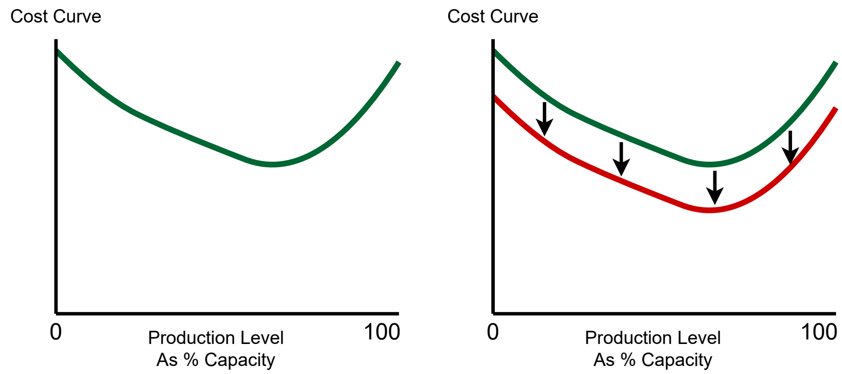

(a) **Pre-merger** (b) **Post-merger**

Figure 11: Improved Capacity Utilization

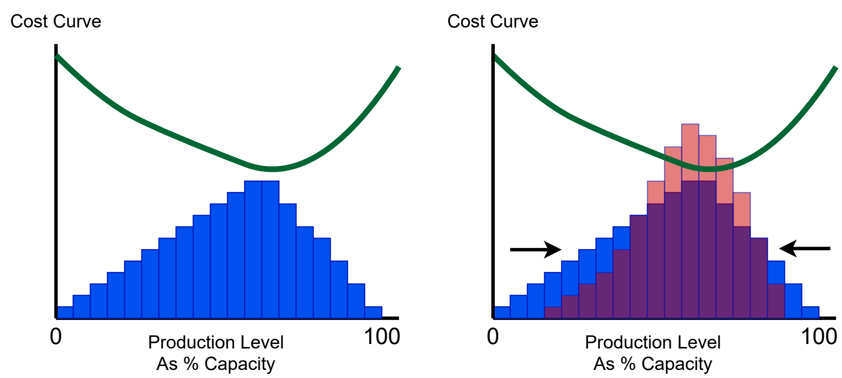

(a) **Pre-merger** (b) **Post-merger**

Figure 12: Portfolio Effects

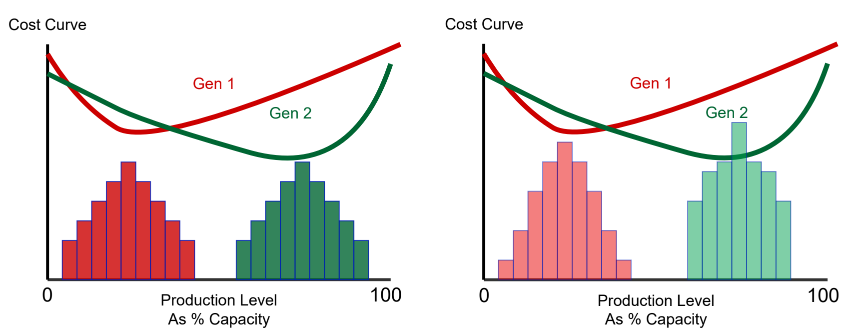

(a) **Pre-merger** (b) **Post-merger**

Note: Illustration of different efficiency gain mechanisms introduced in Section 7.

change in production. As discussed in Section 2.3, there is typically an efficient scale above and below which efficiency declines and firms incur ramp-up and ramp-down costs. Since the demand is stochastic and not perfectly predictable, balancing production by considering ramp-up and ramp-down costs and the shape of the cost curve is not straightforward. For example, it requires coordination between the trading desk personnel, who choose a bidding strategy, and plant personnel, who observes marginal cost and decide on production. Optimizing these margins allows a firm to produce more with less input by adjusting the production profile even if the cost curve remains identical pre- and post-acquisition. Figure 11 demonstrates this effect where production is more concentrated around the efficient scale post-acquisition, implying less ramp-up and ramp-down. A prediction of this mechanism is:

Prediction 2: If firms improve dynamic efficiency, the standard deviation of heat rate goes down.

Portfolio (Allocative) Efficiency. The third mechanism to improve efficiency is portfolio effects. Plant owners solve complex optimization problems with thousands of parameters as they face stochastic demand and time-varying transmission constraints. Having multiple power plants with different production costs in the same market can give firms more flexibility and improve efficiency by allocating production optimally across power plants. This effect is illustrated in Figure 12. Since this mechanism is present only if firms have other plants in the same market, a prediction for portfolio efficiency is:

Prediction 3: The efficiency of the existing plants of the acquirer firm in the same market will improve.

7.2 Quantifying Mechanisms of Efficiency Gains

We start by testing for the presence of productive efficiency using an empirical strategy guided by *Prediction 1*. In particular, we estimate the cost curves of generators non-parametrically by flexibly controlling for ramp-up and ramp-down costs and production level:

$$y_{it} = f_{i\tau}(Q_{it}, r_{it}),$$

where y_{it} is heat rate, Q_{it} is production level as percentage of total capacity of generator i at time t, and r_{it} is the ramp defined as the percentage change in production relative to $t-1$. $f_{i\tau}(Q_{it}, r_{it})$ is generator-specific cost function that depends on production level and ramp. We estimate this cost function for each generator separately using data one year

Figure 13: Estimated Cost Curves

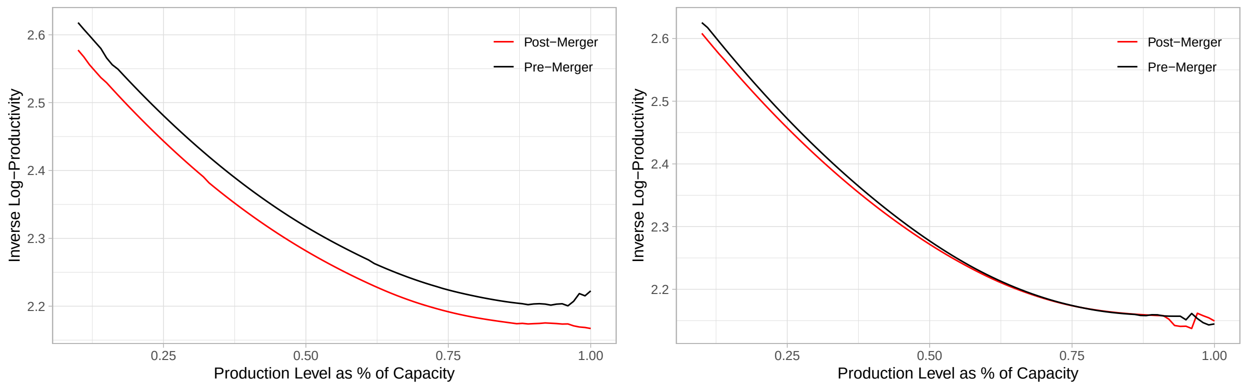

(a) Treated Group

(b) Control Group

Note: These figures shows estimates of average costs curves one year before acquisition and one year after acquisition. Panel (a) shows this for the treated group and panel (b) is for the control group.

prior and one year after acquisition. As a result, f_{i0} corresponds to the cost curve before acquisition, and f_{i1} corresponds cost curve after acquisition.[29]

It is worth highlighting two critical features of this exercise. First, the estimated cost function is generator-specific, as indicated by the index i. Estimating the cost curve at the generator level is important to capture heterogeneity in production technology across generators. Second, different from our main specification, we estimate this regression using hourly data to control for change in production level accurately. Moreover, hourly-level data allows us to estimate generator-specific cost functions since we have thousands of observations from each generator pre- and post-acquisition. This estimation highlights the advantage of the rich data environment, as traditional production function estimation typically requires aggregation at the industry level.

We estimate f_{i0} and f_{i1} for every generator acquired during our sample period and ask how the cost curve changes after acquisition controlling for ramp. In particular, we quantify the productive efficiency gain in the following way:

$$\Delta C(Q) = c_{post}(Q) - c_{pre}(Q) = \sum_i f_{i1}(Q, 0) - f_{i0}(Q, 0)$$

where $Q \in (0, 100)$, $c_{post}(Q)$ is the average cost at production level Q after acquisition, and

[29]We estimate these functions for generators that are operating more than 20 percent of the time in the window one year pre- and post-acquisition. This excludes some generators that produce only during peak demand and others that are not active before acquisition.

Figure 14: Change in Variation of Heat Rate

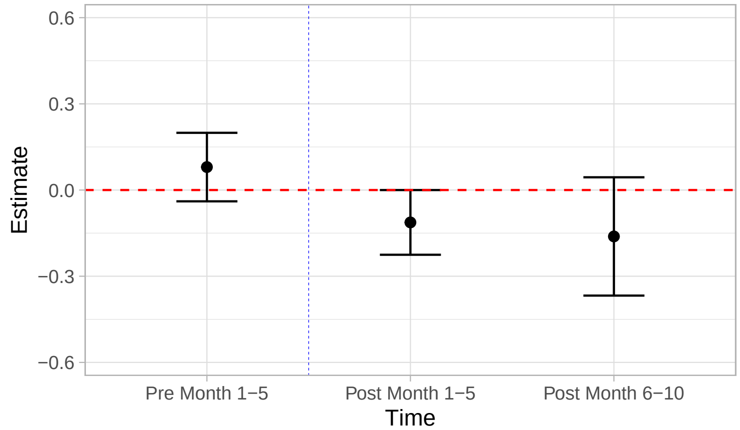

Note: Coefficient estimates from a regression of standard deviation of heat rate on treatment dummies. Error bars show 95 percent confidence intervals.

$c_{pre}(Q)$ is the average cost at production level Q before acquisition. The difference between these two functions gives us changes in productive efficiency at production level Q.

Figure 13a displays $c_{pre}(Q)$ and $c_{post}(Q)$ for the acquired generators, and Figure 13b displays the same curves for the control generators that have never acquired.[30] Comparing the pre- and post-acquisition cost curve demonstrates that the cost curve shifts downs at every production level for the treated plants, and the cost curves are almost identical for the control plants. The difference between the cost curves for the treated group is slightly larger at production levels close to the efficient scale, but this difference is not statistically significant. These results provide direct and strong evidence that the acquirers improve the productive efficiency of the acquired plants.

This analysis also allows us to quantify the role of productive efficiency. To see this, we integrate the difference between the cost curves to quantify the total productivity gain from productivity efficiency:

$$\Delta = \frac{1}{N_{acq}} \sum_{i}^{N_{acq}} \int \big(f_{i1}(Q,0) - f_{i0}(Q,0)\big) dF_i(Q),$$

where N_{acq} is the number of acquired generators and $dF_i(Q_{it})$ is the distribution of production level of generator i before acquisition. This calculation suggests the overall efficiency

[30] Details of how the control units are constructed are given in the Appendix. We also provide the bootstrapped standard errors for the difference between the two cost curves in the Appendix.

Figure 15: Impact of Merger on Other Plants

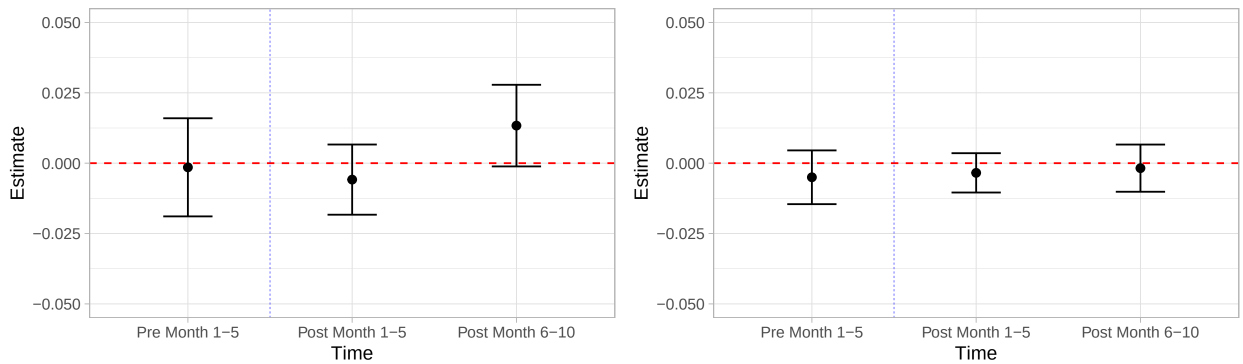

(a) Same Market

(b) Different Market

Note: Panel (a) shows coefficient estimates from a regression of log efficiency on treatment dummies where existing units of the acquirer in the acquisition market are treated. Panel (b) shows the results from the same regression except that existing units of the acquirer in the different markets are treated. Error bars show 95 percent confidence intervals. Standard errors are clustered at the plant level.

gain explained by productive efficiency is 3 percent, corresponding to 75 percent of the total efficiency gain identified in the previous section. Therefore, we conclude that most efficiency gain comes from productive efficiency.

Next, we move to dynamic efficiency. According to Prediction 2, an increase in dynamic efficiency should reduce the standard deviation of the heat rate after acquisition. To test this hypothesis, we estimate the same specification as in Equation (5.1) but use the weekly standard deviation of heat rate as the outcome variable. Figure 14 shows the results from this specification. We find that the average standard deviation of the heat rate goes down after acquisition. Unlike the efficiency results, the decline in the volatility of heat rate is realized more rapidly. From this regression, we conclude that the firms not only improve production efficiency but also dynamically allocate production more efficiently over time to reduce ramp-up and ramp-down costs.

Finally, we test the portfolio efficiency effects of acquisitions. According to Prediction 3, portfolio efficiency occurs only if the acquirer owns other plants in the same market, and plants in other markets are not affected by portfolio efficiency. To test this prediction, we estimate our main specification in Equation (5.1) where we treat the generators owned by the acquirer. We do this estimation separately for the acquirer's generators in the same market where acquisition happens and for those in different markets. We assume power plants are in the same market if they are located in the same balancing authority.

Figure 15 presents results for these regressions. Figure 15a shows the change in the heat rate of the generators owned by the acquirer and located in the same market as the acquired plant. Figure 15b shows the change in heat rate of the plants owned by the acquirer but located in different markets. We find that acquirers' power plants in the same market exhibit efficiency improvements by 1.3 percent, whereas acquirers' power plants in different markets show no change in the average heat rate. On balance, these results suggest an efficiency increase for acquirers' existing power plants, but only if they are in the same market, consistent with portfolio efficiency. Moreover, we observe that the efficiency improvement of the acquirer's plants is much lower than the efficiency improvements of the acquired plants, 1.3 vs. 4 percent, suggesting that the scope for improvement through productive efficiency is larger than the scope of improvements through allocative efficiency.

To summarize this section, we find evidence for all three mechanisms we identified: (i) productive efficiency, (ii) dynamic efficiency, and (iii) allocative efficiency. The mechanism with the largest effect is productive efficiency, explaining 75 percent of the total efficiency gain. The rest is explained by dynamic and allocative efficiency.

7.3 How Do Firms Improve Productive Efficiency?

So far, our results have provided clear evidence of an efficiency increase after acquisition, and this efficiency increase primarily comes from productive efficiency. The next natural question is what firms do to achieve this efficiency gain. In this section, we investigate this question.

In Section 2.3, we posited two potential ways to improve power plants' productive efficiency. The first is that acquirers make operational improvements or hire more skilled personnel after acquisition. This would point transfer of knowledge from the acquirer to acquired generator. The second is large-cost capital investments, where acquirers upgrade the capital. If efficiency improvements occur this way, it would suggest that the previous owner has liquidity constraints and cannot make efficiency-improving capital investments. Disentangling these two sources of efficiency gains is important for merger policy because efficiency gains have to be merger-specific for them to be viewed as cognizable. Efficiency increases due to relaxing capital constraints are not merger-specific, as they can be accomplished by lending or investment subsidy. However, knowledge transfer is merger-specific since it happens in an organization and is unlikely to be accomplished without a merger.

We disentangle the sources of productive efficiency improvements with additional datasets on manager changes, capital expenditures and non-fuel costs. In particular, we ask whether power plants experience personnel changes after the acquisition and whether

Figure 16: Probability of Management Change

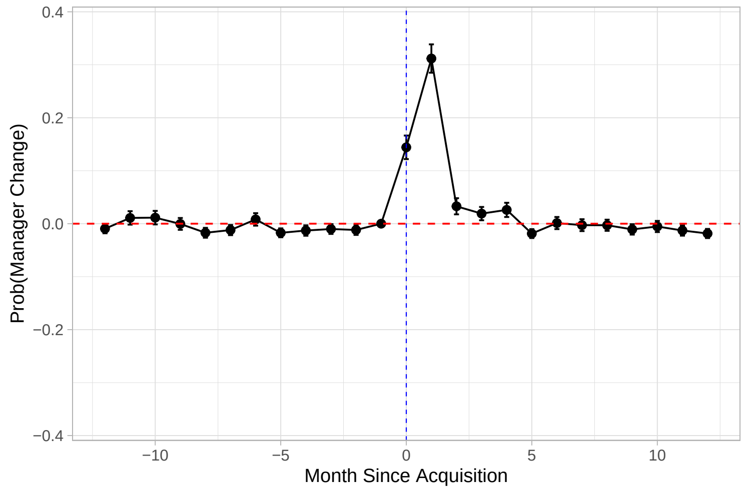

Note: Coefficient estimates from a regression of management change dummy on pre- and post-treatment variables. Error bars show 95 percent confidence intervals.

there is any significant change in capital expenditures and costs. Personnel changes would provide suggestive evidence for significant operational changes after acquisition, and changes in capital expenditures would provide direct evidence for the role of cost investment.

To study whether power plant managers change after an acquisition, we estimate the dynamic difference-in-differences specification given in Equation (5.2). The outcome is an indicator variable that equals one if the power plant manager is replaced after acquisition and zero otherwise. We include the same control variables but estimate the regression at the monthly level. Figure 16 reports coefficient estimates. We see that the probability of management change jumps with acquisition, with 15 percent of acquired power plants changing their managers within one month and 30 percent changing their managers within two months. The cumulative change is 55 percent within 12 months after acquisition.[31] These results suggest that acquired firms make operational changes through new management, potentially affecting efficiency. The potential role of management changes in explaining productivity differences is plausible, given the recent findings of Bloom and Van Reenen (2010), who show that productivity measures correlate with various management practices.

Next, we turn to results from capital expenditures, which are reported in Table 4. Since this variable is at the annual level, we estimate the difference-in-differences specification with yearly data. The results suggest that acquirers do not increase their capital expen-

[31] Note that the unconditional probability of management change in a given year is only 10 percent.

Table 4: Effects of Mergers on Non-fuel Costs

	Non-fuel Cost	Number of Employees	Capital Expenditures
	(i)	(ii)	(iii)
Post-Merger × Treat	-0.068	-0.054	-0.020
	(0.053)	(0.031)	(0.032)
# of Acq	655	584	678
# of Obs	29325	26866	29418
R^2	0.62	0.92	0.86
Unit FE	X	X	X
Year FE	X	X	X

Note: This table presents the coefficient estimates from estimating the effects of mergers on non-fuel cost, number of employees and capital expenditures with annual data. Standard errors are clustered at the plant level.

ditures after acquisition. As additional analyses, we study the number of employees and non-fuel costs to understand whether there are substitutions to other inputs. The results, reported in Columns 2 and 3, do not suggest any evidence for substitution effects.

Finally, the timing of the efficiency effects also provides indirect evidence against the capital expenditure hypothesis. Large capital expenditures often require significant downtime in the power plant; therefore, efficiency gains through investment take a long time to be realized. In contrast, we see efficiency begin to rise right after acquisition and reach a new level in six months. Overall, the evidence in this section suggests that firms achieve productive efficiency through operational improvements rather than costly capital expenditures.

8 Robustness Checks

In this section, we investigate the robustness of our results to alternative specifications.

8.1 All Acquisitions

In our estimations, we only use the first acquisitions of generators if they are acquired multiple times (about 25 percent of all plants). We did this because it was unclear how to properly estimate the event study with generators acquired more than once. In this robustness check, we repeat our estimation procedure by including all acquisitions of generators if they are acquired more than once in our sample period. The results, reported in Appendix E, are broadly similar to our main results.

8.2 Estimation with Daily and Hourly Data

Our main specification estimates the effects of acquisitions with data at the weekly level. Aggregation at the weekly level decreases computation time and reduces noise in the hourly data. To understand how robust our results are to this choice, Appendix E shows the estimation results from daily and hourly data. We see that the results are robust to estimation frequency, with some increase in standard errors.

8.3 Staggered Difference-in-Differences

Recent literature in econometrics has shown that the difference-in-differences method could yield a weighted average of all possible permutations of pairwise difference-in-differences estimators, where a pair is either the never-treated control group compared with the cohort treated at time t, or a cohort treated at time t compared with a cohort treated later (Goodman-Bacon (2021), Callaway and SantAnna (2021), and De Chaisemartin and dHaultfoeuille (2020)). To address this point, we estimate cohort-specific treatment effects using the Callaway and SantAnna (2021) method and report dynamic treatment effects.

8.4 Matching Difference-in-Differences

Our main specification uses standard difference-in-differences estimation estimated with two-way fixed effects. In Appendix E, we also consider a matching estimator. The matching estimator matches acquired generators to similar generator and calculate generator-specific treatment effects by comparing them to match plants. To implement this estimation, we identify the three nearest neighbors from our sample pool of 2,500 control units. We match on capacity, age, and fuel type using a least-squares metric to calculate the distances between generation units, with weights inversely proportional to standard deviation for age and capacity. We use our distance measure to select the three nearest neighbors for each acquired unit, allowing control units to be matched to multiple acquired plants. We exclude the generators in the same market from the pool of potential control plants due to potential spillover effects.

8.5 Weighted Difference-in-Differences

In our main specification, we estimate average treatment effects without considering the different capacity sizes of acquired plants. An alternative estimation would be weighting observations by their production level or capacity, which would be a more accurate measure of total cost savings. We report estimates from this specification in Appendix E, and the results are broadly similar.

9 Concluding Remarks

By allocating resources between firms, mergers and acquisitions affect a significant portion of the economy. Despite this importance, there is limited systematic evidence of their effects on productivity and market power. This study provides detailed empirical analyses of the efficiency effects of mergers by examining a large sample of power plant mergers and acquisitions between 2000 and 2020. Our empirical results can be summarized into three principal findings. First, we find that acquired plants experience an average of four percent efficiency increase five to eight months after acquisition, and most of this productivity increase is explained by improvements in productive efficiency. Second, our findings suggest that acquisitions reallocate assets to more productive uses: we find that high-productivity firms buy underperforming assets from low-productivity firms and make the acquired asset almost as productive as their existing assets after acquisition. Finally, we find that the new owners improve productivity by changing operational processes rather than making costly capital investments.

The underlying source of our findings is using a large number of acquisitions in the power generation industry and taking advantage of high-frequency physical productivity measures obtained from physical input and output quantities. With physical measure and studying a homogeneous product, we can disentangle the productivity effects from other potential merger effects, such as market power, buyer power, and changes in quality. With high-frequency data, we can treat mergers as discrete events and compare firm productivity immediately before and after the acquisition. Finally, by aggregating evidence from a large number of mergers and acquisitions, we have statistical power to uncover many interesting mechanisms that could generate efficiency gains.

Our findings have important policy implications, as they can be a direct input to evaluating the trade-off between market power and efficiency due to mergers. Beyond antitrust, our results have important implications for the role of mergers on aggregate productivity growth. Our finding that mergers reallocate assets to more productive firms suggests that mergers contribute to aggregate productivity growth.

We believe that the availability of this high-frequency and high-quality data on power plant production can be useful for understanding other important issues for antitrust and productivity, such as quantifying total welfare effects by estimating price effects of mergers together with efficiency and studying how strategic firm behavior changes after a merger. These various issues various could be fruitful areas for future research.

References

Atalay, Enghin. 2014. "Materials prices and productivity". *Journal of the European Economic Association* 12 (3): 575–611.

Atalay, Enghin, Alan Sorensen, Wanjia Zhu, and Christopher Sullivan. 2020. "Post-Merger Product Repositioning: An Empirical Analysis".

Bhattacharya, Vivek, Gaston Illanes, and Stillerman. 2022. "Have Mergers Raised Prices? Evidence from US Retail".

Bitzan, John D, and Wesley W Wilson. 2007. "Industry costs and consolidation: efficiency gains and mergers in the US railroad industry". *Review of Industrial Organization* 30 (2): 81–105.

Blonigen, Bruce A., and Justin R. Pierce. 2016. *Evidence for the effects of mergers on market power and efficiency*. National Bureau of Economic Research.

Bloom, Nicholas, and John Van Reenen. 2010. "Why do management practices differ across firms and countries?" *Journal of economic perspectives* 24 (1): 203–24.

Bradley Associates, LLC. 2017. "Powering America: The Economic and Workforce Contributions of the U.S. Electric Power Industry". *Available at: https://www. electric.coop/wp-content/uploads/2017/08/PoweringAmerica.pdf*.

Braguinsky, Serguey, Atsushi Ohyama, Tetsuji Okazaki, and Chad Syverson. 2015. "Acquisitions, productivity, and profitability: evidence from the Japanese cotton spinning industry". *American Economic Review* 105 (7): 2086–2119.

Bushnell, James B., and Catherine Wolfram. 2005. "Ownership change, incentives and plant efficiency: The divestiture of US electric generation plants".

— . 2009. "The guy at the controls: Labor quality and power plant efficiency". In *International differences in the business practices and productivity of firms*, 79–102. University of Chicago Press.

Callaway, Brantly, and Pedro HC SantAnna. 2021. "Difference in differences with multiple time periods". *Journal of Econometrics* 225 (2): 200–230.

Campbell, Richard J. 2013. *Increasing the efficiency of existing coal-fired power plants*.

Cicala, Steve. 2022. "Imperfect markets versus imperfect regulation in US electricity generation". *American Economic Review* 112 (2): 409–41.

— . 2015. “When does regulation distort costs? lessons from fuel procurement in us electricity generation”. *American Economic Review* 105 (1): 411–44.

Craig, Stuart V, Matthew Grennan, and Ashley Swanson. 2021. “Mergers and marginal costs: New evidence on hospital buyer power”. *The RAND Journal of Economics* 52 (1): 151–178.

Dafny, Leemore, Kate Ho, and Robin S Lee. 2019. “The price effects of cross-market mergers: theory and evidence from the hospital industry”. *The RAND Journal of Economics* 50 (2): 286–325.

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.

De Chaisemartin, Clément, and Xavier dHaultfoeuille. 2020. “Two-way fixed effects estimators with heterogeneous treatment effects”. *American Economic Review* 110 (9): 2964–96.

De Loecker, Jan, and Frederic Warzynski. 2012. “Markups and firm-level export status”. *American economic review* 102 (6): 2437–71.

Dranove, David, and Richard Lindrooth. 2003. “Hospital consolidation and costs: another look at the evidence”. *Journal of health economics* 22 (6): 983–997.

Eliason, Paul J., Benjamin Heebsh, Ryan C. McDevitt, and James W. Roberts. 2020. “How acquisitions affect firm behavior and performance: Evidence from the dialysis industry”. Publisher: Oxford University Press, *The Quarterly Journal of Economics* 135 (1): 221–267.

EPA. 2018. “FACT SHEET: Overview of the Clean Power Plan”.

EPRI. 2022. *Proceedings: 2022 EPRI heat-rate improvement conference*. Tech. rep. Electric Power Research Institute.

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.

Farrell, Joseph, and Carl Shapiro. 2010. “Antitrust evaluation of horizontal mergers: An economic alternative to market definition”. *Available at SSRN 1313782*.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?" *American Economic Review* 98 (1): 394–425.

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

Harrison, Teresa D. 2011. "Do mergers really reduce costs? Evidence from hospitals". *Economic Inquiry* 49 (4): 1054–1069.

Jovanovic, Boyan, and Peter L. Rousseau. 2008. "Mergers as reallocation". Publisher: The MIT Press, *The Review of Economics and Statistics* 90 (4): 765–776.

— . 2002. "The Q-theory of mergers". *American Economic Review* 92 (2): 198–204.

Knittel, Christopher R. 2002. "Alternative regulatory methods and firm efficiency: Stochastic frontier evidence from the US electricity industry". *Review of Economics and Statistics* 84 (3): 530–540.

Kulick, Robert B. 2017. "Ready-to-Mix: Horizontal Mergers, Prices, and Productivity". *US Census Bureau Center for Economic Studies Paper No. CES-WP-17-38*.

Lewis, Matthew S, and Kevin E Pflum. 2017. "Hospital systems and bargaining power: Evidence from out-of-market acquisitions". *The RAND Journal of Economics* 48 (3): 579–610.

Luco, Fernando, and Guillermo Marshall. 2020. "The competitive impact of vertical integration by multiproduct firms". *American Economic Review* 110 (7): 2041–64.

Maksimovic, Vojislav, and Gordon Phillips. 2001. "The market for corporate assets: Who engages in mergers and asset sales and are there efficiency gains?" Publisher: Wiley Online Library, *The journal of finance* 56 (6): 2019–2065.

McGuckin, Robert H., and Sang V. Nguyen. 1995. "On productivity and plant ownership change: New evidence from the longitudinal research database". Publisher: JSTOR, *The RAND Journal of Economics*: 257–276.

Mercadal, Ignacia. 2022. "Dynamic Competition and Arbitrage in Electricity Markets: The Role of Financial Players". *American Economic Journal: Microeconomics* 14 (3): 665–699.

Nguyen, Sang V, and Michael Ollinger. 2006. "Mergers and acquisitions and productivity in the US meat products industries: Evidence from the micro data". *American Journal of Agricultural Economics* 88 (3): 606–616.

Nocke, Volker, and Michael D Whinston. 2022. "Concentration Thresholds for Horizontal Mergers". *American Economic Review* 112 (6): 1915–48.

Prager, Elena, and Matt Schmitt. 2021. "Employer consolidation and wages: Evidence from hospitals". *American Economic Review* 111 (2): 397–427.

Rhodes-Kropf, Matthew, and David T. Robinson. 2008. "The market for mergers and the boundaries of the firm". Publisher: Wiley Online Library, *The Journal of Finance* 63 (3): 1169–1211.

Sargent & Lundy, LLC. 2009. "Coal-fired power plant heat rate reductions". *LLC Sargent & Lundy*.

Schmitt, Matt. 2017. "Do hospital mergers reduce costs?" Publisher: Elsevier, *Journal of health economics* 52:74–94.

Schoar, Antoinette. 2002. "Effects of corporate diversification on productivity". Publisher: Wiley Online Library, *The Journal of Finance* 57 (6): 2379–2403.

Staudt, James E., and Jennifer Macedonia. 2014. "Evaluation of heat rates of coal fired electric power boilers". In *Paper was Presented at the Power Plant Pollutant Control "MEGA" Symposium August*, 19–21.

Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49, no. 2 (): 326–365.

A Data Appendix

A.1 Unit Level Data

We use EIA Forms 860 and 923, EPA's Continuous Emissions Monitoring Systems (CEMS), S&P Global, and Velocity Suite to construct a dataset for unit characteristics. The EIA forms and CEMS are public data sources, whereas S&P Global and Velocity Suite are data providers for energy markets that require a subscription. The EIA Forms cover the universe of power plants in the US, whereas CEMS data include power plants with a capacity above 25 MWh that are subject to environmental regulations. Private data providers S&P Global and Velocity Suite compile unit- and plant-level data from various resources, including EIA, EPA, and FERC. We merged these datasets based on unit and plant names available in all datasets. The final dataset includes information on generation, regulation status, technology type, install year, fuel type, coal type, boiler type, boiler model, boiler manufacturer, capacity, fuel cost, prime mover category, and dispatch type. We also know whether a unit is connected to the grid, is an internal generator, is marginal or infra-marginal, and can switch fuel. We identify deregulation cases as changes in the regulation status at the plant level in Form 860. This procedure results in a monthly panel data (except for generation) on unit characteristics. We provide more details about some of the variables below.

Generation Continuous Emission Monitoring Systems (CEMS) are required under some of the EPA regulations for continual compliance determinations of environmental regulation. For this purpose, EPA collects boiler-level hourly data on heat rates from fossil-powered power plants and makes this data publicly available. The coverage corresponds to roughly 95 percent of fossil-powered generation in the US. Even though the data is available starting in 1995, the data quality is poor before 2000. For this reason, we restrict the study period from 2000 to 2020. We remove power plants in Alaska and Hawaii and winsorize the heat rate if they are larger than 16 and lower than 5, which affects less than 1 percent of observations. We also eliminate generators that are not active for more than 10 percent of their lifetime. With these restrictions, the final sample includes all the US fossil-fuel generators under the CEMS program, except those in Alaska and Hawaii and those that are active less than 10 percent of the time during their lifetime. This procedure results in an hourly unit-level dataset on generation fuel input, and heat rate between 2000 and 2020. We aggregate this data to daily to weekly levels in some of the analyses employed in the paper.

We match unit-generation data from CEMS to unit-level data from other data sources mentioned above. While most units are easily matched using the unit name, some do not match as EPA uses boiler names as a unit, whereas EIA uses generator names. For those cases, we rely on the EPA's Power Sector Data Crosswalk available on EPA's website.[32]. This crosswalk does not include units that retire before 2020. For those retired and other unmatched units, we manually match based on capacity, installment year, and retirement year information.

Gross vs. Net Generation The generation and heat rate provided by EPA are based on gross generation, including the ancillary services and other non-market products and consumption at the unit, such as scrubbers. For our study, gross generation is the relevant variable for understanding the overall efficiency of power plants since we study how fuel is transformed into electricity, not revenue obtained from the generation.

Capacity EPA data does not provide capacity data. Therefore, we need to assign a capacity for those units that do not exactly match the EIA and CEMS data. We infer capacity from their generation using the following algorithm for these units. We first keep units that work more than two weeks' worth of hours each year. Then, we obtain the hourly generation distribution each year and use the 99th percentile of the generation as the capacity for the unit every year. This algorithm generates unit capacity that is stable over time for most units. However, there is a significant yearly variation for some units with no apparent capacity change at EIA plant-level data. In that case, we take the unit's maximum capacity for consecutive years. To check the accuracy of this algorithm, we look at how close the capacity provided in EIA is for the units that have a perfect match in the EPA and the EIA. We find that capacity generated from the EPA data align with those provided by EIA.

A.2 Plant Level Data

We use EIA Forms 860 and 923 and Velocity Suite to construct data for characteristics. From these data sources, we obtain information on location, ISO, NERC region, regulation status, capital expenditures, number of personnel, and non-fuel costs. Data on capital expenditures, number of personnel, and non-fuel costs are obtained from Velocity Suite, which is compiled from annual FERC Form 1. FERC Form No. 1 is a comprehensive financial and operating report submitted for Electric Rate regulation and financial audits.

[32] https://www.epa.gov/airmarkets/power-sector-data-crosswalk.

It is mandatory for investor-owned utilities; therefore, the coverage for these variables is lower than the coverage of other variables.

Regulation To identify ownership changes due to deregulation, we use Cicala (2022)'s deregulation list from 2000-2012. After 2012, we rely on EIA 860's regulation status for plants. Using this dataset, we flag ownership changes that coincide with a change in the regulation status, and we exclude those from the merger sample. This results in a total of 181 plants between 2000 and 2020 that we remove from the merger sample.

A.3 Load Data and Market

We collect data and hourly demand and market definition. Market definitions in electricity markets are time-varying due to nodal pricing and congestion. These definitions are especially relevant for market power considerations (Mercadal (2022)). However, this paper only uses ISO market definitions and relies on Velocity data for plants' corresponding ISOs.

Hourly data on electricity usage (load) is obtained from S&P Global either at the Balancing Authority Areas (BAA) or ISO Zone level, depending on data availability. We also use FERC Form 714 to obtain data on the fuel composition of total generation. FERC Form 714 treats Power Control Areas (PCA)s as markets, and the PCA market definition overlaps with our ISO market definitions for the deregulated plants and gives further granularity to regulated plants. Before 2006, data from FERC 714 was not complete, so we assemble data on load after 2006. We rely on EIA's market definitions to match load data with plants, resulting in data on load for roughly 70 percent of the plants.

A.4 Personnel Data

Each power plant subject to at least one of the EPA programs must submit a representative to EPA. This representative information is essential for EPA, as a potential problem like leakage needs to be addressed quickly, and responsible parties should be accountable. This data is self-reported and includes the representative's name, start and end date, and contact information. We use data on plant representatives from EPA between 2000 and 2020 to construct personnel data. Even though this data is complete between these years, it does not include some key information, such as job titles. To obtain this information, we matched representative names to their Linkedin profiles and found about 70 percent of representatives on Linkedin. The match rate improves over time, reaching 80-90 percent in later years. We obtain a history of job titles, employment, and education from

Linkedin profiles. The job title suggests that about 70-80 percent of representatives are plant managers and the rest are engineers or regulatory compliance managers. Given that most of these representatives are plant managers, we treat this personnel information as plant managers in this study.

This procedure results in a monthly plant-level panel data on plant managers. In this data, we know the manager's exact start and end date of tenure, and we have information on the manager's employment and education history.

A.5 Ownership Data

Every acquisition that involves a power plant, however small it is, should be notified to the corresponding state or federal agency for approval. For this reason, very rich data is available on the universe of power plant mergers and acquisitions. To construct this dataset, we use two separate data from S&P Global: ownership and transaction datasets.

The first dataset includes all shareholders (name and company ID) of a generator, how much share is owned by each shareholder, and the exact date of the change in the power plant ownership. For a given generator, this dataset is updated when there is a change in the owners of the generator. In particular, for any given ownership change, we observe the names of the owner before the change and after the change along with the ownership changes. We turned this data into a month-generator panel dataset and included the largest three shareholders into this panel. For each generator, we started this panel from the installment date of the generator and remove the ownership information before the installment data and after the retirement date. Another advantage of this dataset is that it provides the subsidiary that owns a power plant and the parent company that owns the subsidiary. Therefore, we can see the corporate structure of the owner of this power plant. S&P Global backfills any company name change, so firm name changes do not affect the ownership structure significantly over time. Our final ownership results in a month-generator panel with the following information: the largest three shareholders of the generator, the parent company of each shareholder, and the percentage of the power plant owned by each shareholder.

The second dataset is mergers and acquisition data. This dataset provides detailed information for every transaction, such as buyer, sellers, transaction type (divestitures, cash deal, LBO), and deal value. This dataset includes a transaction ID and transaction description. Around 80-85 percent of transactions include transaction descriptions where one can see acquired assets, acquisition motives, and other important information. The rest of the transactions do not have a description. For these transactions, we manually search for companies involved in the transaction and manually classify whether these are true ownership changes or corporate restructuring. Most of these transactions without a description turned out to be false acquisitions.

Next, we merged the two datasets using transaction ID and company names. This gives us a complete picture of ownership changes, including new owners and previous owners, and very rich data on merger characteristics. After merging these datasets, we removed the ownership changes we identified as false ownership changes from the transaction data. This gave us the final sample of data on ownership changes.

A.6 Firm Data

Even though the transaction data provides useful information about buyers and sellers, we used CapitalIQ to obtain more information about companies involved in transactions. For this purpose, we merged company balance sheet data from CapitalIQ with the ownership data from S&P Global. This data merge is straightforward for about 80 percent of the companies because, for those, S&P Global and CapitalIQ use the same company ID. For the rest, we manually searched for company names in CapitalIQ and matched those to the ownership data. Except for a few companies that went bankrupt or were company funds, we could match all company names. This firm data provides information such as industry, year founded, asset size, and various balance sheet information.

A.7 Sample Restrictions

Before employing the empirical analysis, we made some other sample restrictions. This section describes those restrictions. We first take month-unit level ownership data for fossil fuel generators from S&P Global between 2000 and 2020, which rely on EIA data for unit definitions. We drop a few power plants with missing ownership data for the entire study period.

Then we match S&P Global ownership data with EPA CEMS data. We use our crosswalk on top of EPA's crosswalk between CEMS and EIA for the mismatch cases. For plants whose unit owners are the same, mismatch cases do not present any problem, as we can directly impose ownership on all the units on CEMS data. Sometimes, our crosswalk leads to possible one-to-one ownership matches for the plants with different unit owners, even when there are multiple owners. For other cases, we merge ownerships at the CEMS unit level by using capacity-weighted average ownership shares.

After the match, we remove units in S&P Global that are missing in CEMS data. Some units in CEMS data produce only steam. Moreover, some units do not have any production data in our observation period. We remove these units from our sample in both cases.

Heat rate observation is calculated by dividing the total heat input by the total electricity output for an hour. Therefore if there are significant changes in the production within the hour, the heat rate could be very high or very low. This sometimes generates noise in hourly heat rates, especially at small production levels. Therefore, we winsorize heat rates above 16 or below 6 MMBtu per MWh. This winsorization affects less than 1 percent of all observations.

B Estimation Details

B.1 Estimation of Residual Productivity

B.2 Heterogeneity Analysis

This section provides the estimation details for the heterogeneity analysis that are presented in Section 5.3.

To increase the power in detecting heterogeneity, we consider a standard event study setting where we include a post-treatment dummy variable and interact with the variable for which we want to understand heterogeneity. In particular, we estimate the following specification:

$$\log(y_{it}) = \theta_1 D_{it} + \theta_2 \text{Treat}_{it} D_{it} \times Z_{it} + X_{it} + \mu_t + \alpha_i + \eta_{it},$$

where y_{it} is the efficiency of generator i at week t (measured as inverse heat rate given in Equation (2.1), the controls, X_{it}, include state-month fixed effects, time-varying generator characteristics such as age and fuel type (for coal), capacity, indicators for whether the unit is connected to the grid and whether it is an internal generator. α_i is generator fixed effect and μ_t is week fixed effect. Z_{it} is a plant, firm, or transaction characteristic we would like to test heterogeneity. The results report the estimates of θ_3.

B.2.1 Plant Characteristics

- **Gas Plant:** An indicator variable that equals one if the acquired unit is powered by natural gas and zero otherwise. Since most of the acquired natural gas power plants, this variable equals 1 for 90 percent of transactions.

- **Plant Age $>$ Median** An indicator variable that equals one if the age of the acquired unit is larger than the median. We consider all the units in our main specification to calculate the median age and find the median value.

- **Unregulated Plant:** An indicator variable that is one if the plant is unregulated and zero otherwise.

- **Unit Capacity $>$ Median:** An indicator variable that equals one if the age of the acquired unit is larger than the median. To calculate the median capacity, we consider all the units in our main specification and find the median capacity.

- **infra-marginal Plant:** An indicator variable that equals one if the acquired plant is infra-marginal and zero otherwise. This categorization is provided by Velocity Suite, which classifies plants as base load or peaker units.

B.2.2 Firm Characteristics

- **Target Exits:** An indicator variable that equals one if the target firm owns no fossil fuel power plants post-transaction and zero otherwise.

- **Acquirer Enters** An indicator variable that equals one if the acquirer firm owns no fossil fuel power plants pre-transaction and zero otherwise.

- **Financial Aquirer:** An indicator variable that equals one if the acquirer is a financial firm and zero otherwise. The firm classification is obtained from CapitalIQ Pro.

- **Acquirer Size $>$ Median:** An indicator variable that equals one if the total capacity of the acquirer pre-transaction is larger than the median capacity of firms that have been involved in a transaction between 2000-2020.

- **Serial Acquirer:** An indicator variable that equals one if the total capacity acquired by the acquirer between 2000-2020 is larger than the median of the total capacity acquired by firms between 2000-2020.

B.2.3 Deal Characteristics

- **Deal Value $>$ Median:** An indicator variable that equals one if the total transaction value (including assets other than power plants) and zero otherwise.

- **Year$>$2010** An indicator variable that equals one if the acquisition occurs later after 2010 and zero otherwise.

- **Acquirer's Capacity in the Market$>$Median:** An indicator variable that equals one if the acquirer's capacity in the ISO where the plant is located is greater than the median capacity.

- **Acquired Capacity $>$ Median:** An indicator variable that equals one if the total fossil fuel plant acquired capacity in the transaction is larger than the median.

- **Bankrupt Sale:** An indicator variable that equals one if the transaction is a bankrupt sale and zero otherwise.

C Robustness Checks

In this section, we provide the details of the robustness checks we employ in this paper.

C.1 Acquisition Sample

Since our sample covers 20 years, many units are acquired more than once. Out of the total 2500 units that have been acquired, 1200 of them were acquired more than once. In our main specification, we considered only the first acquisitions because, with multiple acquisitions, the post period of the first acquisition overlaps with the other acquisitions. So to simplify the differences and differences analysis, we consider the first acquisitions and drop the observations of the firm's multiple acquisitions 16 months after the acquisition. This section analyzes how robust our results are in considering different acquisition samples. In particular, we will consider all acquisitions, units acquired only once, and acquisitions after the first.

The specification with all acquisitions includes all acquisitions except for those that are within 32 months of each other. We drop these because the post and pre-acquisitions periods overlap. Then we will estimate the same equation but include all acquisitions. In this specification, 16 months following the acquisition, we will have the post-treatment indicators for each month; 16 months before the acquisition, we include pre-treatment indicator variables for each month. We will include a single dummy variable for the periods after 16 months after the last acquisition and single dummy variable 16 months before the last acquisition. We assume control and treated groups follow the same trend for other periods. The results from this estimation are reported in Appendix Figure 30 and Table 14.

In another specification, we include units that are acquired only once. This specification reduces the number of acquired units significantly because 75 percent of acquisition events come from units that are acquired multiple times. Nevertheless, we want to see whether the results are robust. The results from this estimation are reported in Table 12.

C.2 Data Frequency

In our main specification, we considered estimation at the weekly frequency where the efficiency is defined as total electricity output divided by total heat input. The weekly frequency reduces the computational burden and reduces noise in the hourly data. We repeat our estimation with daily and hourly data to understand whether our results are robust to estimation frequency.

Estimation with daily data follows the same steps as the estimation in weekly data. We aggregate input and output to a daily level, define efficiency at the daily level, and include monthly indicator variables for treated units before and after acquisition. We estimate the same specification as in Equation, with an additional control variable: the day of the week. Since the day of the week is a strong driver of demand, estimation with daily data better controls for demand fluctuations. The results from this estimation are reported in Appendix Table 10 and Figure 27. The effect of mergers on efficiency is similar to what we found with the weekly data.

In estimation with hourly data, we use the raw data obtained from CEMS. We have electricity input and fuel input every hour. We consider the same specification as weekly and daily data, but we also add fixed effects hour of the day. Since the hour of the week is a strong predictor of demand, the hourly specification controls for demand much more precisely than daily and hourly data. The results from this estimation are reported in Appendix Table 11 and Figure 28. The effect of mergers on efficiency is similar to the results with weekly data, but estimates are less precise since hourly data is noisier.

Overall, these robustness checks suggest that our results are robust to aggregation of input and output at the daily and weekly levels.

C.3 Observation Weights

In our regressions, we weighted units equally. A natural alternative to this is to weigh them by size, which would be robust to the case where all efficiency gains come from small units. Moreover, it would be more informative about how much total production is affected by efficiency gains. To investigate this, we estimate Equations (5.1) and (5.2) by weighting units by their capacity. The results from this estimation are reported in Appendix Table 13 and Figure 31. We find that the efficiency effect is slightly larger when we weigh units by capacity, which is consistent with the finding reported in Figure 8 that the efficiency effect is larger for larger units. This finding also suggests that acquisitions of small units do not drive our main results.

C.4 Matching Estimators

Our main specification uses standard difference-in-differences estimation estimated with two-way fixed effects. In this section, we also consider a difference-in-differences matching estimator as a robustness check.

We match each of our 2365 acquired units to the three nearest neighbors from the pool of 2500 units that have never been acquired during our sample period. For each treated unit, we first find the unit that is active during the time of the acquisitions with the same fuel type and in a different ISO (to prevent spillovers). This constitutes the pool of candidate control units for that unit. Then, we find the nearest neighbor units on capacity and age using a least-squares metric to calculate the distances between generation units. The weights in the metric are inversely proportional to the standard deviation of the corresponding variable. We allow for control units to be matched to multiple acquired plants. Using these nearest neighbors, we calculate the unit-specific treatment effect as follows:

$$\widehat{\Delta} Y_{it} = Y_{it}(1) - \widehat{Y}_{it}(1) \tag{C.1}$$

where $\widehat{Y}_{it}(1)$ is the average heat rate of the control units that are matched to i and scaled such that the average outcome of the control at the time of acquisition is the same as the outcome of the treated unit. By indexing the levels to a baseline period, we obtain a unit-specific âĂIJdifference-in-differencesâĂİ estimate for any outcome of interest. We take the average of the unit-specific treatment effects to obtain the final estimates.

To construct the confidence intervals, we employ a bootstrap procedure, where we resample without replacement both the never treated plants and treated plants and follow the same matching procedure. We repeat this procedure 100 times and take the 2.5 and 97.5 percentile of the bootstrap distribution to construct the confidence intervals.

The results from this estimation are reported in Appendix Figure 29.

D Additional Tables and Figures

Table 5: Largest 25 Acqusitions

Acquirer	Target	Year	Capacity	# of units
NRG Energy, Inc.	GenOn Energy, Inc.	2012	26174	139
Investor group	Calpine Corporation	2018	22991	127
RRI Energy, Inc.	Mirant Corporation	2000	22748	140
Duke Energy Corporation	Progress Energy, Inc.	2012	19048	134
GC Power Acquisition LLC	CenterPoint Energy, Inc.	2004	13204	43
NRG Energy, Inc.	Texas Genco Inc.	2006	13017	42
Westar Energy, Inc.	Great Plains Energy	2018	12237	66
Investor group	TXU Corp.	2007	11116	45
Exelon Corporation	Constellation Energy Group, Inc.	2012	10790	66
PPL Corporation	E.ON AG	2010	10035	44
NRG Energy, Inc.	Edison Mission Energy	2014	9052	30
FirstEnergy Corp.	Allegheny Energy, Inc.	2011	8631	36
Investor group	Engie SA	2017	8604	39
Reliant Resources, Inc.	Orion Power Holdings, Inc.	2002	8247	85
Carolina Power & Light Company	Florida Progress Corporation	2000	7721	63
Powergen PLC	LG&E Energy Corp.	2000	7445	31
ArcLight Capital Partners, LLC	Tenaska Energy Inc.	2015	7398	79
Dynegy Inc.	Energy Capital Partners LLC	2015	7334	28
MidAmerican Energy Holdings	NV Energy, Inc.	2013	7149	52
Astoria Generating Co.	EBG Holdings LLC	2007	7143	66
Riverstone Holdings LLC	Talen Energy Corporation	2016	6941	12
Dynegy Inc.	LS Power Group	2007	5909	26
Reliant Energy Power Generation Inc.	Sithe Energies Inc.	2000	5852	61
KGen Partners LLC	Duke Energy Corporation	2004	5253	44
Emera Incorporated	TECO Energy, Inc.	2016	5131	23

Note: Largest 25 acquisitions in the fossil fuel power generation industry between 2000–2020. The columns indicate the year the transaction occurs, total production capacity that changed ownership and the total number of units that changed ownership.

Table 6: Regression Results

	Owner's Existing Assets in the Same Market	Owner's Existing Assets in Different Market	Standard Deviation of Heat Rate
	(i)	(ii)	(iii)
Late-Pre	0.004	0	0.08
	(0.006)	(0.003)	(0.061)
Early-Post	-0.004	-0.002	-0.113
	(0.005)	(0.003)	(0.057)
Late-Post	0.015	0.001	-0.161
	(0.006)	(0.004)	(0.105)
R2	0.628	0.625	0.514
# of Obs	1.4M	1.68M	1.22M
# of Acq	897	897	583
Unit FE	X	X	X
State by Month FE	X	X	X
Week FE	X	X	X

Note: This table presents the coefficient estimates from estimating Equation (5.1). Standard errors are clustered at the plant level.

Table 7: Plant Characteristics Heterogeneity Coefficients

	Gas Plant	Plant Capacity > Median	Unregulated Plant	Plant Age > Median	infra-marginal Plant
Post x Treat	0.007	-0.017	-0.023	0.015	-0.01
	(0.022)	(0.007)	(0.016)	(0.009)	(0.007)
Post x Treat x Z	0.026	0.079	0.064	0.034	0.085
	(0.022)	(0.012)	(0.018)	(0.015)	(0.013)
# of Acquired Units	897	897	897	897	897
# of Units with ($Z = 1$)	809	448	777	416	297
# of Obs	1.19	1.19	1.19	1.19	1.19
Adj. R2	0.638	0.638	0.638	0.638	0.638
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: Estimates of θ_2 from Equation (5.3) for plant characteristics. Standard errors are clustered at the plant level.

Table 8: Firm Characteristics Heterogeneity Coefficients

	Financial Acquirer	Serial Acquirer	Target Exits	Acquirer Enters	Acquirer Size > Median
Post x Treat	0.029	0.011	0.039	0.037	0.016
	(0.009)	(0.008)	(0.01)	(0.009)	(0.007)
Post x Treat x Z	0.014	0.053	-0.017	-0.019	0.028
	(0.015)	(0.015)	(0.014)	(0.013)	(0.013)
# of Acquired Units	897	897	897	897	897
# of Units with ($Z = 1$)	199	430	393	236	421
# of Obs	1.19	1.19	1.19	1.19	1.19
Adj. R2	0.638	0.638	0.638	0.638	0.638
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: Estimates of θ_2 from Equation (5.3) for firm characteristics. Standard errors are clustered at the plant level.

Table 9: Transaction Characteristics Heterogeneity Coefficients

	Acquired Capacity > Median	Acquirer's Capacity in the Market > Median	Bankrupt Sale	Year > 2010	Trans. Value > Median
Post x Treat	0.009	-0.001	0.011	0.016	0.023
	(0.008)	(0.008)	(0.007)	(0.009)	(0.007)
Post x Treat x Z	0.056	0.058	0.191	0.026	0.034
	(0.015)	(0.013)	(0.02)	(0.014)	(0.021)
# of Acquired Units	897	897	897	897	897
# of Units with ($Z = 1$)	448	407	60	374	193
# of Obs	1.19	1.19	1.19	1.19	1.19
Adj. R2	0.638	0.638	0.639	0.638	0.638
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: Estimates of θ_2 from Equation (5.3) for transactions characteristics. Standard errors are clustered at the plant level.

Figure 17: Firms with Largest Capacity Increase: 2010-2020

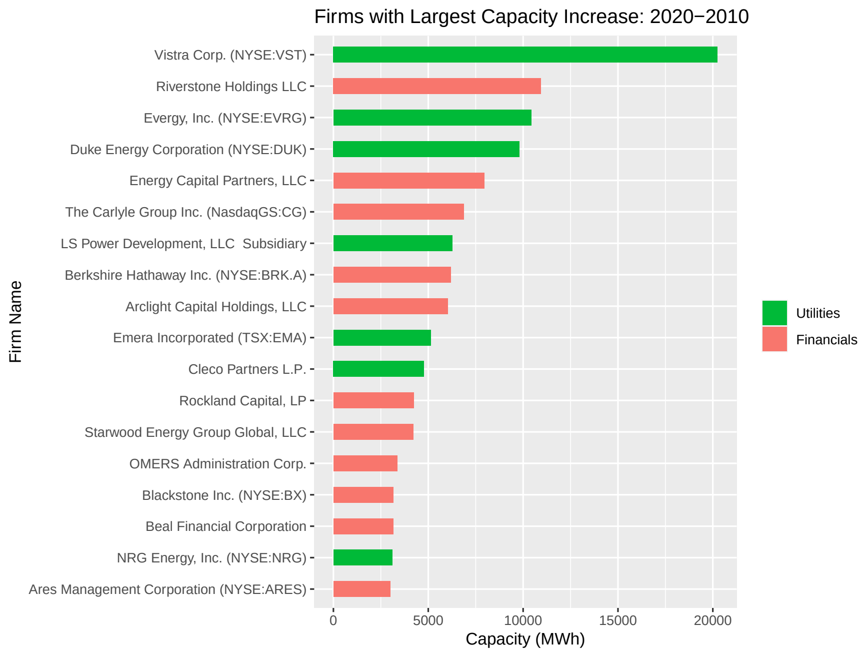

Note: This figure shows firms with the largest capacity increase in fossil fuel plants in the US between 2000 and 2020

Figure 18: Firms with Largest Capacity Decrease: 2010-2020

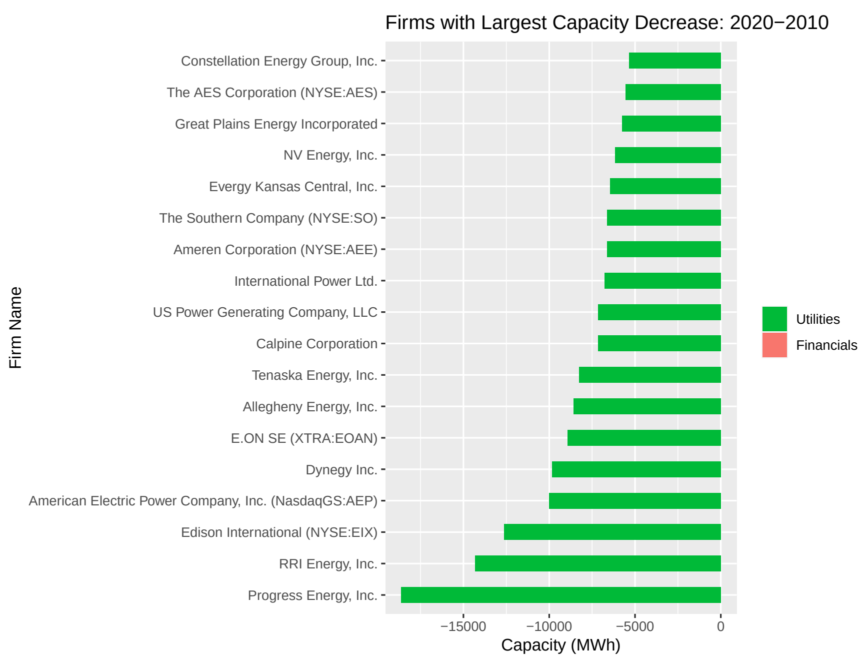

Note: This figure shows firms with the largest capacity decrease in fossil fuel plants in the US between 2000 and 2020

Figure 19: Effects Of Manager Change without Mergers on Efficiency

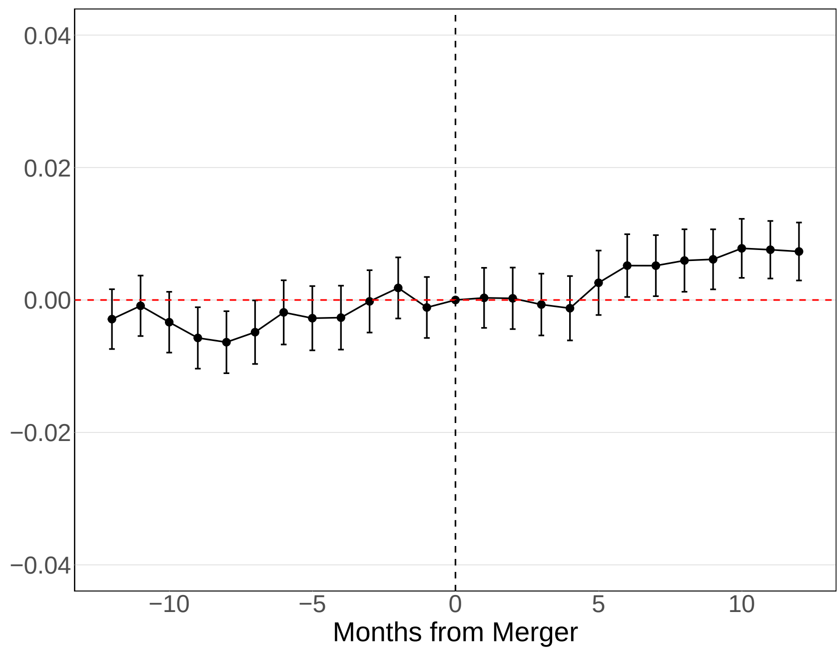

Note:

Figure 20: Case Studies of Heat Rate Improvement

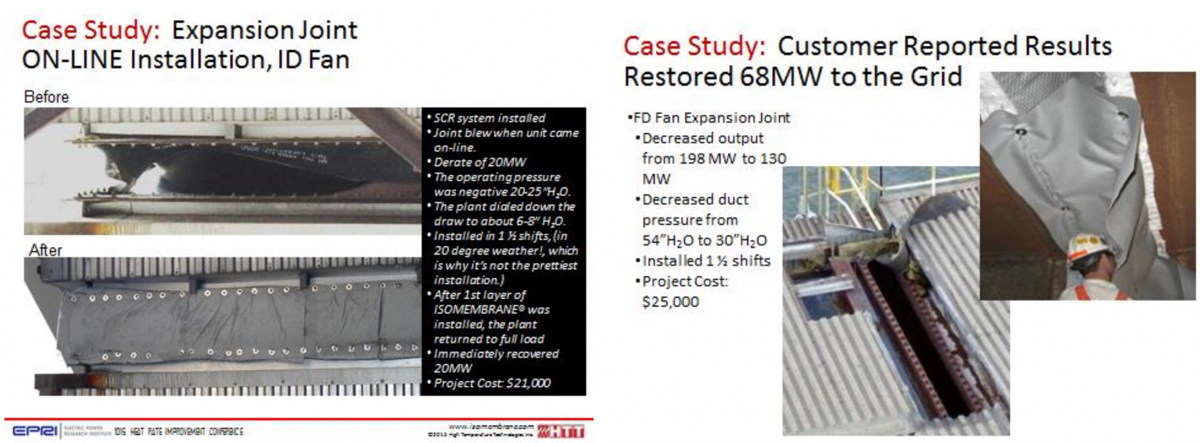

(a) a

(b) a

Note: These pictures demonstrate some methods that were implemented in power plants to improve heat rate.

Figure 21: Confidence Band for Cost Curve Differences

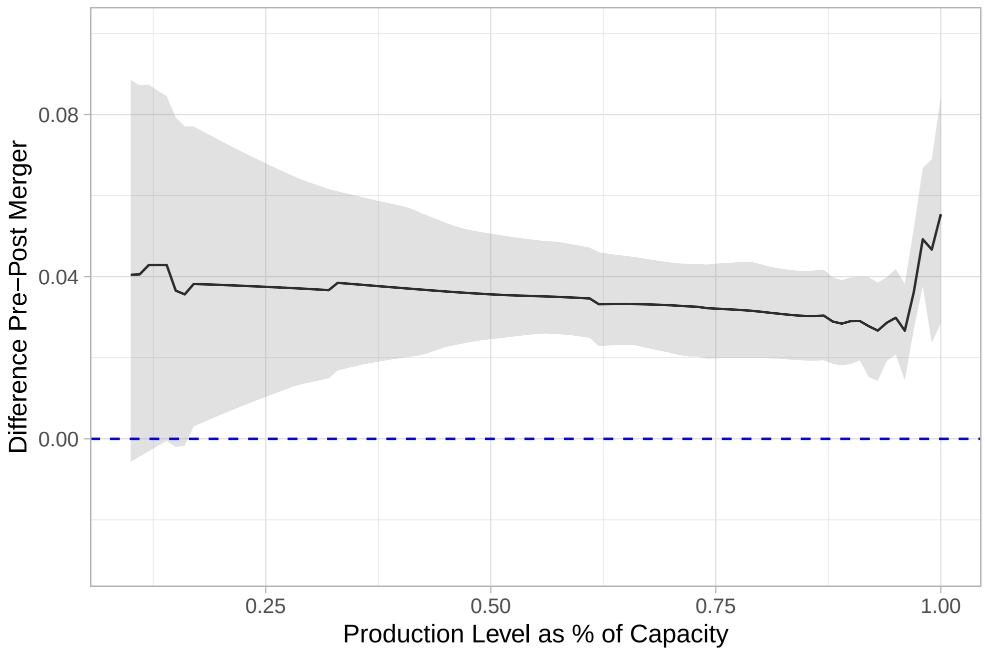

Note: This figure shows average year-to-year within-plant productivity growth for the plants that were not involved in an acquisition.

Figure 22: Average Within-Plant Annual Productivity Change

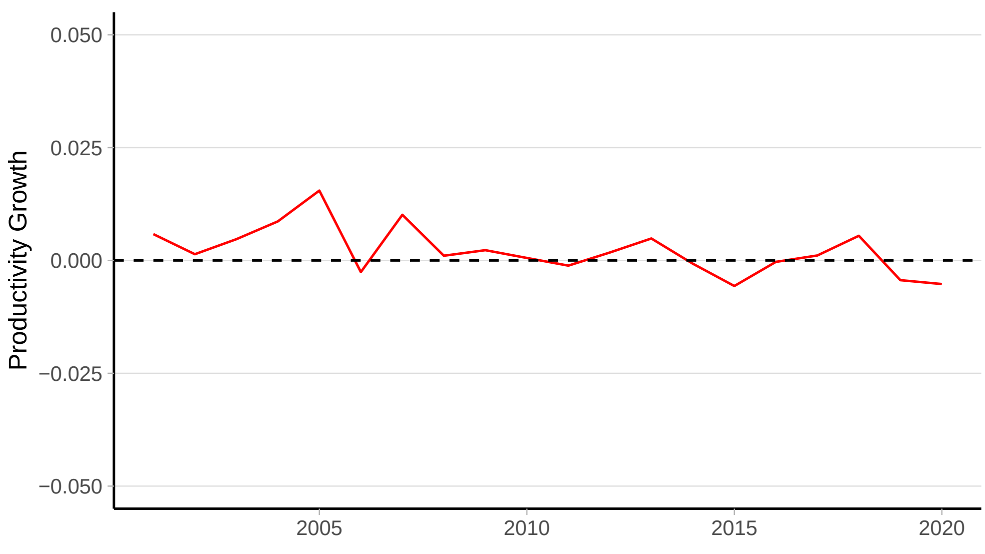

Note: This figure shows average year-to-year within-plant productivity growth for the plants that were not involved in an acquisition.

Figure 23: Change of Percentage of Fossil Fuel Generation Capacity

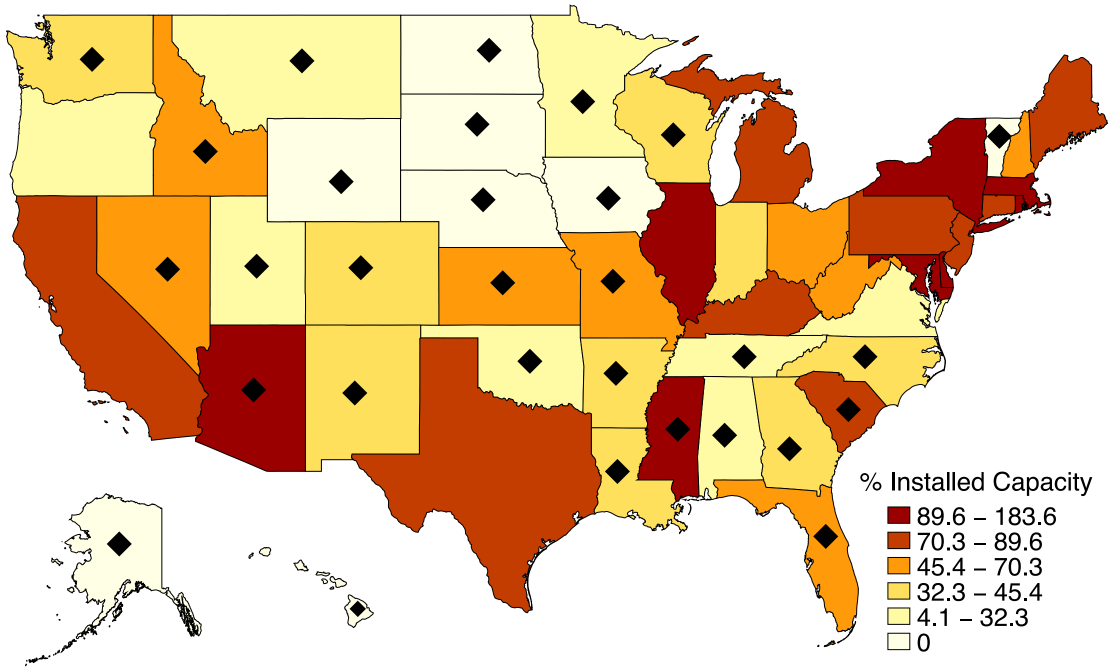

Note: Geographical distribution of power plant acquisitions by capacity. The diamond indicates the regulated states.

Figure 24: Acquiring Firms often make claims about Heat Rate improvement

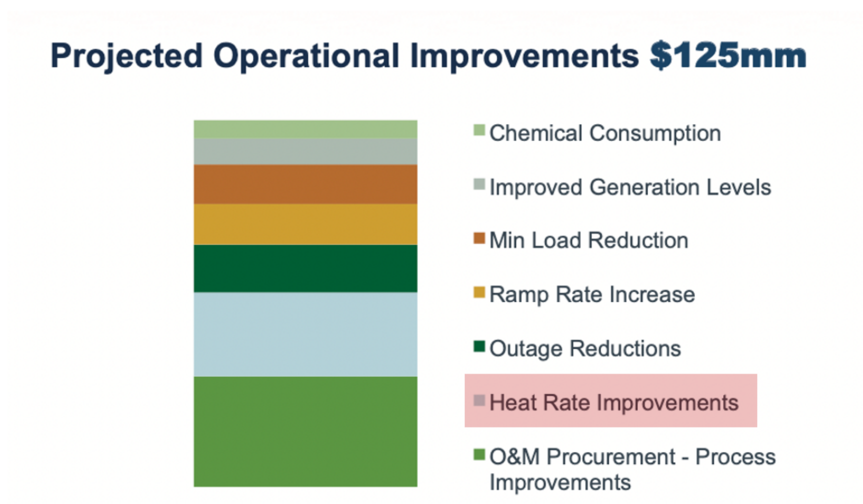

Note: This figure is from the conference call of acquisition of Dynergy by Vista Energy (2018, $1.74 billion deal)

Figure 25: Ownership Change Types

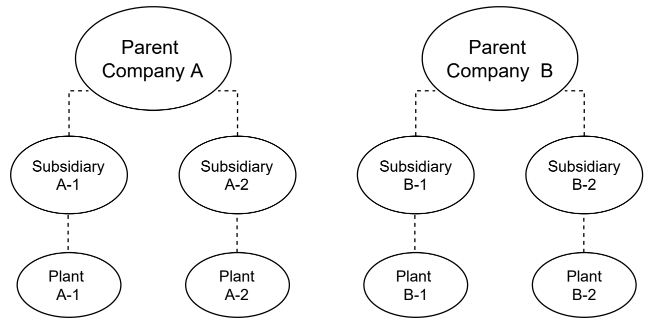

(a) Before

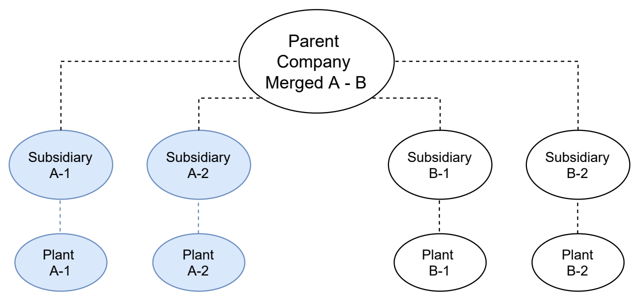

(b) Asset Sales

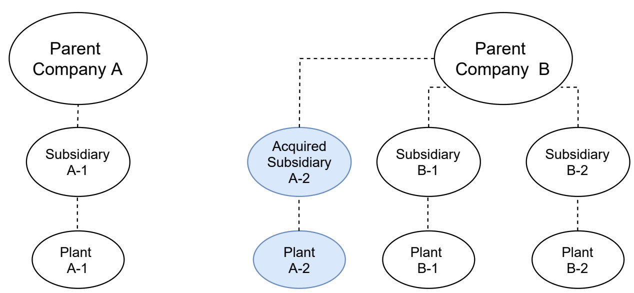

(c) Subsidiary Acquisitions

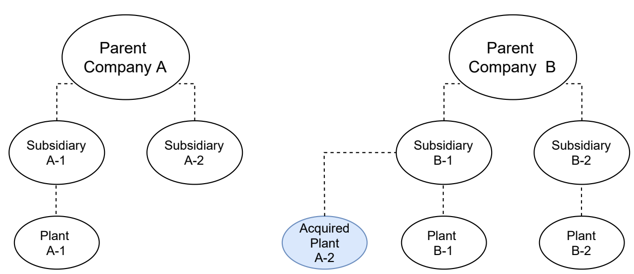

(d) Mergers

Note: These pictures demonstrate some methods that were implemented in power plants to improve heat rate.

Figure 26: Change in Market Concentration

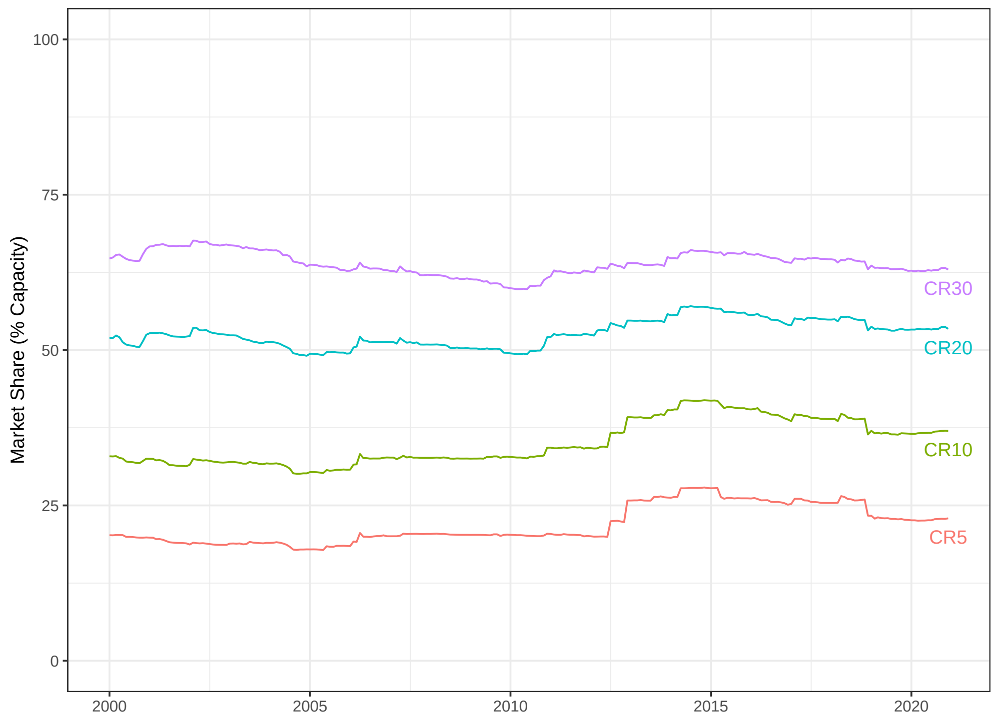

Note: This figures show change of the HHI in the overall US fossil fuel power plants market between 2000 and 2020

E Robustness Checks Results

Table 10: Impact of Merger on Productivity (Daily Data)

	All M&A	Owner/Parent Company Changes	Only Parent Company Changes	Minority Owner Changes (Placebo)	Name Changes (Placebo)
Late pre-acquisition	0.001	-0.003	-0.004	-0.005	-0.006
	(0.005)	(0.007)	(0.006)	(0.007)	(0.009)
Early post-acquisition	0	0.005	-0.003	-0.01	-0.003
	(0.005)	(0.007)	(0.006)	(0.007)	(0.006)
Late post-acquisition	0.012	0.034	-0.004	-0.007	0.006
	(0.006)	(0.012)	(0.006)	(0.01)	(0.01)
Adj. R^2	0.645	0.652	0.642	0.663	0.655
# of Obs.	8.3M	6.26M	6.48M	5.12M	5.77M
# of Acq.	1760	897	921	405	456
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: This table presents the coefficient estimates from estimating Equation (5.1). Standard errors are clustered at the plant level.

Table 11: Impact of Merger on Productivity (Hourly Data)

	All M&A	Owner/Parent Company Changes	Only Parent Company Changes	Minority Owner Changes (Placebo)	Name Changes (Placebo)
Late pre-acquisition	0.001	0	-0.008	-0.01	-0.003
	(0.005)	(0.008)	(0.006)	(0.011)	(0.011)
Early post-acquisition	-0.004	0.005	-0.008	-0.023	-0.001
	(0.006)	(0.008)	(0.008)	(0.011)	(0.008)
Late post-acquisition	0.02	0.042	0.001	-0.021	0.005
	(0.008)	(0.017)	(0.007)	(0.013)	(0.012)
Adj. R^2	0.353	0.365	0.348	0.358	0.373
# of Obs.	146.06M	111.61M	122.03M	99.15M	109M
# of Acq.	1760	897	921	405	456
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: This table presents the coefficient estimates from estimating Equation (5.1). Standard errors are clustered at the plant level.

Table 12: Impact of Merger on Productivity (Single Acquisition Units)

	All M&A	Owner/Parent Company Changes	Only Parent Company Changes	Minority Owner Changes (Placebo)	Name Changes (Placebo)
Late pre-acquisition	0.001	0.003	-0.005	0.006	-0.028
	(0.007)	(0.012)	(0.009)	(0.018)	(0.014)
Early post-acquisition	-0.001	-0.007	0.002	0.012	-0.013
	(0.006)	(0.009)	(0.008)	(0.016)	(0.012)
Late post-acquisition	-0.002	0.025	-0.017	0.001	0.002
	(0.008)	(0.016)	(0.01)	(0.019)	(0.022)
Adj. R^2	0.632	0.646	0.631	0.646	0.65
# of Obs.	1.21M	1.03M	1.13M	1.02M	1.03M
# of Acq.	489	133	311	139	112
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: This table presents the coefficient estimates from estimating Equation (5.1). Standard errors are clustered at the plant level.

Table 13: Impact of Merger on Productivity Weighted Regressions

	All M&A	Owner/Parent Company Changes	Only Parent Company Changes	Minority Owner Changes (Placebo)	Name Changes (Placebo)
Late pre-acquisition	0.001	0.002	-0.009	-0.007	-0.012
	(0.005)	(0.008)	(0.007)	(0.008)	(0.009)
Early post-acquisition	0.002	0.01	-0.006	-0.015	-0.003
	(0.005)	(0.009)	(0.006)	(0.009)	(0.007)
Late post-acquisition	0.014	0.045	-0.01	0.001	0.012
	(0.007)	(0.013)	(0.008)	(0.01)	(0.01)
Adj. R^2	0.599	0.62	0.599	0.631	0.613
# of Obs.	1.79	1.38	1.4	1.12	1.22
# of Acq.	1760	897	921	405	456
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: This table presents the coefficient estimates from estimating Equation (5.1). Standard errors are clustered at the plant level.

Table 14: Impact of Merger on Productivity All Acquisitions

	All M&A	Owner/Parent Company Changes	Only Parent Company Changes	Minority Owner Changes (Placebo)	Name Changes (Placebo)
Late pre-acquisition	0.002	0.003	-0.003	-0.005	-0.011
	(0.004)	(0.007)	(0.006)	(0.008)	(0.008)
Early post-acquisition	0.002	0.007	0.002	-0.009	0
	(0.004)	(0.007)	(0.006)	(0.007)	(0.005)
Late post-acquisition	0.016	0.036	0.003	0.001	0.011
	(0.01)	(0.009)	(0.01)	(0.008)	
Adj. R^2	0.622	0.635	0.622	0.652	0.635
# of Obs.	1.79	1.38	1.4	1.12	1.22
# of Acq.	2913	1202	1198	430	677
Unit FE	X	X	X	X	X
State by Month FE	X	X	X	X	X
Week FE	X	X	X	X	X

Note: This table presents the coefficient estimates from estimating Equation (5.1). Standard errors are clustered at the plant level.

Figure 27: Impact of Merger on Productivity (Daily Data)

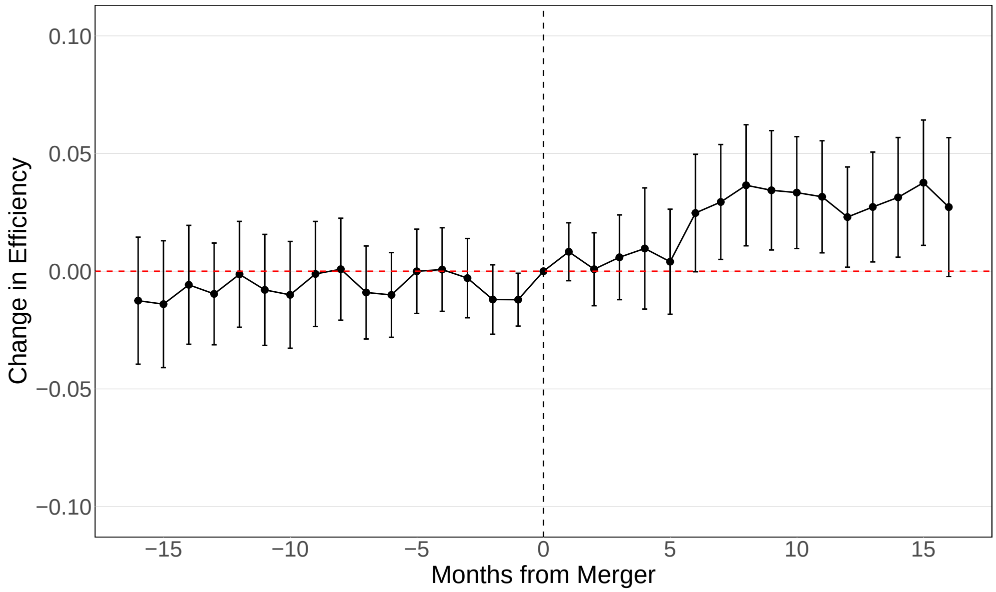

Note: The dynamic effects of acquisitions estimated from Equation (5.2). Standard errors are clustered at the plant level.

Figure 28: Impact of Merger on Productivity (Hourly Data)

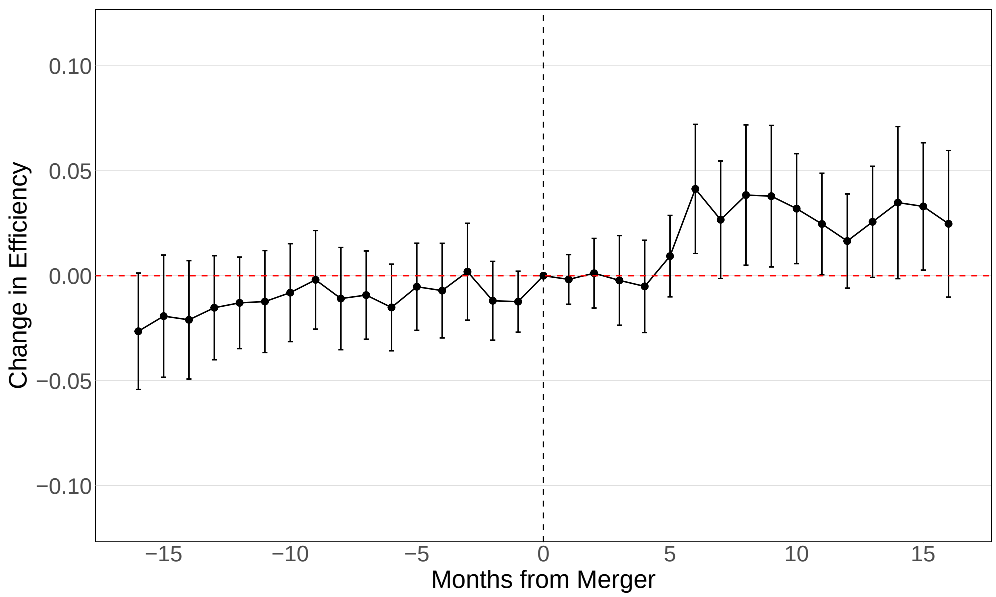

Note: The dynamic effects of acquisitions estimated from Equation (5.2). Standard errors are clustered at the plant level.

Figure 29: Impact of Merger on Productivity (Matching Estimators)

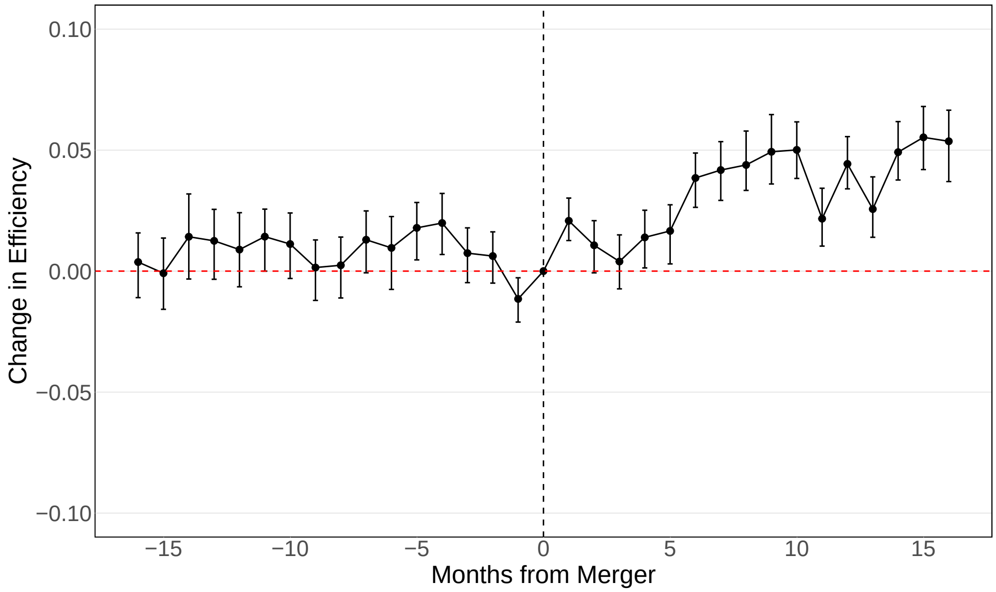

Note: The dynamic effects of acquisitions estimated from Equation (5.2). Standard errors are clustered at the plant level.

Figure 30: Impact of Merger on Productivity (All Acquisitions)

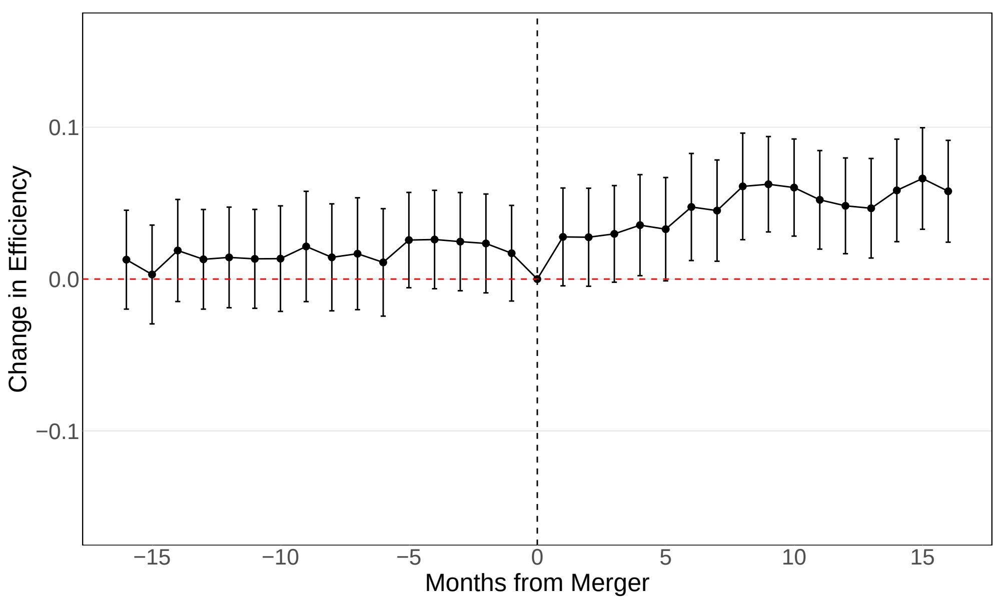

Note: The dynamic effects of acquisitions estimated from Equation (5.2). Standard errors are clustered at the plant level.

Figure 31: Impact of Merger on Productivity (Weighted By Capacity)

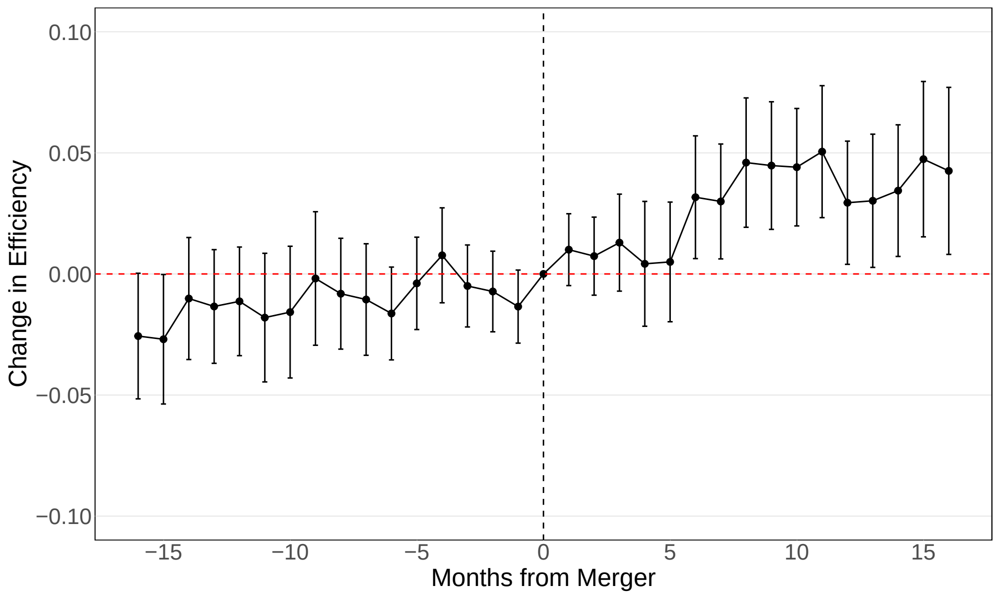

Note: The dynamic effects of acquisitions estimated from Equation (5.2). Standard errors are clustered at the plant level.