Do Mergers and Acquisitions Improve Efficiency? Evidence from Power Plants*

Mert Demirer[†] Ömer Karaduman[‡]

June 22, 2025

Abstract

Using rich data on hourly physical productivity and thousands of ownership changes from US power plants, we study the effects of acquisitions on efficiency and underlying mechanisms. We find a 2% average increase in efficiency for acquired plants, beginning five months after acquisitions. Efficiency gains rise to 5% under direct ownership changes, with no significant change when only parent ownership changes. Investigating the mechanisms, three-quarters of the efficiency gain is attributed to increased productive efficiency, while the rest comes from dynamic efficiency through changes in production allocation. Our evidence suggests that *high-productivity* firms buy *underperforming* assets from *low-productivity* firms and make them *as productive as* their existing assets through operational improvements. Finally, acquired plants improve their performance beyond efficiency by increasing output and reducing outages.

JEL: L22, L25, G34, L40

Keywords: mergers, efficiencies, productivity, antitrust, power generation

*We thank Seabron Adamson, Nikhil Agarwal, Charles Angelucci, Vivek Bhattacharya, Severin Borenstein, David P. Byrne, Steve Cicala, Lucas Davis, Natalia Fabra, Silke Forbes, Bob Gibbons, Ali Hortaçsu, Paul Joskow, Chris Knittel, Charlie Kolstad, Jacob LaRiviere, Alex Mackay, Nathan Miller, Ariel Pakes, Mathias Reynaert, Dennis Rickert, Nancy Rose, Richard Schmalensee, Chad Syverson, Mike Whinston, Frank Wolak, Catherine Wolfram, Ali Yurukoglu, and various seminar participants for helpful comments. Anna Simmons, Michel Gutmann, David Scolari, and Sam Pomerantz provided excellent research assistance. We thank Matt Barmack, Sang Ha Gang, Charlie Gates, Robert Kasle, Peter Rider, and Jon Sepich for insightful discussions on plant operations and acquisitions in the power generation industry. All remaining errors are our own.

[†]MIT Sloan. Email: mdemirer@mit.edu

[‡]Stanford GSB. Email: omerkara@stanford.edu

1 Introduction

A fundamental issue in antitrust policy is the trade-off between the market power and efficiency effects of mergers. The increase in market power raises prices for consumers; however, potential efficiency gains can counteract this effect, rendering the net impact of mergers on welfare ambiguous (Williamson, 1968). Although there is extensive literature on the price effects of mergers, we have limited evidence on how mergers affect efficiency (Whinston, 2007; Asker and Nocke, 2021).[1] With little guidance from empirical evidence, researchers analyzing the competitive effects of mergers often rely on hypothetical efficiency gains (Farrell and Shapiro, 2010; Nocke and Whinston, 2022; Berger et al., 2023).[2]

A major challenge in analyzing the efficiency effects of mergers is distinguishing true efficiency gains from other potential merger effects, such as changes in market power, buyer power, and product quality. Due to the limitations of common production datasets, most productivity studies rely on revenue-based productivity (TFPR), derived from revenues and input expenditures, rather than quantity-based measures (Foster et al., 2008; Atalay, 2014). Using TFPR is particularly problematic in merger analysis because changes in market power, buyer power, or quality can affect TFPR even without any efficiency gains. This makes it difficult to identify the true efficiency effects of mergers.[3]

In this paper, we provide large-scale evidence on the efficiency effects of mergers while tackling these challenges. We focus on acquisitions in the US electricity generation industry between 2000 and 2023. Four key features of this industry allow us to overcome the difficulties in quantifying merger efficiencies. First, we observe, at the hourly frequency, the physical quantities of both output and the primary input, the consumption of fuel, which makes up 79% of operational costs. Using this high-frequency data, we construct an efficiency measure (heat rate) and analyze how it changes around the time of acquisition. Second, electricity is a homogeneous product, eliminating potential quality changes that could confound our analysis. Third, the power generation industry experienced a significant number of acquisitions during the sample period. Our sample includes 505 transactions with 3,515 generator ownership changes, representing an average of 4.5% of the industry's annual capacity. These ownership changes exhibit significant heterogeneity in transaction, firm, and plant characteristics, which we leverage to study the mechanisms

[1] Weinberg (2008), Ashenfelter and Hosken (2010), and Kwoka (2014) provide reviews of the literature on the price effects of mergers.

[2] 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"; "...there is remarkably little solid empirical evidence on this point.".

[3] The examination of efficiencies is a standard part of merger review; see Section 3.3 of the Merger Guidelines (DOJ and FTC, 2023).

of efficiency gains. Finally, electricity generation is an important industry, with efficiencies leading to positive externalities through decreased emissions (EPA, 2018b).

Our analysis begins with a difference-in-differences estimation, comparing the efficiency of acquired plants to those not involved in acquisitions. We find that, on average, the efficiency of acquired plants increases by 2.0% after acquisitions. However, this average effect masks significant heterogeneity from the two types of ownership changes observed in the data: parent and subsidiary level ownership changes. Ownership changes only at the parent level do not change efficiency on average, whereas changes in subsidiary ownership result in a 4.9% increase. This finding highlights the influence of the direct owner in power plant operations. Finally, examining the timing of efficiency gains, we observe improvements beginning five months post-acquisition, stabilizing after eighteen months. This suggests that new owners require time to implement efficiency-improving changes.

Efficiency increases in electricity generation can manifest in various ways, not all of which are necessarily welfare-improving. For instance, generators incur additional costs when adjusting production levels, known as ramp costs, so acquirers might improve efficiency by reducing production and ramping (Borrero et al., 2024). Alternatively, they can use generators more intensively at the expense of reliability (Borenstein et al., 2023). To understand the nature of efficiency effects, we study other generator outcomes indicative of performance, including output, capacity utilization, outages, and emissions. We find that acquired generators increase generation by 7.3%, raise capacity utilization by 2.2%, and experience a 33.5% reduction in outages following the acquisition. These results suggest that acquirers improve other dimensions of plant performance beyond efficiency, and efficiency gains occur without deterioration in other performance indicators.

While evidence of efficiency gains after mergers is important, understanding the underlying mechanisms is essential for informing antitrust policy and generalizing findings from this industry to others. With this motivation, the remainder of the paper conducts a comprehensive mechanism analysis by investigating heterogeneity in the efficiency effects and modeling sources of efficiency gains in power plants.

We start by analyzing the characteristics of generators, firms, and transactions that may be informative about efficiency effects. We find that, on average, generators with above-median capacity experience a 3.3 percentage points (pp) larger efficiency increase than those with below-median capacity. This difference perhaps reflects greater incentives to improve efficiency in larger-capacity units, as any improvement in efficiency would yield higher returns. Regarding firm characteristics, efficiency improvement is 4.1 pp higher when the acquirer is larger than the median and 5.8 pp higher when the acquirer is a serial acquirer. These results suggest that a firm's experience in plant operation and acquisition is an important predictor of post-acquisition efficiency gains. Finally, we examine the differential impact of cross-market acquisitions on efficiency, finding that they result in efficiency gains 3.9 pp lower than within-market acquisitions.

We then proceed to a more structural analysis to uncover the mechanisms behind efficiency gains. We focus on two mechanisms through which a firm can increase the efficiency of electricity generation: (i) improving the performance of individual generators (productive efficiency) and (ii) optimizing production allocation dynamically within a generator to reduce ramping (dynamic efficiency). We develop a testable prediction for each mechanism and quantify their contributions by modeling the efficiency of generators. In particular, relying on a Leontief electricity production function as a microfoundation, we model the heat rate as a function of output, ramp rate, and weather conditions (temperature and humidity). The availability of hourly production data allows us to estimate this function for each generator separately for pre- and post-acquisition periods, thereby directly measuring the change in production function due to acquisitions.

We test the role of the first mechanism, productive efficiency, by quantifying the efficiency change due to shifts in the heat rate curve, controlling for ramp and weather conditions. For dynamic efficiency, we analyze the variation of an acquired plant's production, with less variation over time indicating dynamic efficiency. To quantify its contribution, we use the estimated production function model and calculate the efficiency increase due to changes in the post-acquisition production distribution while keeping the heat rate curve as in the pre-acquisition technology.

We find that productive efficiency accounts for most (three-quarters) of the total efficiency gain. The average heat rate curve of acquired generators shifts downward after acquisition at every production level, suggesting that acquirers improve generators' internal efficiency. We also find evidence supporting an increase in dynamic efficiency. Following acquisitions, generators' coefficient of variation (CoV) of production tends to decrease, explaining the remainder of the efficiency increase.

Having established the role of productive efficiency, the next natural question is how acquirers improve productive efficiency. There are two potential channels: (i) operational improvements, which involve, for example, installing control software, implementing effective maintenance, providing personnel training, or adopting best practices, and (ii) capital investments, which involve equipment upgrades. Process improvements suggest transfers of intangible capital post-acquisition (Atalay et al., 2014), whereas capital upgrades indicate either credit constraints faced by the former owner or insufficient incentives to make efficiency-improving capital investments (Midrigan and Xu, 2014). Although the large-scale nature of our study precludes us from directly observing what changes inside the plant, we supplement our efficiency data with two additional datasets to distinguish between these mechanisms: (i) data on plant managers and (ii) data on capital expenditures and non-fuel costs.

Starting with the manager data, we find that 55% of acquired power plants change managers within one year of acquisition. These new managers are 5.6 pp more likely to have a master's degree and 4.3 pp more likely to have a bachelor's degree than managers of non-acquired plants. In contrast, we find no evidence of increased capital expenditures or non-fuel costs after acquisitions. Therefore, the new owners appear to improve efficiency through low-cost operational improvements rather than high-cost capital investments. This analysis contributes to the growing body of evidence suggesting that acquisitions serve as a mechanism for transferring within-organization knowledge to newly acquired assets (Hortaçsu and Syverson, 2007; Bloom et al., 2012; Atalay et al., 2014; Eliason et al., 2020).

Efficiency gains through operational improvements point to superior capabilities of acquirers in plant operation and utilization compared to target firms. To further explore this and understand how acquisitions reallocate assets within the economy, our final analysis estimates and compares the productivity of target and acquirer firms. We find that high-productivity firms buy underperforming assets from low-productivity firms and make the acquired assets as productive as their existing assets after acquisitions. On average, acquirers are 1.7% more productive than the targets, and assets sold by the target firms underperform their other assets by 3%. These findings suggest acquisitions allocate assets to firms with both relative and absolute advantages in utilizing those assets, providing evidence for both the "high-buys-low" and "like-buys-like" theories of merger gains in the literature (Jovanovic and Rousseau, 2002; Rhodes-Kropf and Robinson, 2008).

As with all retrospective merger analyses, an identification challenge in our paper is the potential endogeneity of mergers. To address this concern, we implement several strategies and robustness checks. First, our specification incorporates a rich set of controls along with flexible time trends (fuel, technology, vintage, and state), accounting for factors that potentially influence selection into acquisitions. Second, we analyze the timing of the effect, demonstrating parallel trends between the treated and control groups three years before acquisitions and an increase in efficiency starting a few months after acquisitions. Third, we show that our results are robust to the empirical method, sample period, acquisition definition, and data frequency in measuring efficiency.

We conclude the introduction by noting that our results do not characterize the full impact of acquisitions, as we have identified only one component of the merger welfare analysis. Furthermore, the magnitude of the efficiency effect identified in this paper might not extend to industries with significantly different production techniques or acquisition motives compared to electricity generation. Although we focus on a single sector to leverage the available data and numerous acquisitions, we provide detailed evidence of mechanisms to draw broader lessons from this sector.

Contribution to the Literature This article contributes to the literature on the effects of mergers on productivity. As noted by Whinston (2007) and Asker and Nocke (2021) in the two latest IO handbooks, there is a limited number of papers on the productivity effects of mergers. Among these studies, Blonigen and Pierce (2016) apply the De Loecker and Warzynski (2012) method to separately identify the effects of mergers on market power and productivity in US manufacturing plants. Their findings suggest an increase in market power but no evidence of a productivity effect. Kulick (2017) studies mergers in the US ready-mix concrete industry, finding that prices rose due to increased market power despite a 6% productivity increase in acquired plants. Braguinsky et al. (2015) examine the effects of consolidation in the early 20th-century Japanese cotton spinning industry. They find that although acquirers are not more productive conditional on operation, they are more profitable due to better inventory management and higher capacity utilization. After acquisitions, the acquirers improve capacity utilization in the acquired plants, raising both productivity and profitability.[4]

This article contributes to the literature studying efficiency in the power generation industry, which has focused mainly on the effects of deregulation that began in the 1990s (Knittel, 2002; Bushnell and Wolfram, 2005; Fabrizio et al., 2007; Davis and Wolfram, 2012; Hausman, 2014; Cicala, 2015, 2022). These studies compare the performance of plants in states that pursued restructuring to those in states that did not, generally finding a positive impact of restructuring on plant operations.[5] Our paper differs from this literature, as we analyze the effects of ownership changes rather than deregulation on productivity.[6] We primarily focus on the post-deregulation period and exclude forced divestitures due to

[4]Evidence of cost savings from other industries includes meat products (Nguyen and Ollinger, 2006), railroads (Bitzan and Wilson, 2007; Chen, 2024), electricity distribution (Kwoka and Pollitt, 2010; Clark and Samano, 2022), radio (Jeziorski, 2014), banking (Focarelli and Panetta, 2003), and healthcare (Dranove and Lindrooth, 2003; Harrison, 2011; Schmitt, 2017). These studies typically analyze firm costs, which include both input prices and firm productivity. Another strand of literature provides evidence on efficiency effects by analyzing a single merger. Some examples are the Molson and Coors merger (Grieco et al., 2018) and the Miller and Coors merger (Ashenfelter et al., 2015) in the brewing industry, and the Boeing-McDonnell Douglas merger (An and Zhao, 2019) in the aerospace industry.

[5]MacKay and Mercadal (2024) find that, despite decreasing generation costs, wholesale prices increased due to market imperfections.

[6]Bushnell and Wolfram (2005) also study the impact of ownership changes on power plant efficiency. Their study focused on utility divestitures in the context of industry deregulation. By contrast, our study examines ownership changes that occurred after deregulation.

deregulation from our sample. In the literature studying electricity markets, our paper also relates to Hortaçsu and Puller (2008) and Hortaçsu et al. (2019), who study the heterogeneity of firms' strategic bidding ability in wholesale electricity auctions and how ownership changes potentially affect strategic ability.

Finally, this paper contributes to the recent wave of retrospective merger research that examines the impact of mergers on various firm outcomes. This literature has advanced our understanding of cross-market mergers (Lewis and Pflum, 2017; Dafny et al., 2019), vertical mergers (Luco and Marshall, 2020), monopsony power (Prager and Schmitt, 2021), buyer power (Craig et al., 2021), price effects (Bhattacharya et al., 2024; Brand et al., 2023; Brot-Goldberg et al., 2024), quality (Eliason et al., 2020; La Forgia, 2024), product availability (Atalay et al., 2024), firm entry (Fan and Yang, 2025), capacity utilization (Kalnins et al., 2017), employment (Geurts and Van Biesebroeck, 2019), and political influence (Moshary and Slattery, 2024). We complement this body of work by studying how mergers affect efficiency and providing evidence of the mechanisms.

2 Institutional Background and Plant Productivity

This section begins with the institutional background of the power generation sector, followed by an overview of mergers and acquisitions in the industry. We then discuss power plant operation and our approach to measuring plant productivity.

2.1 The Power Generation Sector in the US

The US electric power sector accounts for roughly 2% of the US GDP (Bradley & Associates, LLC, 2017). Before the 1990s, US electricity generation was predominantly supplied by regulated and vertically integrated utilities. These entities typically served a specific territory and controlled all components of the sector, including generation, transmission, and distribution. The returns of these utilities were regulated through rate-of-return on capital investments and cost-of-service regulation (Joskow et al., 1989). This highly regulated market structure provided minimal incentives for efficiency improvements, leading to significant inefficiencies in electricity generation (Fabrizio et al., 2007; Cicala, 2015).

In the 1990s, the industry underwent significant deregulation. In many states, electricity generation was decoupled from transmission and distribution, with most generators transitioning to market-based compensation. This shift coincided with the establishment of independent system operators (ISOs), which manage the electricity grid and organize the wholesale market for electricity. By 2020, about 70% of US electricity demand was served through seven ISOs (EIA, 2020).[7]

2.2 Acquisitions in the Power Sector

Large power companies are often structured into multiple subsidiaries under a single parent company, each serving distinct locations and segments of the power sector. These parent companies frequently own assets in generation, transmission, and distribution within the same region, although some operate subsidiaries across various parts of the country. Following the deregulation wave in the 1990s, there was a notable increase in mergers and acquisitions among these entities (Davis and Wolfram, 2012). Moreover, financial firms, particularly private equity firms and bank funds, began investing heavily in the power generation sector (Andonov and Rauh, 2023).

Acquisitions in the power sector can be categorized into two types: (i) asset acquisitions and (ii) subsidiary acquisitions. Asset acquisitions involve a firm selling parts of its power plant portfolio, with the acquired assets placed under a subsidiary of the acquiring company. Subsidiary acquisitions occur when a parent company acquires another company's subsidiary, including all its assets. In asset acquisitions, both parent and subsidiary owners change, whereas in subsidiary acquisitions, only the parent owner changes.[8] For a visual explanation of these acquisition types, see Figure OA-3.

Proposed power plant acquisitions in the US electricity sector are subject to review by the Federal Energy Regulatory Commission (FERC), the Department of Justice (DOJ), and state Public Utility Commissions (PUC) (Niefer, 2012). FERC conducts its review under Section 203 of the Federal Power Act to determine if the merger aligns with the public interest (FERC, 2012). The DOJ's review focuses on the potential anticompetitive effects. If either the DOJ, FERC, or PUC finds consumer harm, they may challenge the merger or require remedies.[9] Despite reviews by three government agencies, most proposed power plant mergers over the past two decades have gained approval (Hempling, 2018).

Firms cite various motives for acquisitions, including synergies, financial benefits, and complementarities between different asset types.[10] Since fuel represents a large portion of

[7] We use ISO as an umbrella term for both ISOs and regional transmission organizations.

[8] In some cases, two companies merge to form a new entity, and power plants become part of this new entity. These cases typically fall under subsidiary acquisitions.

[9] To give some examples, in 2005, the Exelon-PSEG merger was not completed after failing to get approval from the New Jersey PUC (Morris and Oska, 2008; Wolak and McRae, 2008). In 2012, following the DOJ's request, Exelon Corporation and Constellation divested three plants in Maryland (Bushnell et al., 2012).

[10] For many acquisitions in our sample, we accessed investor presentations and conference calls, which allowed us to identify the stated motives. Examples include (i) improvements in management (AES-DPL merger), (ii) cost synergies of $175 million per year (NRG-GenOn merger), (iii) annual cost savings of $150 million (Mirant-RRI Energy merger), and (iv) benefits of geographic, fuel, market, and earnings diversification (Vistra-Dynegy merger). Other motives include increasing the consumer base, diversifying the portfolio

operational costs, fuel efficiency improvements are often cited as a primary source of cost savings post-acquisition.[11]

2.3 Electricity Production and Efficiency Measure

A major challenge in analyzing merger efficiencies is the scarcity of suitable data, as most industries lack reliable cost and physical productivity measures (Asker and Nocke, 2021). The US power generation industry is unusual in this respect because of the availability of high-frequency fuel efficiency data. This section describes the efficiency measures used in this study and explains the production process at power plants.

A power plant is an industrial facility that generates electricity. As of 2020, there were 11,070 utility-scale electric power plants in the US (EIA, 2020). Typically, power plants include multiple generators, transforming different forms of energy (primarily heat, wind, or solar) into electricity using various production technologies. Our research focuses on fossil fuel generators as their efficiency is more easily measured with available data.

Fossil fuel generators produce electricity using the heat energy released from burning fuels (coal, natural gas, and oil).[12] In this process, the input is measured as the heat content of the fuel used in generation, while the output is measured as the electricity generated. This leads to a natural efficiency metric, called *heat rate*, which indicates how efficiently a generator converts fuel into electricity. Heat rate is calculated as the ratio of the fuel's heat content, in million British thermal units (MMBtu), to the generator's electricity output in megawatt-hours (MWh). Our measure of efficiency is the inverse of heat rate:

$$\text{Fuel Efficiency (Inv. Heat Rate)} = \frac{\text{Energy Output (MWh)}}{\text{Energy Input (MMBtu)}}. \tag{1}$$

Heat rate is the critical determinant of generator efficiency since fuel is the major input, representing 79% of operating costs.[13] For this reason, it is a standard efficiency metric in the industry, commonly used by regulatory agencies and firms (EPRICA, 2014; EIA, 2015).

Most importantly for this paper, fuel efficiency provides key advantages in analyzing the efficiency impacts of acquisitions. First, fuel efficiency is a quantity-based measure derived from input and output quantities rather than from revenues and input expenditures. Consequently, it is not directly affected by changes in input or output prices due to

across technologies and regions, and accelerating efforts to comply with potential future environmental regulations.

[11] As an example, Figure OA-1 shows a slide from an investor presentation for the 2018 Dynegy and Vistra Energy merger, where firms claim that heat rate improvements will lead to $30 million cost savings.

[12] In thermal power plants, water is heated to generate steam, which moves through a turbine attached to a shaft. As the steam flows, it causes the shaft to spin, driving a generator that produces electricity.

[13] Based on the authors' calculations; see Section B.5 for methodology.

Figure 1: Representative Heat Rate Curve

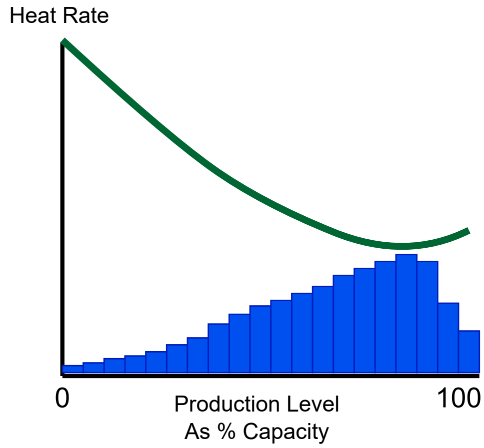

Note: The green line represents the typical heat rate curve in electricity generation, showing how heat rate changes with the production level. The blue bars represent a hypothetical distribution of production as a function of generator capacity.

buyer and market power effects, allowing us to distinguish efficiencies from other merger-induced changes. Second, electricity is a homogeneous product, precluding any potential impacts on quality.[14] Finally, the efficiency measure relies primarily on sensor data rather than survey responses, as is common in many other industries.[15]

Several factors can influence the heat rate in a power plant. Figure 1 displays a hypothetical example of a heat rate curve, where the green line represents the heat rate and the blue bars represent a typical production distribution as a percentage of capacity. As suggested by the heat rate curve, a power plant's efficiency depends on its production level, typically reaching its peak when operating near full capacity. Moreover, fluctuations in production significantly affect efficiency. Given that electricity cannot be stored on a large scale and demand varies over time, power plants must frequently adjust their production in response to changing market conditions. These adjustments, known as ramp costs, reduce the overall efficiency of electricity generation (Borrero et al., 2024).

Although electricity generation may appear relatively mechanical, the efficiency of generators in the US shows notable variation as in other manufacturing industries (Syverson, 2011). Figure 2 shows the distribution of annual residual log productivity of generators after controlling for a rich set of observables, including ramp, generator age, fuel type,

[14]Some post-acquisition changes, such as reliability and environmental performance, might be viewed as aspects of the 'quality' of electricity generation. We will analyze these aspects later in the paper.

[15]It is worth noting that our efficiency measure is fuel efficiency rather than TFP and does not account for non-fuel inputs. While non-fuel inputs play a less significant role in electricity generation compared to other manufacturing industries, and substitution from fuel to other inputs is limited (Fabrizio et al., 2007), we explore them in Section 6. We also provide a theoretical foundation of fuel efficiency based on a Leontief production function in Section 5.2.

Figure 2: Distribution of Residual Log Productivity

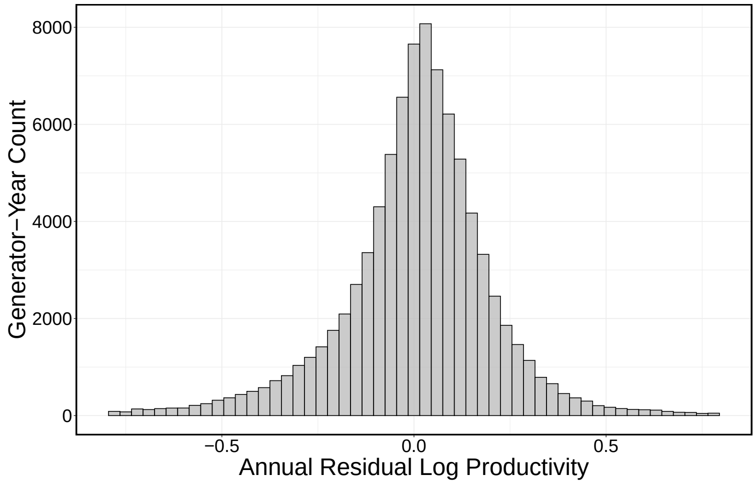

Note: This figure shows the distribution of residual yearly log productivity of fossil fuel generators in the US between 2000 and 2023, after controlling for year, standard deviation of heat rate, generator age, fuel type, technology, capacity, boiler manufacturer and model.

technology, capacity, boiler manufacturer, and model.[16] The difference between the 90th and 10th percentiles of log productivity is 0.46, indicating that generators in the top decile are 58% more productive than those in the bottom decile.[17] The large dispersion in residual productivity highlights the role of unobserved heterogeneity in efficiency and suggests that there is substantial room for efficiency improvements in many power plants.

Improving the heat rate performance of a generator can be achieved in two main ways: (i) operational improvements and (ii) capital upgrades. Operational improvements, generally lower cost than capital upgrades, include a range of practices such as installing control software, continuously monitoring unit and equipment performance, promptly repairing equipment impacting heat rate, training personnel, and implementing effective maintenance.[18] Every year, power plant managers convene at the Heat Rate Improvement Conference to discuss these practices (EPRI, 2022).[19] The second approach to improving

[16] We provide the details of this estimation procedure in Section B.1.

[17] The 90-10 percentile ratio of 1.58 is smaller than typical findings in other manufacturing sectors (Syverson, 2011), likely because we condition on a richer set of observables than in other settings. Other researchers have also observed the heterogeneity in power plant productivity. Sargent & Lundy, LLC (2009), in a study commissioned by the EPA, finds that the heat rates of coal-fired power plants range from 5 to 32.7 MMBtu/MWh. Staudt and Macedonia (2014) examine factors contributing to heat rate variation, including facility size, capacity factor, and coal type, and find considerable unexplained variability in heat rate.

[18] Several software products are available to monitor and improve power plant performance, such as PI Data Historian, EtaPRO/Virtual Plant, and Emerson Enterprise Data Server. Heat rates can also be improved with turbine enhancements such as blade and seal repairs, cycle control optimization, boiler improvements, and deposit removal. Boiler improvements involve heat transfer surface maintenance, burner system inspection, and intelligent soot blower utilization. For methods of heat rate improvements and other examples, see EPRI (2009); EIA (2015); Emerson Process Management (2016); Environmental Defense Fund (2017).

[19] Figure OA-4 highlights a few case studies of heat rate improvements from the 2015 conference, as reported

plant efficiency is by upgrading critical equipment, such as boilers, fuel feeders, and cooling systems, as old equipment deteriorates and new technology becomes available.

A critical factor influencing operational practices in a power plant is managerial and engineering input. As documented in detail in Bushnell and Wolfram (2009), the skills of key personnel can profoundly impact power plant performance. These personnel are responsible for continuously monitoring unit and equipment performance, conducting periodic tests to assess equipment condition, and planning production and maintenance schedules. Bushnell and Wolfram (2009) notes the operator's impact as follows: "the act of balancing all of these input parameters was described by one manager as playing the piano and one star operator was considered a virtuoso on the instrument".[20]

Improving power plant efficiency is also crucial for environmental reasons. The higher a plant's efficiency, the less fuel it requires, directly leading to reduced emissions of local pollutants and greenhouse gases. As a result, enhancing fuel efficiency can be an effective method to mitigate emissions, a fact acknowledged by policymakers in the EPA's Clean Power Plan (EPA, 2018b).

3 Data and Summary Statistics

Our primary objective is to compile a dataset that allows us to construct a measure of generator efficiency and identify ownership changes. In this section, we outline our data sources and present summary statistics.

3.1 Data Sources

We combine data from the FERC, the Environmental Protection Agency (EPA), the Energy Information Administration (EIA), the North American Electric Reliability Corporation (NERC), S&P Global Market Intelligence (GMI, formerly SNL Financial), and Velocity Suite at the firm, plant, and generator levels for fossil fuel-fired power plants in the continental US from January 2000 to March 2023. This section briefly describes the datasets, while Appendix A provides more detailed information on the data sources, variable construction, and descriptive statistics.

Generator and Plant Level Data. We use data from EIA, EPA, FERC Form 1, Velocity Suite,

in Fitzgerald and Gelorme (2015). The following quote is particularly noteworthy: "For years we've talked about heat rate, but let's be honest, in reality, it hasn't driven maintenance and operational activities to a great degree".

[20] As another example of the importance of personnel, PacifiCorp Energy states in their 2016 Heat Rate Improvement Plan Document that "Continuous improvement and management of unit heat rates is the responsibility of all plant personnel" and "good management of heat rate requires that plant management make optimizing heat rate a priority each day" (PacifiCorp Energy, 2016).

and GMI to construct generator- and plant-level datasets. For generators, the information includes the installation year, fuel type, technology type, capacity, boiler model, and boiler manufacturer. For plants, we construct data on regulation status, location, ISO, and FERC region. In addition, for approximately 35% of power plants, we have information on the number of employees, non-fuel costs, and capital expenditures.[21]

Production Data. We use the EPA's Continuous Emissions Monitoring Systems (CEMS) to obtain hourly input and output data. This dataset provides generation, heat input, and various emissions for nearly all fossil fuel units in the US.[22] Additionally, CEMS provides information on the environmental programs each generator is subject to and the scrubbers used for various pollutants. We merge this dataset with our generator- and plant-level data as detailed in Appendix A.1.

Ownership and Acquisition Data. We construct a dataset on fossil fuel generator acquisitions by combining two separate datasets on ownership and transactions from GMI, as well as information from company press releases and newspaper articles.[23].The ownership data includes all shareholders and their shares at both subsidiary and parent company levels. The transaction data provides details on transferred assets, transaction size, buyer and seller, announcement and closing dates, conference call transcripts, and descriptions. Given that virtually all power plant acquisitions in this industry require notification to regulatory agencies, this dataset provides comprehensive coverage of transactions during the study period.[24] It is well-known that ownership datasets may misidentify acquisitions by interpreting firm name changes and restructurings as ownership changes (Davis et al., 2024; Arora et al., 2021). We address this issue by cross-referencing transaction and ownership data and reviewing transaction descriptions, press releases, and news articles as detailed in Appendix A.5.

Maintenance and Outage Data. We obtain event-level data on outages, capacity reductions (derates), and maintenance from the Generating Availability Data System (GADS)

[21]The data sources for capital expenditures and non-fuel inputs are FERC Form 1 and Rural Utilities Service (RUS) Form 12, which are available only for major electric utilities as defined by FERC.

[22]Every fossil-fuel generator in the US with a capacity greater than 25 MW must comply with the EPA CEMS program. This sample represents approximately 95% of the US fossil fuel generating capacity.

[23]GMI, previously known as SNL Financial, collects data for the US electricity sector using regulatory filings from agencies like the Securities and Exchange Commission (SEC), FERC, Rural Utilities Service (RUS), EIA, and state-regulated utilities (GMI, 2024). Additionally, it uses news aggregators to capture news articles, press releases, and corporate announcements. GMI has been widely used by researchers to study electricity markets (Davis and Hausman, 2016; Jha, 2020; Abito et al., 2022; Borenstein and Bushnell, 2022)

[24]Before 2019, all power plant transactions required FERC approval regardless of size. On February 21, 2019, the FERC issued a rule setting a $10 million threshold for approval and a $1 million threshold for notification within 30 days. See https://www.ferc.gov/news-events/news/ferc-issues-final-rules-revising-utility-merger-hydropower-regulations, accessed on June 30, 2024.

database through a data-sharing agreement with NERC. This dataset covers all generators with a capacity over 20 MW, which are required to report events affecting their generation capabilities to NERC. Available from 2013 to 2021, the data includes each event's start and end times, type, and cause. The generator names are anonymized in this dataset, but information on capacity, state, fuel type, and monthly production hours is available. Using this information, we match this data to CEMS units using an algorithm described in Section A.7, achieving a match rate of 92.8% based on capacity.

Personnel Data. We compile monthly data on plant personnel from 2000 to 2020 using an EPA database of plant representatives, including names, tenure start and end dates, and contact information. We successfully matched about 70% of the personnel names to their LinkedIn profiles, thereby obtaining their title and education. Using LinkedIn data, we verified that 78% of the listed personnel are plant managers, while the remainder are primarily environmental compliance personnel and engineers. Thus, we consider plant representatives to be plant managers for the purposes of this study.

Other Datasets. We collect hourly data on ambient temperature and humidity from Velocity Suite for power plants in our sample, as weather conditions can affect generation performance. We also obtain firm-level industry classifications and the publicly listed status from GMI.

3.2 Construction of the Generator and Acquisition Sample

Our initial sample includes all electricity generators in the contiguous US that operated between January 2000 and March 2023 and are subject to CEMS regulations (5,876 generators). From this set, we exclude cogenerators that produce both steam and electricity, reducing the sample to 5,264 generators.

For acquisitions, we start with 5,216 generator acquisitions involved in any firm-to-firm transaction between January 2000 and March 2023. We eliminate acquisitions before a unit becomes operational and after its retirement (534) and minority acquisitions where less than 50 percent of the shares change ownership (864).[25] Next, we eliminate the ownership changes due to restructuring that happened mostly in the early 2000s because the source of efficiency improvements in these cases relates to incentive changes instead of acquisitions (Cicala, 2015; Fabrizio et al., 2007). To do this, we primarily use data from Cicala (2015), Abito et al. (2024), and EIA Electricity Monthly Reports. We further supplement this data with the regulatory status from EIA, Velocity Suite, and S&P Global and eliminate

[25] A retired power plant may change ownership due to the value of its land or salvageable equipment or to transfer environmental cleanup responsibilities to the new owner.

Figure 3: Summary Statistics on Power Plant Transactions

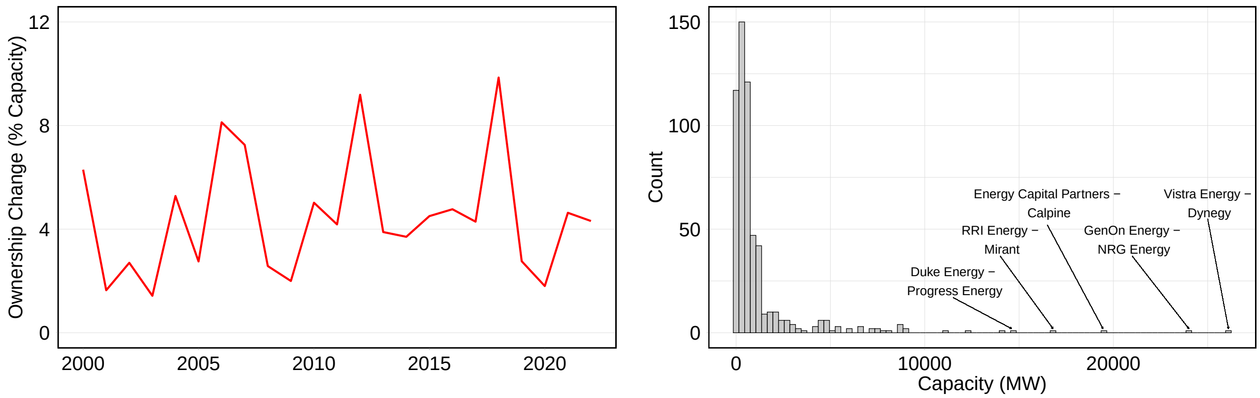

(a) Share of Capacity with Ownership Change

(b) Distribution of Transaction Size

Note: Panel (a) shows the annual percentage of fossil fuel electricity generation capacity that changed ownership in the US from 2000 to 2023. Panel (b) displays the distribution of transaction sizes based on fossil fuel generation capacity in the US during the same period. In Panel (b), the unit of observation is a transaction, with the five largest transactions labeled.

all ownership changes that result in a change in regulatory status within a 15-month window. Appendix A.4 provides further details of this procedure. This step eliminates 615 acquisitions, reducing our sample to 3,515 generator acquisition events.

3.3 Descriptive Statistics on US Power Plant Acquisitions

This section presents descriptive statistics on fossil fuel power plant acquisitions. We demonstrate that the industry has undergone a substantial number of acquisitions, with significant heterogeneity in transaction, firm, and plant characteristics.

Figure 3(a) shows the share of fossil fuel electricity generation capacity that changed ownership between 2000 and 2023.[26] On average, 4.5% of the industry capacity changes ownership annually, with some year-to-year fluctuations. As seen in Figure 3(b), these transactions vary widely in generation capacity. While most transactions include a few plants, there are some moderately-sized transactions involving 5,000–10,000 MW capacity, as well as several large ones over 10,000 MW capacity.[27] This variation indicates that our evidence is not solely from a few large mergers, and we can test the heterogeneity of the effect by different transaction characteristics.[28]

[26]We define an acquisition as a change in ownership when a different firm gains the majority of the generator's shares post-acquisition. In a small number of cases where no firm holds more than 50% of the shares, an acquisition is defined as a change in the largest shareholder.

[27]Table OA-1 lists the 25 largest transactions during the sample period.

[28]Despite many acquisitions in the study period, there has been no significant change in market concentration in the US as shown in Figure OA-5, which reports the national market shares of the largest 5, 10, 20, and 30 firms by capacity owned. The concentration fluctuates but remains broadly stable due to significant entry and exit in the industry. Some examples can be seen in Figures OA-7 and OA-8, where we report firms with

Table 1 presents summary statistics on generators, firms, and transaction characteristics. Our sample consists of 505 transactions, covering 3,515 distinct acquisition events that involve 2,048 unique generators. Most of these generators are gas-fired (82%) and operate within an ISO (77%). About half of the acquisitions are cross-market transactions, where the acquirer does not own existing capacity in the acquisition market. The sample includes 244 unique acquirers and 224 unique target firms, with acquirers owning slightly more units than targets on average.

In Column (3), we present the same statistics, but this time for the first acquisition of each generator, which forms our baseline sample in the empirical analysis. The observable unit characteristics are broadly similar between this subsample and all acquisitions (Column (2)). Comparing acquired generators in Column (3) with all generators in our sample (Column (1)), we find no meaningful differences in capacity, average installation year, and whether they operate in an organized market. However, we note differences in fuel type, with acquired generators more likely to be gas-fired (77% vs 71%). This trend primarily comes from the large number of coal power plant retirements in the 2010s, fewer acquisition opportunities due to coal plants being more likely to be in regulated states, and the uncertainty about the future of coal power plants (Davis et al., 2022). To address potential identification challenges arising from this and other potential differences, we control for monthly trends by fuel type, technology, capacity, and installation year in our empirical specifications.

Finally, the last two columns categorize the generators into two acquisition types we identified: those involving both subsidiary and parent ownership changes and those involving only parent ownership changes. Typically, a subsidiary is the legal entity that owns the power plant, while the parent company owns the subsidiary. Some transactions (asset acquisitions) involve changes in both subsidiary and parent ownership, whereas others (subsidiary acquisitions) involve changes only in parent ownership. Columns (4-5) list summary statistics for each generator's first acquisition of each type. We observe that these transaction types differ mainly in size, with parent-only ownership changes being significantly larger (an average of 15 vs 5 units). This is consistent with the nature of parent-only ownership changes, which often involve taking over a large part of the target's portfolio.

We next document the firm composition in the industry. Figure 4 displays the evolution of ownership shares by the primary activity of the company (utilities, industrials, and financials) and by company type (publicly listed, private, government-owned). Panel (a)

the largest capacity increase and decrease between 2010 and 2023.

Table 1: Summary Statistics

	All Units (1)	All Acquisitions (2)	First Acquisitions (3)	Subsidiary/Parent Change (4)	Only Parent Change (5)
			Panel A. Generator Characteristics		
# of Units	5264	3515	2048	1089	1142
# of Plants	1581	1223	744	380	405
# of Unique Units	5264	2048	2048	1089	1142
# of Unique Plants	1581	726	726	373	400
% Gas	0.71	0.82	0.77	0.89	0.68
% Coal	0.18	0.09	0.12	0.04	0.17
% Other Fuels	0.10	0.09	0.11	0.07	0.15
% Cross-Market	-	0.57	0.58	0.51	0.57
% in ISO	0.69	0.77	0.74	0.81	0.70
Avg. Unit Capacity	173.04	164.43	171.28	155.77	171.16
	(184.75)	(159.01)	(173.39)	(145.02)	(179.41)
Avg. Installation Year	1986.37	1990.59	1989.29	1994.60	1984.43
	(20.08)	(16.24)	(17.27)	(14.25)	(17.86)
			Panel B. Firm Characteristics		
# of Acquirer Firms	-	244.00	182.00	126.00	61.00
# of Target Firms	-	224.00	159.00	111.00	70.00
Avg. # of Units Acquirer Owns	-	45.81	45.94	50.72	39.35
	-	(53.40)	(49.20)	(54.88)	(39.54)
Avg. # of Units Target Owns	-	32.13	33.91	37.51	38.07
	-	(47.23)	(49.57)	(50.64)	(53.30)
Avg. Acquirer Firm Capacity	-	5244	5595	6369	6391
	-	(8698)	(9112)	(9605)	(9507)
Avg. Target Firm Capacity	-	7314	7466	8312	6200
	-	(9862)	(9335)	(9948)	(7532)
			Panel C. Transaction Characteristics		
# of Transactions	-	505	318	213	72
Avg. Transaction Size in # of Units	-	7.0	6.4	5.1	15.9
	-	(12.9)	(11.2)	(7.8)	(19.9)
Avg. Transaction Size in Capacity	-	1191	1164	812	2909
	-	(2378)	(2039)	(1491)	(3510)

Note: This table includes summary statistics on acquisitions of fossil fuel-generating US units between 2000 and 2023. A description of the sample's construction can be found in Section 3.2. Each column reports the counts and characteristics of the data at varying sample restriction levels. Column (1) reports statistics from all generators in the data. Column (2) reports data from acquired generators. Column (3) restricts the acquisition sample to the first acquisition of each generator. Column (4) reports statistics for first acquisitions of each unit that involve both subsidiary and parent owner changes, while Column (5) focuses on the first parent-only owner change for each generator. The numbers in parentheses represent the standard deviation. The market definition for the cross-market calculations is the power control area. In Columns (2-5), the number of unique plants may differ from the total plant count, as in rare cases, units within the same plant were acquired at different times. Average acquirer and target characteristics report information before acquisitions. All capacity information is reported in MW.

Figure 4: Share of Generation Capacity by Firm Type

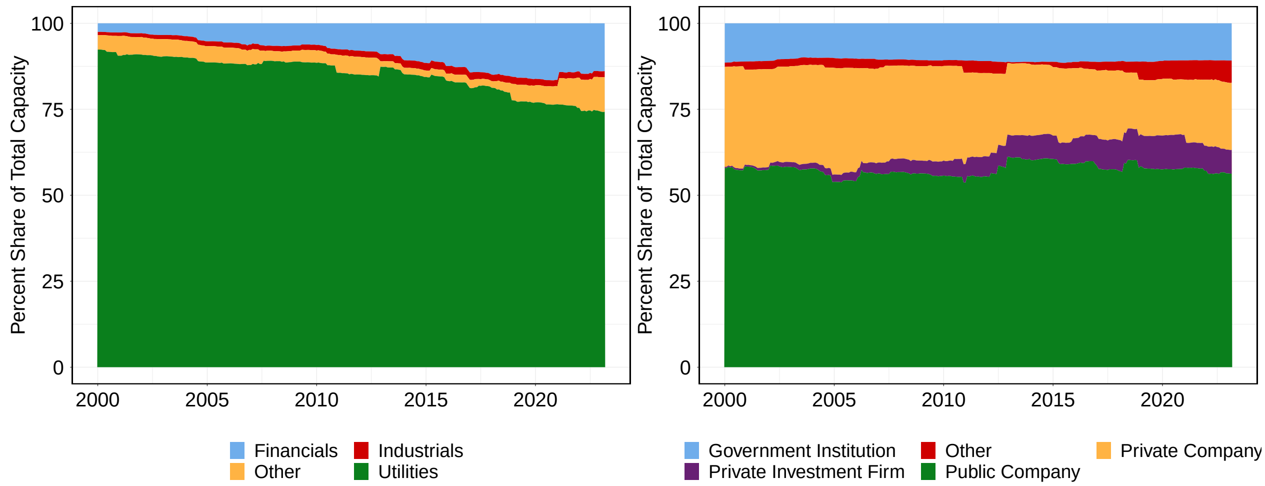

(a) Primary Activity of Owner **(b)** Company Type of Owner

Note: Panel (a) shows the percent share of fossil fuel generation capacity in the US between 2000 and 2023 by the primary industry type of the parent company. The "Utilities" category includes the generation capacity of both regulated and non-regulated power plants owned by energy companies. Panel (b) shows the same statistics by categorizing parent owners into Public Companies, Private Investment Firms, Private Companies, and Government Institutions. These financial firms primarily include private equity firms, pension funds, and bank funds. The classification is taken from GMI.

indicates an increasing presence of financial firms in the industry. The share of total capacity owned by financial firms rose from 2.5% in 2000 to 13.9% in 2023, suggesting a substantial reallocation of assets from utilities to financial firms. Panel (b) highlights that publicly listed firms own slightly more than half of the industry capacity, with their share remaining stable over time. Finally, government institutions—primarily local governments in rural areas, except the Tennessee Valley Authority—own 10.8% of total capacity.[29]

4 Empirical Results

Our empirical strategy aims to identify the effects of acquisitions on power plant productivity and other key operational outcomes. For this purpose, we implement a difference-in-differences research design that compares productivity trends of acquired generators to those that were never or not-yet acquired. The main advantage of our empirical setting is the availability of a high-frequency measure of generator productivity, which enables us

[29] An important institutional detail in electricity markets is variation in market structure and regulation across states. As shown in Table 1, most acquisitions (77%) occur in organized markets where electricity prices are determined through competitive auctions. This trend is also reflected in the geographic distribution of acquisitions in Figure OA-9, where states with nonregulated wholesale generation have significantly higher numbers of acquisitions relative to their size (see also Table OA-2). This highlights the potential role of regulatory institutions and market characteristics in shaping firms' selection into mergers, which should be taken into account when interpreting our results.

to track changes in productivity immediately before and after acquisitions.

We find that acquisitions increase the productivity of power plants by 2.0% on average, but with significant heterogeneity across acquisition types. In particular, ownership changes at both the subsidiary and parent owner levels lead to a 4.9% average efficiency gain, whereas ownership changes only at the parent company level suggest no statistically significant efficiency effect. We conclude this section by examining the heterogeneity of the efficiency effect and studying how acquisitions affect other plant outcomes, such as generation and outages.

4.1 Effects of Acquisitions on Efficiency

We estimate the effects of acquisitions on efficiency using a regression of the following form:

$$y_{it} = \delta_1 \mathbb{1}_{\{\text{Pre-year 1}\}} + \delta_2 \mathbb{1}_{\{\text{Post-year 1}\}} + \delta_3 \mathbb{1}_{\{\text{Post-year 2}\}} + \delta_4 \mathbb{1}_{\{\text{Post-year 3}\}} + X_{it} + \alpha_i + \mu_t + \epsilon_{it}, \qquad (2)$$

where y_{it} is the log efficiency of generator i at week t (measured as inverse heat rate), α_i and μ_t are generator and week fixed effects, respectively.[30] The controls, X_{it}, in our preferred specification include ambient temperature and humidity, a dummy variable for each environmental regulation indicating whether the generator is subject to that regulation, and pollution control device (scrubber) indicators for NO_x, SO_2 and PM.[31] Controlling for factors related to environmental regulations is important, as policy changes over the past two decades may influence firms' acquisition decisions or directly affect plant efficiency due to scrubber installations.

In addition to these variables, we control for monthly time trends by state, installation year, fuel type, capacity bins, and technology type.[32] By incorporating state-specific time trends, we account for changes in electricity demand and supply of non-fossil fuel generation. Furthermore, the time trends for generator characteristics allow for different efficiency trajectories based on generator type. For example, generators might experience a

[30] Even though the underlying data are hourly, we estimate this specification at the weekly frequency to alleviate computational complexity and reduce noise in the hourly data. We later perform a robustness check with daily frequency.

[31] These programs are Clean Air Interstate Rule NO_x Program, Nitrogen Oxides Budget Trading Program, Cross-State NO_x Program, Ozone Transport Commission Program, State Implementation Plan NO_x Program, Regional Greenhouse Gas Initiative, Clean Air Interstate Rule Ozone Season Program, Cross-State Ozone Season Program (Group 1-2), New Hampshire NO_x Program, Mercury and Air Toxics Standards, Clean Air Interstate Rule SO_2 Program, Cross-State Ozone Season Program (Group 1-3), Cross-State SO_2 Program, New Source Performance Standards for Toxics, Texas SO_2 Program.

[32] Capacity bins are categorized as follows: 0-50MW, 50-100MW, 100-250MW, 250-500MW, 500-2000MW; fuel types include gas, coal, and other; and technology types distinguish between combined cycle and other technologies.

decline in efficiency over their lifespans, which can be nonlinear and vary by their vintage. We capture this variation by including installation-year-by-month fixed effects.

The model in Equation (2) includes coefficients of interest, δ_1 to δ_4, to estimate efficiency effects from one-year pre-acquisition to three years post-acquisition.[33] We include δ_1 to examine potential pre-acquisition productivity effects, which could arise due to anticipation effects or disruptions in the production process. The regression coefficients are normalized relative to the period two years before the acquisition, and standard errors are clustered at the acquisition level. We exclude acquired generators from the sample three years after their first acquisition to ensure that their post-treatment periods are not used as controls for other units. This means that we use only the first acquisition of each generator in our baseline empirical model, with a robustness check that includes all acquisitions presented in Section 7.[34]

It is worth noting that the unit of analysis is a generator rather than a plant. Although the same firm usually owns all the generators within a plant, generators often have distinct production profiles, maintenance schedules, and even retirement years (Gowrisankaran et al., 2025). Therefore, we think the generator is the right level of analysis, and it is maintained throughout the paper unless otherwise stated.

Table 2 presents results with various sets of control variables (Columns 1-4) and different acquisition types based on subsidiary and parent owner changes (Columns 5-6). Our preferred specification with the full set of control variables in Column (4) demonstrates efficiency increases following ownership changes. The efficiency of acquired generators increases by 0.6% one year after acquisition and reaches 2% after two years on average. The efficiency increase is robust to including a rich set of controls and time trends, and there is no efficiency change in the year leading up to the acquisition. Overall, these findings suggest that acquisitions lead to some improvements in generator efficiency.

Columns (5-6) of Table 2 test whether the efficiency effect differs by the type of ownership change.[35] Column (5) shows the estimates only for acquisitions with ownership

[33] Specifically, $\mathbb{1}_{\{\text{Pre-year 1}\}}$ is an indicator variable for 1 to 12 months pre-acquisition; $\mathbb{1}_{\{\text{Post-year 1}\}}$ for 0 to 12 months post-acquisition, $\mathbb{1}_{\{\text{Post-year 2}\}}$ for 13 to 24 months post-acquisition, and $\mathbb{1}_{\{\text{Post-year 3}\}}$ for 25 to 36 months post-acquisition.

[34] One potential concern is that never-treated units operating in the same markets as treated units may be affected by acquisition through competitive spillovers, thereby violating the Stable Unit Treatment Value Assumption (SUTVA). To address this concern, we implement a matching approach in Section C.4 that matches each treated unit only with those from different markets. Moreover, Table OA-4 presents results based solely on small acquisitions—those involving less than 10% of market capacity and thus unlikely to generate significant spillovers. These estimates closely align with our main findings, suggesting that spillover effects, on average, are likely not substantial enough to impact our estimates.

[35] When estimating the effects of acquisition on one subsample of acquired units, we exclude the other acquired generators from the regression rather than grouping them with the never-acquired units so that they are not

Table 2: Effects of Acquisitions on Generator Productivity

	All Acquisitions Types (1)	All Acquisitions Types (2)	All Acquisitions Types (3)	All Acquisitions Types (4)	Subsidiary and Parent Changes (5)	Only Parent Changes (6)
	Dependent Variable: Log of Efficiency					
Pre-acquisition (1 Year)	0.002	0	-0.003	-0.003	-0.002	-0.007
	(0.004)	(0.003)	(0.003)	(0.003)	(0.005)	(0.003)
Post-acquisition (1 Year)	0.018	0.015	0.006	0.006	0.015	-0.01
	(0.006)	(0.006)	(0.005)	(0.005)	(0.007)	(0.005)
Post-acquisition (2 Years)	0.035	0.034	0.02	0.02	0.039	-0.003
	(0.009)	(0.008)	(0.007)	(0.007)	(0.01)	(0.007)
Post-acquisition (3 Years)	0.039	0.036	0.02	0.02	0.049	-0.008
	(0.011)	(0.01)	(0.009)	(0.009)	(0.012)	(0.007)
Ambient Temp. & Humidity	X	X	X	X	X	X
Unit & Week FE	X	X	X	X	X	X
State by Month FE		X	X	X	X	X
Unit Characteristic by Month FE			X	X	X	X
Scrubber & Enviro. Prog. FE				X	X	X
R^2	0.707	0.726	0.752	0.753	0.763	0.764
# of Observations	1.838M	1.838M	1.838M	1.838M	1.494M	1.575M
# of Never-Treated Units	2311	2311	2311	2311	2311	2311
# of Treated Units	2046	2046	2046	2046	1089	1142

Note: This table presents the coefficient estimates of δ_1, δ_2, δ_3, and δ_4 from estimating Equation (2). Columns (4-6) present our baseline specification, where we allow for time trends to vary flexibly by unit characteristic and include weather, scrubber, and environmental program controls. Unit characteristic fixed effects include installation year, fuel, technology, and unit capacity bins. The dependent variable is the logarithm of the inverse heat rate. Standard errors are acquisition at the plant level. Table OA-3 presents the same analysis results but for the subsample of acquisitions with both subsidiary and parent company changes.

changes at both the parent and subsidiary levels. By contrast, Column (6) includes ownership changes at only the parent level. The results suggest significant heterogeneity in the efficiency change based on acquisition type. When only the parent owner changes, the estimate is small and not statistically significant, whereas for both subsidiary and parent ownership changes, it indicates an efficiency increase of 4.9%. One might expect the efficiency effects to differ in these two cases because the subsidiary owners typically exert direct control over power plant operations and personnel, whereas the parent owners exercise indirect control through actions such as appointing directors, approving capital expenditures, and setting performance targets (Akey and Appel, 2021). Furthermore, changes at the parent level are more likely to be financial acquisitions, potentially driven by motivations such as diversification and environmental policy considerations rather than efficiency gains (Andonov and Rauh, 2023). Overall, our results highlight that efficiency gains are influenced by the level of ownership change in the corporate structure and whether the direct owner changes.

used as control units.

Figure 5: Dynamic Effects of Acquisitions on Productivity

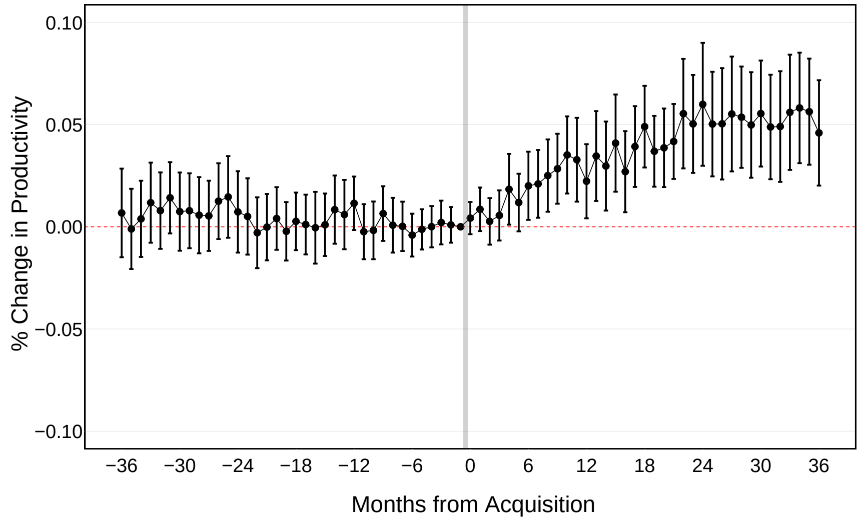

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with 95% confidence intervals. The dependent variable is the logarithm of the inverse weekly heat rate. The unit of observation is generator-week. Standard errors are clustered at the acquisition level.

After demonstrating the impact of acquisitions on generator efficiency, we shift our focus to the dynamic effects. Our goal is to determine the timing of efficiency changes and to test for different pre-treatment trends between the acquired and other generators. To this end, we estimate the change in efficiency around the time of acquisition using the following specification:

$$y_{it} = \sum_{s \in (-36,36)} \hat{\delta}_s D_{i(t'+s)} + X_{it} + \alpha_i + \mu_t + \epsilon_{it}, \tag{3}$$

where $D_{i(t'+s)}$ is a monthly indicator variable equal to 1 for generator i if it is acquired in month t', and zero otherwise. X_{it} includes the same control variables as before. Since we find efficiency effects only in acquisitions where both the subsidiary and parent owner change, we focus exclusively on those acquisitions hereafter.

The estimates of dynamic effects are shown in Figure 5. The pre-acquisition coefficients are small and statistically insignificant, indicating similar productivity trends between acquired and non-acquired generators before acquisition. The difference between these groups remains small until five months post-acquisition, at which point the efficiency of acquired plants begins to diverge. On average, the efficiency of acquired generators increases by 5% eighteen months post-acquisition and then stabilizes. Not observing efficiency gains immediately after acquisitions suggests that the new owner requires time to implement efficiency improvements.

To interpret our results on efficiency gain as causal, we rely on the assumption that an acquisition creates a discontinuous change in power plant behavior, and any unobservable efficiency trends that might lead to selection would be gradual enough to be distinguishable from the more discrete acquisition effect. Our data-rich setting offers key advantages for this assumption to hold, as we observe production at short intervals and incorporate flexible time trends that account for factors likely to influence selection into acquisitions. Additionally, parallel trends holding three years pre-acquisition, coupled with the productivity increase beginning just a few months post-acquisition, further suggest that efficiency gains are not likely caused by unobserved confounding factors.

Still, ownership changes are, of course, not random, and unobservable factors could influence efficiency without acquisitions. If acquirers observe these factors, that might lead to reverse causality, with acquisitions made in anticipation of efficiency gains. Although we cannot eliminate all potential identification threats or account for every unobservable factor, we conduct several robustness checks to ensure our results are robust to various specification choices and identification threats. For example, one possible scenario is that the acquirer observes that the target plant's manager will retire soon and decides to buy the plant, anticipating that the new manager will improve efficiency. To address such a concern, we estimate the effects of manager changes on efficiency in the absence of mergers and find no efficiency increase (Figure OA-10). In addition, we do a battery of robustness checks, including matching estimators, the Callaway and Sant'Anna (2021) estimator, estimation with daily data, estimation with net generation, weighted estimation, and placebo tests with minority acquisitions. We find that the results are robust to these specification choices.[36] See Section 7 for a summary of robustness checks and Appendix F for the corresponding results.

The results so far suggest that the efficiency of power plants improves following ownership changes. Yet, it is important to recognize that efficiency gains in power plants can occur in various ways, not all of which are socially beneficial. For instance, generators might improve their average efficiency by decreasing production and reducing ramping, but this could lead to increased production from a high-cost generator. Alternatively, new owners might operate generators more intensively, increasing their short-term efficiency but potentially causing increased outages and declining long-term performance. In the

[36] We emphasize that our estimates report the average treatment effect on the treated (ATT), specifically the efficiency effects of the proposed acquisitions. In our setting, the ATT, not the average treatment effect (ATE), is the primary and policy-relevant object of interest because we want to learn the effects of actual acquisitions, not hypothetical ones that would occur at random. However, this requires interpreting our results with an understanding of the circumstances under which mergers are proposed and the existing regulatory approach to merger review. As a result, our results do not directly allow for a counterfactual analysis under specific changes in antitrust policy, such as those considered in Bhattacharya et al. (2024).

Table 3: Effects of Acquisitions on Generator Performance Measures

Dep. Var.	Total Generation (1)	Capacity Utilization (2)	Operating Hours (3)	Forced Outages/Derates (4)	Log CO_2 Intensity (5)
Pre-acquisition (1 Year)	-169.221	0.003	-1.556	-0.003	0.006
	(186.918)	(0.004)	(0.882)	(0.012)	(0.006)
Post-acquisition (1 Year)	192.18	0.006	-0.096	-0.026	-0.009
	(267.302)	(0.005)	(1.137)	(0.015)	(0.007)
Post-acquisition (2 Years)	457.981	0.013	0.584	-0.034	-0.037
	(346.979)	(0.005)	(1.427)	(0.019)	(0.01)
Post-acquisition (3 Years)	527.331	0.015	0.985	-0.063	-0.046
	(383.545)	(0.006)	(1.593)	(0.02)	(0.013)
Ambient Temp. & Humidity	X	X	X	X	X
Unit & Week FE	X	X	X	X	X
Unit Characteristic by Month FE	X	X	X	X	X
Scrubber & Enviro. Prog. FE	X	X	X	X	X
Pre-acquisition Mean	7207.429	0.665	45.116	0.188	-0.419
R^2	0.797	0.595	0.695	0.243	0.842
# of Observations	2.612M	1.494M	2.612M	0.705M	1.418M
# of Controls	2311	2311	2311	1383	2026
# of Treated Units	1089	1089	1089	409	977

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3$, and δ_4 from estimating Equation (2). Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. The unit of observation is generator-week. Standard errors are clustered at the acquisition level. The number of observations in Column (4) is lower than the rest because the outage and maintenance data begin in 2013. Some units in Column (5) are missing because their CO_2 emissions always equal zero in the data. The corresponding event study figure for each regression is reported in Figure OA-13. The units of the dependent variables are, respectively: (i) MWh, (ii) a unitless ratio between 0 and 1, (iii) hours, (iv) a unitless ratio between 0 and 1, (v) tons of carbon per MWh.

rest of this section, we provide additional analyses to gain insights into efficiency gains while reserving a more formal investigation of underlying mechanisms for the next section.

We examine the effect of ownership changes on various operational outcome measures, including generation, capacity utilization, operating hours, outages, and the carbon intensity of production. Capacity utilization is defined as the average hourly production as a proportion of capacity over a week, conditional on operation. Operating hours are calculated as the total hours a unit is operational in a given week. For outages, which are available between 2013 and 2021, we calculate the share of hours in a given week a unit experiences a forced outage or derate. Finally, the CO_2 intensity is calculated by dividing CO_2 emissions by generation. Using these outcome measures, we estimate the same specification as in Equation (2).[37]

[37]We estimate the first four outcomes in levels rather than logarithms to account for the presence of zeros. For generations, its distribution shown in Figure OA-6 indicates a thin upper tail, so a small number of

The coefficient estimates in Table 3 indicate improvements in plant performance across multiple dimensions. In Column (1), we find that acquired generators increase their generation by 7.3% compared to the baseline following the acquisition, so the efficiency improvements do not come at the expense of a decline in production.[38] Columns (2-3) suggest that generation increases at both the intensive and extensive margins. We observe that acquired power plants increase capacity utilization by 1.5 pp and operating hours by about 1 hour, though the latter is not statistically significant at the 5% level. The increase in utilization can be viewed as another source of productivity gain, as the generator produces more output conditional on the existing capital and labor stock, as argued in Braguinsky et al. (2015). Moving to Column (4), the results indicate improvements in reliability, with a 6.3 pp reduction in average probability of forced outages and derates, suggesting that efficiency gains are achieved without compromising reliability. Finally, we note a 4.6% decrease in CO_2 intensity, mirroring the results on efficiency gains, as CO_2 emissions are inversely proportional to heat input.[39]

4.2 Discussion of Results

Our findings in this section suggest that acquisitions lead to a 5% average increase in efficiency, but only when both subsidiary and parent owners change. Additionally, acquired generators tend to increase production and utilization, reduce outages, and improve emission intensity. How large is the average 5% efficiency gain? To interpret this finding, it is helpful to compare our estimates to the average within-generator productivity growth in this industry, which is only 0.3% annually.[40] Given this modest within-generator productivity growth, the efficiency gains due to ownership changes are particularly noteworthy.

It is important to note that our analysis focuses on short-run to medium-run efficiency and performance, up to three years. Power plants are long-lived assets, and heat rate efficiencies may come at the cost of increased mechanical stress or wear and tear, which

high-output units is not likely to drive the results. Nevertheless, we estimate the same specification using a Poisson regression in Table OA-8 and find a 4.1% increase in generation and a 2.2% increase in capacity utilization.

[38]This regression also offers indirect evidence that acquirers do not exert market power by withholding the output of the acquired generators. Moreover, this result does not necessarily imply that total market quantity increases. The capacity of acquired generators is typically small compared to the overall market, as illustrated in Figure OA-2. Because short-term electricity demand tends to be inelastic, any unit-level generation increase typically occurs due to changes in the unit's position on the dispatch curve—either through reduced downtime or a decrease in marginal cost—which reallocates generation from other units.

[39]The effects on the other pollutants—SO2 and NOx—are also similar, with both falling by roughly 6%, as shown in Table OA-9 and Figure OA-14. This suggests that efficiency improvements do not come at the expense of worsened environmental performance.

[40]Refer to Figure OA-11, which illustrates the average year-to-year within-generator productivity growth for generators not involved in acquisitions. The productivity growth fluctuates around zero, averaging a 0.3% annual increase over the sample period.

might appear in a horizon longer than three years. These dynamics suggest that heat rate efficiencies may not fully capture the long-term trade-offs related to productivity. Although the reduced precision of the estimates limits our ability to examine very long horizons, we report the 5-year post-acquisition estimates for our main results in the Online Appendix.[41] The results clearly show a persistent improvement in heat rate over a five-year horizon, whereas for other outcomes, the patterns are similar to our baseline results but estimated with less precision.

We also estimate the reduction in CO_2 emissions attributable to acquisitions. As detailed in Appendix B.3, our analysis assumes that efficiency gains begin after each unit's first acquisition and that their production levels remain unchanged post-acquisition. Under these assumptions, we calculate a cumulative decrease of approximately 360 million tons in CO_2 emissions due to acquisitions from 2000 to 2023. This reduction is equivalent to the savings from replacing 800 TWh of gas-fired electricity generation with renewables.

4.3 What Predicts Efficiency Gains: Heterogeneity Analysis

This section explores whether efficiency gains are associated with observable plant or firm characteristics. While these findings do not establish causality, they help derive insights applicable to other industries by documenting transaction characteristics that could predict efficiency gains. For this estimation, we modify Equation (2) by interacting treatment indicators with observable variables Z_{it}:

$$\begin{aligned} y_{it} = {} & \delta_1 \mathbb{1}_{\{\text{Pre-year 1}\}} + \delta_2 \mathbb{1}_{\{\text{Post-year 1}\}} + \delta_3 \mathbb{1}_{\{\text{Post-year 2}\}} + \delta_4 \mathbb{1}_{\{\text{Post-year 3}\}} + \bar{\delta}_1 \mathbb{1}_{\{\text{Pre-year 1}\}} \times Z_{it} + \\ & \bar{\delta}_2 \mathbb{1}_{\{\text{Post-year 1}\}} \times Z_{it} + \bar{\delta}_3 \mathbb{1}_{\{\text{Post-year 2}\}} \times Z_{it} + \bar{\delta}_4 \mathbb{1}_{\{\text{Post-year 3}\}} \times Z_{it} + X_{it} + \alpha_i + \mu_t + \epsilon_{it}. \end{aligned} \tag{4}$$

We estimate this equation separately for a set of generator, firm, or transaction characteristics that might be indicative of efficiency gains. In particular, we consider generator capacity, generator age, whether the acquirer is a serial acquirer, acquirer size, and whether the acquisition is a cross-market acquisition. Details on the construction of these variables are provided in Appendix B.[42]

Results, reported in Table 4, reveal that the efficiency increase is 3.4 pp larger when the generator capacity is higher than the median of acquired generator capacity. This suggests that acquirers might have stronger incentives to improve efficiency in larger plants, where the returns on such improvements are potentially higher. We do not find any significant

[41] See Table OA-11 and Figure OA-16 for heat rate estimates, and Figure OA-17 for the other measures.

[42] We also estimate the heterogeneous effects in a single regression that includes all interaction terms and report the estimates for the 10 most common acquisition types based on these observables in Table OA-10.

Table 4: Heterogeneous Effects of Acquisitions on Productivity

Interaction Var. (Z)	Capacity >Median (1)	Age >Median (2)	Serial Acquirers (3)	Firm Size >Median (4)	Cross-Market Acquisitions (5)
	Dependent Variable: Log of Efficiency				
Post-acquisition (1 Year) × Z	0.023	-0.001	0.014	0.012	0.002
	(0.011)	(0.012)	(0.012)	(0.012)	(0.012)
Post-acquisition (2 Years) × Z	0.035	0.004	0.059	0.049	-0.021
	(0.015)	(0.016)	(0.016)	(0.017)	(0.015)
Post-acquisition (3 Years) × Z	0.034	-0.011	0.058	0.041	-0.039
	(0.018)	(0.02)	(0.02)	(0.02)	(0.019)
Ambient Temp. & Humidity	X	X	X	X	X
Unit & Week FE	X	X	X	X	X
Unit Characteristic by Month FE	X	X	X	X	X
Scrubber & Enviro. Prog. FE	X	X	X	X	X
R^2	0.763	0.763	0.763	0.763	0.763
# of Observations	1.494M	1.494M	1.494M	1.494M	1.494M
# of Units	2311	2311	2311	2311	2311
# of Acquisitions	1089	1089	1089	1089	1089

Note: This table presents the coefficient estimates of $\bar{\delta}_2, \bar{\delta}_3$, and $\bar{\delta}_4$ from estimating Equation (4). Each column reports results from a different regression by varying the interaction variable, Z. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. The unit of observation is generator-week, and the dependent variable is the logarithm of the inverse weekly heat rate. Standard errors are clustered at the acquisition level. Appendix B.4 provides details about the heterogeneity variables. See Table OA-5 for the full set of estimates, including δ_1 through δ_4 and $\bar{\delta}_1$ through $\bar{\delta}_4$.

differential effect with respect to generator age, as shown in Column (2). Next, we turn to firm characteristics: whether the acquirer is a serial acquirer and acquirer size (total owned pre-acquisition fossil fuel generation capacity). The results, reported in Columns (3-4), indicate that efficiency improvements are 5.8 pp higher when the acquirer is a serial acquirer and are 4.1 pp higher when the acquirer firm is larger than the median acquirer. These findings suggest that a firm's experience in plant operation and acquisitions could explain efficiency gains. They also align with and complement the findings of Hortaçsu et al. (2019) that large power firms are more sophisticated in bidding in wholesale electricity market auctions.[43] Finally, in Column (5), we explore whether the efficiency effects differ for cross-market acquisitions. We categorize a generator acquisition as a cross-market acquisition if the acquirer owns no fossil fuel generation capacity in the acquisition market (defined as a power control area). We find that cross-market acquisitions exhibit 3.9 pp lower efficiency gains compared to within-market acquisitions.

[43] Hortaçsu et al. (2019) explore a counterfactual scenario in which large firms acquire smaller ones and improve their bidding operations. The evidence presented in this paper essentially validates the counterfactual hypothesis proposed by Hortaçsu et al. (2019). For further discussion, see Section 6.2.

We note that the heterogeneous effects observed in this analysis could arise from the types of acquisitions that select into each category or from the inherent characteristics of those acquisition types. For example, the lower efficiency effects in cross-market acquisitions may occur because merging parties do not require strong efficiency gains to propose the merger, as regulators are less likely to challenge it given the absence of market power concerns. Alternatively, cross-market mergers might inherently lack the synergies or firm-specific specializations that typically occur within a single market, resulting in lower efficiency effects. Although our results in this section do not identify the exact mechanisms, the following section provides further insights into the specific sources of efficiency gains.[44]

5 Mechanisms

This section proposes mechanisms of efficiency gains, tests them empirically, and quantifies their role using a model of production in power plants. The key finding is that the majority of efficiency gains come from increasing productive efficiency within a generator.

5.1 Mechanisms of Efficiency Improvements

Two mechanisms could explain the estimated efficiency gains: (i) productive efficiency and (ii) dynamic efficiency. We first define these mechanisms and then develop a testable prediction for each one.

Productive Efficiency. Productive efficiency arises when the plant's new owner implements operational processes or invests in new equipment that improves efficiency. This mechanism occurs solely through increasing the generator's efficiency, enabling it to produce more with less fuel for a given production level. Therefore, it is independent of changes in the ramp profile or synergies with other plants in the same market. As illustrated in Figure 6(a), an implication of productive efficiency is a lower heat rate curve, leading to the following testable prediction:

<u>*Prediction 1*</u>: If acquirers improve productive efficiency, the generator's heat rate curve

[44] Another important potential source of heterogeneity is market and state characteristics, as we showed that most acquisitions happen in organized markets and deregulated states. To analyze this, Table OA-6 examines heterogeneity based on three factors: (i) units operating in organized markets (ISO), (ii) units located in states with high acquisition activity (relative to the median), and (iii) units in deregulated states. We find that if a unit is in an ISO or in a deregulated state, the efficiency impact is about 4-5% larger. Additionally, the efficiency impact is 1.8% larger in high acquisition activity states (though not significant at the 5% level). These results stress the importance of regulation and market structure in this industry, as markets and restructuring are associated with increased acquisition activity, which has a higher impact on efficiency.

Figure 6: Illustration of Mechanisms of Efficiency Gains

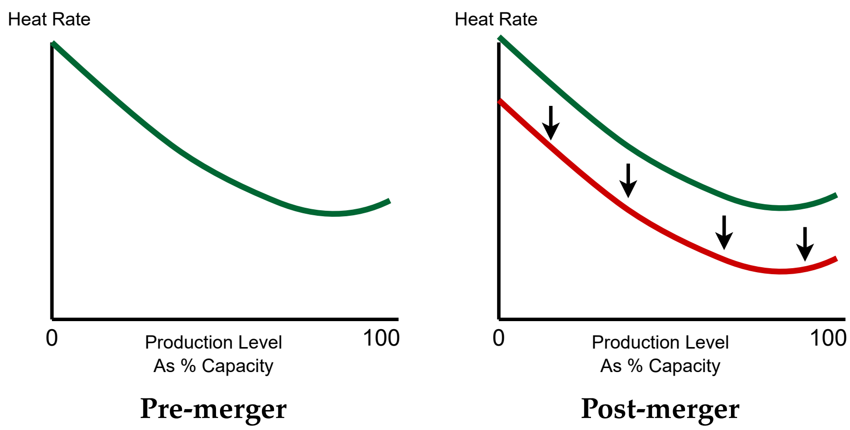

(a) Productive Efficiency

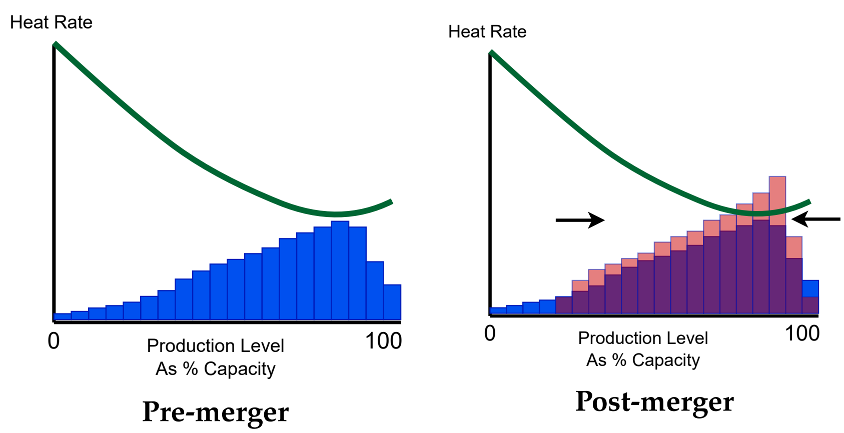

(b) Dynamic Efficiency

Note: This plot illustrates the two mechanisms of efficiency gains studied in Section 5.1. Panel (a) shows efficiency gains through productive efficiency, where the generator's heat-rate curve shifts downward, producing greater efficiency at every capacity level. Panel (b) shows the dynamic-efficiency mechanism, where the heat-rate curve remains unchanged, but the production distribution becomes more concentrated because ramping declines.

shifts down.

Dynamic Efficiency. Dynamic efficiency arises from changes in the generation level over time. As discussed in Section 2.3, a key aspect of power generation is that efficiency is influenced by both the level of production and changes in production. Generators that experience significant production shifts incur ramp costs, which reduce overall efficiency. Power plants must manage these ramp costs due to the stochastic nature of electricity demand, which requires coordination between trading desk personnel responsible for submitting supply bids and plant operators overseeing production. Jha and Leslie (2023) notes that uncertainty in residual demand or mismanagement in production can significantly increase ramp costs.[45] Figure 6(b) illustrates the dynamic efficiency effect, showing a more concentrated production distribution and, therefore, lower ramp costs post-acquisition. A testable hypothesis derived from this mechanism is:

Prediction 2: If acquirers improve dynamic efficiency, the variation of generation goes down.

5.2 Quantifying Productive Efficiency Using Production Functions

We start by testing for productive efficiency using an empirical strategy guided by Prediction 1. In particular, we estimate a production function for generators, where we model heat rate with the following equation:

$$y_{it} = f_{i\tau}(Q_{it}, R_{it}, X_{it}) + \epsilon_{it}, \tag{5}$$

where $y_{it} = \log(\text{Fuel}_{it}/Q_{it})$ is log heat rate, Q_{it} is production of generator i at hour t and, R_{it} is the ramp rate defined as the hourly change in production, $(Q_{it} - Q_{it-1})/C_i$, where C_i denotes generator capacity. The other variables, X_{it}, include ambient temperature and ambient humidity. Subscript i denotes the generator, t denotes the hour, and τ indicates the pre- or post-acquisition period.

As described in Bushnell and Wolfram (2005), this form of heat rate modeling can be micro-founded from a Leontief electricity production function. To see this, assume electricity is produced according to the following production function:

$$Q_{it} = \min(g_i(F_{it}, R_{it}, X_{it})\epsilon_{it}, h_i(K_{it}, L_{it})\omega_{it}), \tag{6}$$

[45] One source of dynamic efficiency could be portfolio effects, where a firm operating multiple power plants in the same market could achieve portfolio-level efficiencies through ramp synchronization and efficient production allocation. Separating this mechanism from dynamic efficiency is challenging; therefore, we consider it part of dynamic efficiency effects.

where F_{it}, K_{it}, L_{it} are fuel, capital and labor inputs, ϵ_{it} is unobserved, time-varying fuel efficiency, X_{it} are observable factors affecting fuel efficiency, and ω_{it} is total factor productivity. This Leontief production function, under a cost minimization assumption, implies that $Q_{it} = g_i(F_{it}, R_{it}, X_{it})\epsilon_{it}$. Assuming $g_i(\cdot)$ is strictly monotone in F_{it}, it can be inverted to write $F_{it} = g_i^{-1}(Q_{it}, R_{it}, X_{it})\epsilon_{it}$. Dividing both sides by Q_{it} and taking the logarithm yields the functional form in Equation (5).

Importantly, the production function in Equation (5) is indexed by i and τ, where τ equals 1 in the post-acquisition periods and 0 in the pre-acquisition periods. Therefore, we estimate a generator-specific production function separately for the pre- and post-acquisition periods, with f_{i0} representing the production technology of generator i before the acquisition and f_{i1} representing it afterward.

It is worth highlighting the benefits of estimating generator-specific production functions. The form in Equation (5) accommodates heterogeneity in production technology across generators through the generator-specific and time-varying production function $f_{i\tau}$. Since $f_{i\tau}$ captures productivity differences across generators and over time, the production function literature generally interprets ϵ_{it} as an ex-post shock (or measurement error) to output that is orthogonal to inputs. Thus, our model is likely to be robust to transmission bias, which creates a correlation between productivity level and inputs (Marschak and Andrews, 1944; Ackerberg et al., 2015). Furthermore, through a time-varying production function, we model the effects of acquisitions not only on the productivity level but also on the production technology.

We can estimate a flexible production model due to the availability of hourly data, as it provides a large number of observations for each generator, even within a limited time frame around acquisitions. This highlights the advantages of a data-rich environment, contrasting with the production function literature, which often imposes an industry-level functional form due to data limitations (De Loecker and Syverson, 2021).

We use a nonparametric local polynomial regression to estimate the functions f_{i0} and f_{i1} for each acquired generator as detailed in Appendix B.2. To estimate f_{i1}, we use three years of post-acquisition data, while f_{i0} is estimated using data from three years prior to the acquisition. We then measure the changes in productivity by calculating the difference between the post-acquisition and pre-acquisition heat rate curves for each generator and then averaging these differences. Specifically, we calculate:

$$\Delta C(Q) = c_{post}(Q) - c_{pre}(Q) = \frac{1}{N_{acq}} \sum_{i=1}^{N_{acq}} \left(f_{i1}(Q, \bar{R}_i, \bar{X}_i) - f_{i0}(Q, \bar{R}_i, \bar{X}_i) \right),$$

Figure 7: Estimates of Average Heat Rate Curves

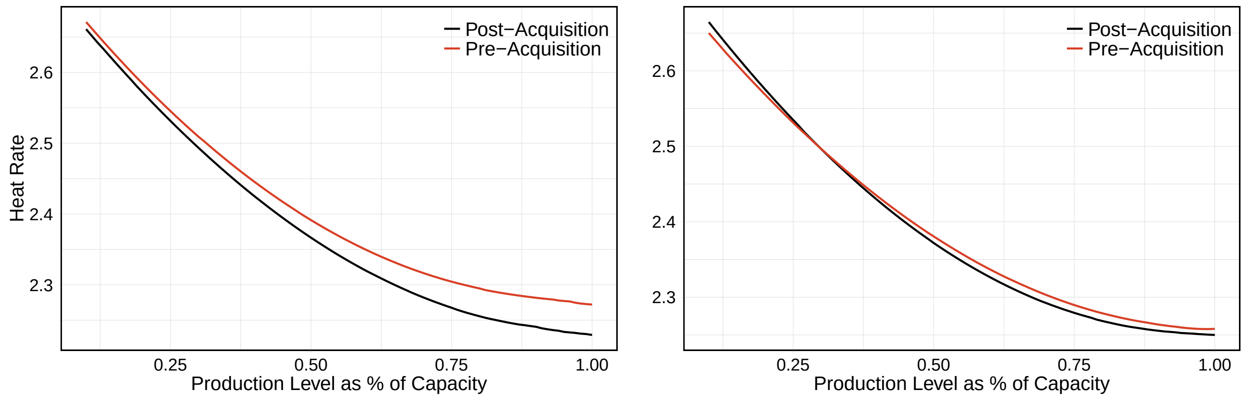

(a) Acquired Generators **(b)** Control Generators

Note: This figure shows estimates of average heat rate curves three years before acquisition and three years after acquisition. Panel (a) shows this for the acquired generators group, and Panel (b) shows this for the control group constructed by a matching procedure detailed in Section C.4. The treated group sample is the same as Column (5) of Table 2. Figure OA-12 reports the confidence band for the difference between the two heat rate curves obtained from a bootstrap procedure. Further details of the estimation procedure are provided in Section B.2.

where N_{acq} represents the number of acquired generators and $Q \in [10, 100]$ is the production level as a percentage of capacity.[46] The terms $c_{pre}(Q)$ and $c_{post}(Q)$ denote the average heat rate at production level Q before and after acquisition, respectively. The control variables are set to $\bar{R}_i$ and $\bar{X}_i$, which is 0 for ramp rate, and the pre-acquisition medians for temperature and humidity to isolate the effects of post-acquisition changes in these variables. Thus, $\Delta C(Q)$, known as the average structural function (Blundell and Powell, 2003), represents the change in the average heat rate at each production level controlling for ramp and weather conditions.

We also construct a control group by matching each acquired generator to those never acquired in a different market based on capacity, age, fuel, and technology type, as detailed in Appendix B.2. We then apply the same estimation procedure to these control generators to quantify changes in the heat rate curves without acquisitions.

Figure 7(a) reports $c_{post}(Q)$ and $c_{pre}(Q)$ for the acquired generators, while Figure 7(b) displays these curves for the control group. Comparing pre- and post-acquisition heat rate curves reveals a downward shift in the heat rate curve for acquired generators at every production level, with larger effects near the generator's capacity. In contrast, the control group's heat rate curve remains stable.[47] We also calculate a confidence band for the difference between the pre- and post-acquisition heat rate curves of acquired generators,

[46]The utilization values start at 10% because production at lower capacity levels is rare and tends to yield noisy estimates.

[47]The slight shift in the heat rate curve of control generators is consistent with the within-generator aggregate efficiency growth documented in Figure OA-11.

as presented in Figure OA-12, confirming that the difference is statistically significant. These results provide direct evidence that the acquirers increase the productive efficiency of the acquired generators by improving their heat rates.

Having estimated the heat rate curves, we can now quantify the total efficiency gain from the downward shift in the heat rate curve. To do this, we integrate the difference between the post- and pre-acquisition curves as follows:

$$\Delta = \frac{1}{N_{acq}} \sum_{i=1}^{N_{acq}} \int \big(f_{i1}(Q, \bar{R}_i, \bar{X}_i) - f_{i0}(Q, \bar{R}_i, \bar{X}_i)\big) dF_{i0}(Q),$$

where $F_{i0}(Q)$ represents the pre-acquisition production distribution of generator i. This calculation maintains the production distribution from the pre-acquisition period and quantifies efficiency gains solely from changes in the heat rate curve. The result indicates a 3.9% (CI: 2.9%, 4.8%) increase in efficiency, accounting for approximately three-quarters of the total efficiency gain observed in the event study. Therefore, most of the efficiency gain stems from increased productive efficiency attributable to the acquirers' improvements to the generator's internal operations.

5.3 Quantifying Dynamic Efficiency Mechanism

We next assess the role of dynamic efficiency. Prediction 2 posits that increased dynamic efficiency results in reduced production variability post-acquisition. To test this, we consider three measures of production variability: the CoV of heat rate, the CoV of utilization, and the number of ramps.[48] These metrics collectively provide insights into how acquisitions influence the production dynamics of a generator.

We estimate our baseline regression using these measures as outcome variables and report the estimates in Table 5. Post-acquisition, we observe significant reductions in all measures of production variability. Specifically, the CoV of heat rate decreases by an average of 0.029 from a pre-acquisition mean of 0.235, and the CoV of utilization drops by 0.029 from a pre-acquisition mean of 0.364. We also find a significant decline in the number of ramps, showing a 12% decrease from the pre-acquisition level.

We can also quantify the contribution of the dynamic efficiency effect using the pro-

[48] We define a ramp event as a change in production where the output increases from below 20% to above 80% of the plant's capacity or decreases from above 80% to below 20% within a period of less than three days.

Table 5: Regression Results on Dynamic Efficiency Mechanism

Dep. Var.	CoV of Heat Rate (1)	CoV of Utilization (2)	Number of Ramps (3)
Pre-acquisition (1 Year)	-0.001	0	-0.059
	(0.005)	(0.004)	(0.1)
Post-acquisition (1 Year)	-0.016	-0.015	-0.265
	(0.007)	(0.005)	(0.138)
Post-acquisition (2 Years)	-0.026	-0.026	-0.369
	(0.008)	(0.008)	(0.159)
Post-acquisition (3 Years)	-0.029	-0.029	-0.435
	(0.009)	(0.008)	(0.167)
Ambient Temp. & Humidity	X	X	X
Unit & Week FE	X	X	X
Unit Characteristic by Month FE	X	X	X
Scrubber & Enviro. Prog. FE	X	X	X
Pre-acquisition Mean	0.235	0.364	3.499
R^2	0.195	0.528	0.452
# of Observations	1.476M	1.476M	1.476M
# of Never-Treated Units	2309	2309	2309
# of Treated Units	1089	1089	1089

Note: This table presents coefficient estimates of $\delta_1, \delta_2, \delta_3$, and δ_4 in Equation (2) from a regression of the CoV of heat rate, CoV of utilization, and number of ramps on treatment dummies. The CoVs are calculated from hourly data every week; thus, the regressions use weekly data. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. The number of observations in Columns (1-2) is smaller because CoV cannot be calculated for some weeks due to a small sample size. Figure OA-15 reports each regression's corresponding event study figure. Standard errors are clustered at the acquisition level.

duction model developed in Section 5.2 as follows:

$$\frac{1}{N_{acq}} \sum_{i=1}^{N_{acq}} \left(\int f_{i1}(Q, R, \bar{X}_i) dF_{i1}(Q, R) - \int f_{i1}(Q, R, \bar{X}_i) dF_{i0}(Q, R) \right),$$

where $F_{i1}(Q, R)$ and $F_{i0}(Q, R)$ denote the distributions of production and ramp rate post- and pre-acquisition. This calculation essentially poses the following counterfactual question: What would be the efficiency difference if the generator had the post-acquisition heat rate curve (f_{i1}) in both the pre- and post-acquisition periods while only changing the production profile from F_{i0} to F_{i1}? In other words, it controls for production technology and captures efficiency effects only due to changes in the production distribution. We caveat that this calculation reflects both changes in production variability and the effect of shifting output toward the plant's efficient scale, even though the latter is not strictly a dynamic efficiency. This calculation yields an efficiency gain of 1.7% (CI: 0.0%, 4.2%), corresponding to around 30% of the total effect.[49]

It is important to recognize that improvements in dynamic efficiency can arise from various factors. One potential factor is increased productive efficiency: a marginal generator in the dispatch curve that becomes more efficient after acquisition may operate infra-marginally more often, leading to reduced ramping. Another explanation could be decreased outages and forced maintenance, which would reduce ramping between inactive and operational modes. Furthermore, the acquirer may change the power plant's operations or improve coordination between the bidding desk and plant operators. Although our analysis does not separate the impact of these individual sources, it highlights the importance of ramping costs in improving power plant efficiency.

The analysis in this section focused only on marginal cost gains from fuel efficiency. Acquisitions may also reduce fixed costs or result in non-fuel cost savings. For example, decreased ramping can reduce wear and tear, thus lowering maintenance expenses and prolonging the lifespan of capital. Additionally, acquisitions may generate economies of scale in maintenance and bidding (Haldi and Whitcomb, 1967; Hortaçsu and Puller, 2008). Although these fixed cost efficiencies could be large, they are generally not considered in merger analysis (Röller et al., 2006) and cannot be accurately measured with our data. Therefore, they fall outside the scope of this paper.

6 How Do Acquirers Improve Productive Efficiency?

So far, our analysis has demonstrated efficiency improvements following ownership changes, mainly due to increased productive efficiency. This result raises a natural follow-up question: How do acquirers achieve these efficiency gains? We will now address this question.

6.1 Productive Efficiency: Operational Improvements or Investment?

In Section 2.3, we proposed two potential mechanisms to improve a power plant's productive efficiency. The first mechanism involves implementing low-cost operational improvements, such as personnel training, efficient production management, best practices, and improvements in repairs and maintenance. Such improvements would indicate a knowledge transfer from the acquirer to the acquired plant. The second mechanism entails high-cost capital investments by acquirers to upgrade existing equipment, suggesting that the previous owner faced credit constraints or lacked the incentives to make efficiency-improving capital investments.

Disentangling these two sources is useful not only for understanding the nature of

[49]The productive and dynamic efficiency effects do not sum to exactly 100% due to noise in estimation.

Figure 8: Effects of Acquisitions on Manager Change

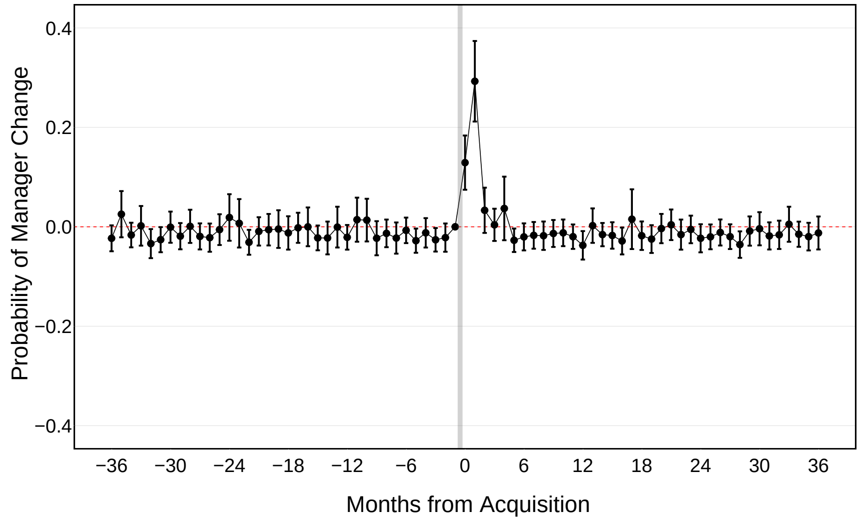

Note: This figure shows coefficient estimates of a manager change dummy on $\hat{\delta}_s$ where $s \in (-36, 36)$. The unit of observation is generator-week. Error bars show 95% confidence intervals. Standard errors are clustered at the plant level.

efficiency gains but also for informing antitrust policy. For efficiencies to be recognized in merger evaluations, they must be merger-specific.[50] Efficiencies from relaxing capital constraints would not be merger-specific, as they could also be attained through alternative means, such as raising new capital or minority investment. In contrast, knowledge transfers can be considered merger-specific because they involve exchanging organizational knowledge and intangible capital between the merging entities, a process that is unlikely to occur outside of a merger (Atalay et al., 2014).

We aim to disentangle the sources of productive efficiency improvements using additional data on manager changes, capital investments, non-fuel inputs, and maintenance. Specifically, we investigate whether power plants undergo personnel changes and increase capital expenditures post-acquisition. The former would suggest operational changes, whereas the latter would provide evidence for the role of capital investment. Moreover, by analyzing non-fuel inputs and maintenance, we evaluate the possibility of substituting fuel with other inputs to achieve efficiency gains.

We use the dynamic difference-in-differences specification in Equation (3) to explore whether acquired plants experience more managerial changes than non-acquired plants. The dependent variable is set to 1 if the power plant manager is replaced in a given month and 0 otherwise. Results shown in Figure 8 reveal a significant increase in managerial changes post-acquisition: acquired plants are 15 pp more likely to experience a change

[50] The 2023 Horizontal Merger Guidelines state, "the merger will produce substantial competitive benefits that could not be achieved without the merger under review". (DOJ and FTC, 2023)

within one month and 45 pp within two months, relative to the non-acquired plants. Using LinkedIn data, we also analyze the qualifications of new managers and find that new managers are 5.58 pp (s.e.=2.51) more likely to hold a master's degree and 4.26 pp (s.e.=1.79) more likely to have a bachelor's degree than those involved in changes without mergers.

The results on manager changes raise an important question: Can the efficiency gains be solely attributed to manager changes? To explore this, we estimate the efficiency effects separately for acquisitions with and without manager changes (reported in Table OA-7) and for manager changes without acquisitions (reported in Figure OA-10). The findings indicate that manager changes without acquisitions have no significant efficiency effects, whereas acquisitions without manager changes still lead to large efficiency improvements. Although these results are not conclusive on their own, they suggest that manager changes alone are not sufficient to generate efficiency gains, and firms cannot achieve efficiencies simply by replacing their manager. Our interpretation is that manager turnover indicates significant operational changes, and it must be accommodated by organizational changes to generate efficiency gains. These insights echo the common findings in the literature on the role of management practices and organization in explaining productivity differences (Bloom and Van Reenen, 2010; Macchiavello and Morjaria, 2022).[51]

Next, we examine the changes in capital expenditures and non-fuel inputs after acquisitions, acknowledging that this analysis relies on a different and more limited dataset. Specifically, data on capital expenditures, number of employees, and non-fuel intermediate input costs are available only for a subset of plants reporting to FERC, and they are annual, unlike the hourly heat rate data. Therefore, while these findings provide useful insights, they warrant cautious interpretation given these data limitations.

The coefficient estimates in Table 6 suggest that acquired plants do not increase capital expenditures. The coefficient estimate for capital expenditures is -24%, but it is imprecise due to the small sample size. Nevertheless, it is still possible to reject the hypothesis that capital expenditure increases by more than 5% at the 10% significance level.[52] The estimates for non-fuel materials costs and labor in Columns (2-3) are also noisy, but they

[51] To further understand the role of managerial changes in efficiency gains, Table OA-12 reports the additional increase in the probability of a manager change across different types of acquisitions, using the observable characteristics from the heterogeneity analysis in Section 4.3. This analysis suggests that the probability of a managerial change tends to be positively correlated with the likelihood of observing higher efficiency gains, providing suggestive evidence for the interpretation that managerial changes signal operational changes in the plant.

[52] Further evidence against the capital expenditure hypothesis comes from the timing of efficiency gains and operating hours. Significant capital investments typically require more than five months to implement and usually involve considerable downtime, neither of which we observe.

Table 6: Effects of Acquisitions on Non-fuel Costs and Maintenance

Dep. Var.	Log Capital Expenditures (1)	Log Non-fuel Costs (2)	Log Number of Employees (3)	Maintenance Probability (4)
Pre-acquisition (1 Year)	-0.214	-0.335	-0.22	-0.013
	(0.161)	(0.4)	(0.111)	(0.009)
Post-acquisition (1 Year)	-0.052	-0.112	-0.326	-0.024
	(0.163)	(0.211)	(0.124)	(0.012)
Post-acquisition (2 Years)	-0.236	0.095	-0.06	-0.038
	(0.17)	(0.267)	(0.136)	(0.014)
Post-acquisition (3 Years)	-0.236	-0.304	-0.003	-0.048
	(0.176)	(0.297)	(0.156)	(0.014)
Ambient Temp. & Humidity	X	X	X	X
Unit & Week FE	X	X	X	X
Unit Characteristic by Month FE	X	X	X	X
Scrubber & Enviro. Prog. FE	X	X	X	X
Pre-acquisition Mean	-	-	-	0.096
R^2	0.896	0.704	0.946	0.15
# of Observations	0.018M	0.018M	0.017M	0.705M
# of Controls	1472	1643	1553	1383
# of Treated	176	203	148	409

Note: This table presents the coefficient estimates from estimating the effects of acquisitions on capital expenditures, non-fuel intermediate input costs, number of employees (all observed at the annual frequency), and maintenance (observed at the weekly frequency). Standard errors are clustered at the acquisition level. Note that the capital expenditure information is available only for major electric utilities as defined by the FERC.

similarly provide evidence against large increases. These findings suggest that efficiency increases do not come from capital expenditures; instead, operational improvements are the key drivers of increases in productive efficiency.

In our final analysis, we examine how maintenance changes after acquisition, as it could also be viewed as an input in electricity generation. Moreover, maintenance is important in its own right to understand plant performance because decreased forced maintenance might indicate better equipment management by new owners, which would increase production, as the generator would go offline less often for maintenance. We analyze the probability that a generator undergoes maintenance in a given week. The results in Column (4) suggest that maintenance probability decreases after acquisitions, indicating that more maintenance duration is not the primary means of improving efficiency.

6.2 Who Acquires Whom: Productivity of Acquirer and Target Firms

This section estimates the productivity levels of acquirer and target firms to determine whether (i) acquirers are more productive than target firms and (ii) acquirers have a comparative advantage in utilizing acquired assets. This analysis not only provides evidence on the mechanisms of efficiency improvements but also offers insights into the broader economic implications of ownership changes. Acquisitions, as a key mechanism of resource reallocation among firms, can lead to allocative efficiency gains in the economy by transferring assets from less productive to more productive firms or enabling better utilization of these assets.

We modify our baseline specification in Equation (2) by including three sets of indicator variables to estimate the efficiency levels of three distinct asset types: (i) acquired generators, (ii) the acquirer's existing generators not involved in the transaction, and (iii) the target's existing generators not involved in the transaction. Formally, we estimate the following specification:

$$y_{it} = \sum_{j=1}^{3} \theta_{1j} \mathbb{1}_{\{\text{Pre-year, 1-3}\}j} + \theta_{2j} \mathbb{1}_{\{\text{Post-year, 1-3}\}j} + X_{it} + \mu_t + \epsilon_{it}, \tag{7}$$

where j represents the asset types listed above and y_{it} is log productivity.[53] This specification estimates the efficiency of the target's assets, acquirer's assets, and acquired assets around the time of acquisition. Note that this regression does not include generator fixed effects, as we aim to estimate level differences in productivity rather than changes. However, we account for generator characteristics by controlling for generator age, capacity, technology, fuel type, and scrubbers. We restrict the sample to transactions where both the acquirer and target own generators not involved in the transaction.[54] We normalize the efficiency of the acquirer's generators to zero in the pre-acquisition period.

Figure 9 presents the estimated coefficients for three groups: the acquirer's existing assets (in red), the target's existing assets (in blue), and the acquired assets (in black). First, we observe that acquisitions do not significantly impact the productivity of existing assets, with the productivity levels of both acquirer and target remaining stable around the time of acquisition. Notably, however, acquirers have a productivity level 1.7% higher than target firms. As a result, acquisitions reallocate assets from less to more productive firms, although the difference in productivity is relatively modest.

[53] For the acquired plants, we exclude the first year after acquisition to estimate the long-term effects of acquisitions.

[54] This subset accounts for 67% of all acquisitions.

Figure 9: Efficiency of Acquirer and Target Firms

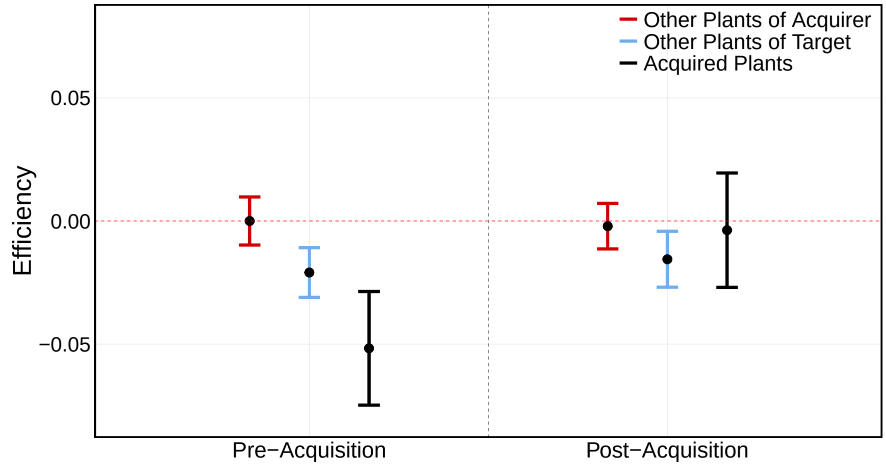

Note: Regression estimates from Equation (7). Red, blue, and black bars show the change in the acquirer's existing assets, the target's assets, and the acquired assets, respectively. Error bars indicate 95% confidence intervals. The efficiency of the acquirer's assets in the pre-acquisition periods is normalized to zero. Standard errors are clustered at the plant level.

Next, we compare the productivity of acquired plants with the acquirer's and target's existing plants. A key observation is that target firms tend to sell their underperforming assets: the sold plants are 3% less productive than other plants in the target's portfolio. What happens to these underperforming plants after acquisition? The efficiency of these plants improves by 5%, reaching the same efficiency level as the acquirer's other plants.

The findings in this section indicate that *high-productivity* firms buy underperforming assets of *low-productivity* firms and make the acquired asset *as productive as* its existing assets after the acquisition. This pattern corroborates our earlier conclusion that efficiency improvements come primarily from operational improvements through knowledge transfers. Furthermore, these results also provide empirical evidence about the theories of merger gains in the literature. One common theory, the Q theory of mergers (Jovanovic and Rousseau, 2002), posits that there are inherent productivity differences between firms, and acquisitions transfer assets from low- to high-productivity firms. This implies a "high-buys-low" pattern. According to another theory proposed by Rhodes-Kropf and Robinson (2008), assets and firms could be complementary, with firms having varying degrees of capability in operating different assets. This implies a "like-buys-like" pattern. Our results lend support to both theories of mergers by demonstrating that assets are allocated to firms with relative and absolute advantages in utilizing them.

This analysis also serves as an important input for merger analysis, particularly in determining post-merger marginal costs for firms with different efficiencies (Farrell and Shapiro, 1990). A key question in this context is the transferability of efficiency between firms, as prior research has highlighted that organizational challenges in integrating firms can hinder the transfer of productivity-improving practices (Weber and Camerer, 2003; Malmendier et al., 2018). Our empirical analysis contributes to this question by providing evidence that efficiency could be transferable in the context of power plant acquisitions.

A natural question arising from this section's findings and the paper's overall conclusions is why previous owners do not implement the operational improvements. Given that our study is an industry-level analysis rather than a firm-level case study, we cannot provide a definitive answer to this question. Nonetheless, it is important to note that our results align with substantial evidence of persistent firm-level productivity differences in various industries (Syverson, 2011; Gibbons and Henderson, 2012). We interpret our evidence to suggest that some firms develop intangible capital over time for more efficient power plant operation, and this within-organization knowledge is transferable through ownership changes. This can occur in many forms, for example, by transferring asset-specific expertise (Hortaçsu and Syverson, 2007; Atalay et al., 2014), operational strategies (Eliason et al., 2020), or managerial practices (Bloom et al., 2012). Therefore, acquisitions provide a channel for spreading intangible capital across firms, which is less likely to be achieved through other means.[55]

7 Robustness Checks

In this section, we explore the robustness of our findings by considering alternative specifications. Detailed descriptions of this analysis and the corresponding results are provided in Appendix C and Appendix F.

Estimation Frequency: Our main analysis uses weekly data to estimate the effects of acquisitions, as this aggregation reduces noise in the hourly data and is computationally convenient. To assess the robustness of our findings, we conduct the same estimation using daily frequency. The results, reported in Figure OA-18, remain consistent at the daily frequency, although there is a slight increase in standard errors.

Acquisition Sample: In our baseline specification, we focus only on each generator's first acquisition to avoid using data from post-acquisition periods. As a robustness check, we extend our analysis to include all acquisitions of generators during the sample period. The findings, reported in Column (4) of Table OA-13, Figure OA-20, and Table OA-16, suggest a slightly smaller effect than our baseline result, indicating that the efficiency gains may be lower with subsequent acquisitions.

[55]This source of efficiency gains differs from the technology-related synergies studied in the merger literature. Some examples include economies of density in the ride-hailing industry (Rosaia, 2020), congestion-related efficiency in the telecommunications industry (Elliott et al., 2025), and reduction in shipping distance in the beer industry (Miller and Weinberg, 2017).

Weighting by Capacity: Our main specification estimates the average effects without accounting for the varying capacity sizes of acquired generators. In a robustness check, we weigh observations by capacity, which provides a more accurate measure of total cost savings. The results from this specification suggest similar efficiency effects, indicating that the evidence does not primarily come from small units (Column (2) of Table OA-13, Table OA-15, and Figure OA-23).

Estimation with Net Generation: While our primary analysis uses gross generation due to its high-frequency availability, we also conduct a robustness check using net generation data from EIA. The results, reported in Figure OA-22, Column (3) of Table OA-13, and Table OA-17, are broadly similar to our main findings, though the effect is slightly lower.

Estimation after 2010: A potential concern in our analysis is the impact of deregulation, which overlaps with our sample period for a few years in the early 2000s. Although we exclude ownership changes corresponding to divestitures, we conduct a robustness check by restricting our analysis to acquisitions after 2010. The results are reported in Column (3) of Table OA-13 and Figure OA-19.

Matching Difference-in-Differences: For the matching specification, we use the sample of first acquisitions that experience both subsidiary and parent change. We match each acquired generator with five never-acquired comparable units. For each unit, we first create a pool of potential control units that share the same fuel type and technology but operate in different markets (ISO) to prevent spillover effects. We then match these generators based on capacity and age using a least-squares distance metric, with weights inversely proportional to each variable's standard deviation. Results are presented in Column (5) of Table OA-13 and Figure OA-22.

Staggered Difference-in-Differences: Recent developments in econometrics suggest that the two-way fixed effects difference-in-differences approach might produce a weighted average of all potential combinations of pairwise difference-in-differences estimators, where the control unit in the pair could be a unit that is treated at a different time (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021). To address this issue, we estimate cohort-specific treatment effects using the Callaway and Sant'Anna (2021) method. The results, reported in Figure OA-21, are similar to our baseline results.

Placebo Tests: We use minority acquisitions as a placebo test against potential unobservable characteristics driving both acquisitions and efficiency changes. If such unobservables exist, they would likely influence minority acquisitions as well. The results, reported in Column (6) of Table OA-13 show no change in power plant efficiency following minority acquisitions.

8 Concluding Remarks

By reallocating resources between firms, acquisitions affect a significant portion of the economy. Despite their importance, there is limited systematic evidence of their effects on productivity. This study provides detailed empirical analyses of the efficiency effects of ownership changes by examining a large sample of power plant acquisitions between 2000 and 2023 in the US.

Our empirical results can be summarized into three principal findings. First, acquired plants experience, on average, a 2% increase in fuel efficiency within five to eighteen months after acquisitions. This effect is more pronounced, rising to 5%, for acquisitions involving changes at both the subsidiary and parent owner levels. Second, acquired generators tend to demonstrate improved operational performance: they produce more, increase their capacity utilization, and decrease their outage frequency and emission intensity. Finally, our evidence suggests that the new owners improve productivity by changing operational processes rather than by making capital investments.

Our findings draw on a large number of acquisitions in the power generation industry and high-frequency data on physical productivity. Using physical measurements in this homogeneous product setting allows us to disentangle the productivity effects from other potential merger effects, such as changes in market power, buyer power, or product quality. With high-frequency data, we can treat mergers as discrete events and compare generator productivity immediately before and after acquisitions. Finally, by aggregating evidence from numerous acquisitions, we have the statistical power to uncover the mechanisms that generate efficiency gains.

The results of this paper have important policy implications, as they provide direct input for evaluating the trade-off between market power and efficiency resulting from mergers. Our results present a mixed view of whether mergers generate efficiencies. On the one hand, we document that mergers in the electricity generation sector can generate efficiencies that are large and through a mechanism that could be considered merger-specific. On the other hand, not all mergers necessarily generate efficiencies; we find no detectable average efficiency increase from parent ownership changes, which tend to be larger and are less likely to influence plants' operations. In conclusion, the main message of our paper is that while efficiency effects in mergers should not be ruled out, they necessitate careful analysis tailored to the circumstances of each merger.

References

Abito, J. M., J. S. Han, J.-F. Houde, and A. A. van Benthem (2024). Agency Frictions and Procurement: New Evidence from US Electricity Restructuring. *The Journal of Industrial Economics 72*(1), 20–48.

Abito, J. M., C. R. Knittel, K. Metaxoglou, and A. Trindade (2022). The Role of Output Reallocation and Investment in Coordinating Environmental Markets. *International Journal of Industrial Organization 83*, 102843.

Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification Properties of Recent Production Function Estimators. *Econometrica 83*(6), 2411–2451.

Akey, P. and I. Appel (2021). The Limits of Limited Liability: Evidence from Industrial Pollution. *The Journal of Finance 76*(1), 5–55.

An, Y. and W. Zhao (2019). Dynamic Efficiencies of the 1997 Boeing-McDonnell Douglas Merger. *The RAND Journal of Economics 50*(3), 666–694.

Andonov, A. and J. D. Rauh (2023). The Shifting Finance of Electricity Generation. *Stanford University Graduate School of Business Research Paper No. 4287123*, Available at SSRN 4287123.

Arora, A., S. Belenzon, and L. Sheer (2021). Matching Patents to Compustat Firms, 1980–2015: Dynamic Reassignment, Name Changes, and Ownership Structures. *Research Policy 50*(5), 104217.

Ashenfelter, O. and D. Hosken (2010). The Effect of Mergers on Consumer Prices: Evidence from Five Mergers on the Enforcement Margin. *The Journal of Law and Economics 53*(3), 417–466.

Ashenfelter, O. C., D. S. Hosken, and M. C. Weinberg (2015). Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry. *The RAND Journal of Economics 46*(2), 328–361.

Asker, J. and V. Nocke (2021). Collusion, Mergers, and Related Antitrust Issues. In *Handbook of Industrial Organization*, Volume 5, pp. 177–279. Elsevier.

Atalay, E. (2014). Materials Prices and Productivity. *Journal of the European Economic Association 12*(3), 575–611.

Atalay, E., A. Hortaçsu, and C. Syverson (2014). Vertical Integration and Input Flows. *American Economic Review 104*(4), 1120–48.

Atalay, E., A. Sorensen, C. Sullivan, and W. Zhu (2024). Product Repositioning by Merging Firms. *The Journal of Industrial Economics 72*(2), 868–908.

Berger, D. W., T. Hasenzagl, K. F. Herkenhoff, S. Mongey, and E. A. Posner (2023). Merger Guidelines for the Labor Market. *National Bureau of Economic Research, No. 31147*.

Bhattacharya, V., G. Illanes, and D. Stillerman (2024). Merger Effects and Antitrust Enforcement: Evidence from US Consumer Packaged Goods. *National Bureau of Economic Research Working Paper No. 31123*.

Bitzan, J. D. and W. 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, B. A. and J. R. Pierce (2016). Evidence for the Effects of Mergers on Market Power and Efficiency. *National Bureau of Economic Research, No. 22750*.

Bloom, N., R. Sadun, and J. Van Reenen (2012). Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review 102*(1), 167–201.

Bloom, N. and J. Van Reenen (2010). Why Do Management Practices Differ Across Firms and Countries? *Journal of Economic Perspectives 24*(1), 203–224.

Blundell, R. and J. L. Powell (2003). Endogeneity in Nonparametric and Semiparametric Regression Models. In *Advances in Economics and Econometrics*, pp. 312–357. Cambridge University Press.

Borenstein, S., J. Bushnell, and E. Mansur (2023). The Economics of Electricity Reliability. *Journal of Economic Perspectives 37*(4), 181–206.

Borenstein, S. and J. B. Bushnell (2022). Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency. *American Economic Journal: Economic Policy 14*(4), 80–110.

Borrero, M., G. Gowrisankaran, and A. Langer (2024). Ramping Costs and Coal Generator Exit. *University of Arizona Department of Economics Working Paper No. 24-04*.

Bradley & Associates, LLC (2017). *Powering America: The Economic and Workforce Contributions of the US Electric Power Industry*. National Rural Electric Cooperative Association.

Braguinsky, S., A. Ohyama, T. Okazaki, and C. Syverson (2015). Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry. *American Economic Review 105*(7), 2086–2119.

Brand, K., C. Garmon, and T. Rosenbaum (2023). In the Shadow of Antitrust Enforcement: Price Effects of Hospital Mergers from 2009 to 2016. *The Journal of Law and Economics 66*(4), 639–669.

Brot-Goldberg, Z., Z. Cooper, S. V. Craig, and L. Klarnet (2024). Is There Too Little Antitrust Enforcement in the US Hospital Sector? *American Economic Review: Insights 6*(4), 526–542.

Bushnell, J. B., E. T. Mansur, and F. A. Wolak (2012). Vertical Position and Price Effects of Mergers: The Exelon Constellation Merger. *Working Paper*.

Bushnell, J. B. and C. Wolfram (2005). Ownership Change, Incentives and Plant Efficiency: The Divestiture of US Electric Generation Plants. *Center for the Study of Energy Markets, CSEM Working Paper No. 140*.

Bushnell, J. B. and C. Wolfram (2009). The Guy at the Controls: Labor Quality and Power Plant Efficiency. In *International Differences in the Business Practices and Productivity of Firms*, pp. 79–102. University of Chicago Press.

Callaway, B. and P. H. C. Sant'Anna (2021). Difference in Differences with Multiple Time Periods. *Journal of Econometrics 225*(2), 200–230.

Callaway, B. and P. H. Sant'Anna (2020). did: Treatment Effects with Multiple Periods and Groups. *Comprehensive R Archive Network, Feb*.

Chen, Y. (2024). Network Structure and the Efficiency Gains from Mergers: Evidence from US Freight Railroads. *Available at SSRN 4742990*.

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

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

Clark, R. and M. Samano (2022). Incentivized Mergers and Cost Efficiency: Evidence from the Electricity Distribution Industry. *The Journal of Industrial Economics 70*(4), 791–837.

Craig, S. V., M. Grennan, and A. Swanson (2021). Mergers and Marginal Costs: New Evidence on Hospital Buyer Power. *The RAND Journal of Economics 52*(1), 151–178.

Dafny, L., K. Ho, and R. 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, L. and C. Hausman (2016). Market Impacts of a Nuclear Power Plant Closure. *American Economic Journal: Applied Economics 8*(2), 92–122.

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

Davis, R. J., J. S. Holladay, and C. Sims (2022). Coal-Fired Power Plant Retirements in the United States. *Environmental and Energy Policy and the Economy 3*(1), 4–36.

Davis, S. J., J. C. Haltiwanger, K. Handley, B. Lipsius, J. Lerner, and J. Miranda (2024). The (Heterogeneous) Economic Effects of Private Equity Buyouts. *Management Science 70*(3), 1491–1513.

De Loecker, J. and C. Syverson (2021). An Industrial Organization Perspective on Productivity. In *Handbook of Industrial Organization*, Volume 4, pp. 141–223. Elsevier.

De Loecker, J. and F. Warzynski (2012). Markups and Firm-Level Export Status. *American Economic Review 102*(6), 2437–71.

DOJ and FTC (2023). 2023 Merger Guidelines.

Dranove, D. and R. Lindrooth (2003). Hospital Consolidation and Costs: Another Look at the Evidence. *Journal of Health Economics 22*(6), 983–997.

EIA (2015). Analysis of Heat Rate Improvement Potential at Coal-Fired Power Plants. https://www.eia.gov/analysis/studies/powerplants/heatrate/pdf/heatrate.pdf. Last accessed on 2024-01-17.

EIA (2020). Annual Energy Outlook 2020. https://www.eia.gov/outlooks/aeo/pdf/AEO2020%20Full%20Report.pdf. Last accessed on 2024-01-17.

EIA (2024a). EIA Frequently Asked Questions. https://www.eia.gov/tools/faqs/faq.php?id=74&t=11. Last accessed on 2024-07-09.

EIA (2024b). Electric Power Monthly. https://www.eia.gov/tools/faqs/faq.php?id=74&t=11. Last accessed on 2024-07-09.

Eliason, P. J., B. Heebsh, R. C. McDevitt, and J. W. Roberts (2020). How Acquisitions Affect Firm Behavior and Performance: Evidence from the Dialysis Industry. *The Quarterly Journal of Economics 135*(1), 221–267.

Elliott, J. T., G. V. Houngbonon, M. Ivaldi, and P. Scott (2025). Market Structure, Investment, and Technical Efficiencies in Mobile Telecommunications. *Journal of Political Economy 133*(5), 1401–1459.

Emerson Process Management (2016). Using Automation to Improve Plant Heat Rate. https://www.emerson.com/documents/automation/white-paper-using-automation-to-improve-plant-heat-rate-rosemount-en-178352.pdf. Last accessed on 2024-01-17.

Environmental Defense Fund (2017). Improving Heat Rate on Combined Cycle Power Power Plants. https://www.andovertechnology.com/wp-content/uploads/2021/03/C_18_EDF_FINAL.pdf. Last accessed on 2024-01-17.

EPA (2018a). CAMD's Power Sector Emissions Data Guide. https://www.epa.gov/power-sector/camds-power-sector-emission-data-guide. Last accessed on 2024-01-17.

EPA (2018b). Fact Sheet: Overview of the Clean Power Plan. https://archive.epa.gov/epa/cleanpowerplan/fact-sheet-overview-clean-power-plan.html. Last accessed on 2024-01-17.

EPRI (2009). Capital and Maintenance Projects for Efficiency Improvements. http://mydocs.epri.com/docs/CorporateDocuments/Newsletters/GEN/2009.Q1/GENupdate_1Q09.pdf.

EPRI (2022). Proceedings: 2022 EPRI Heat-Rate Improvement Conference. https://www.epri.com/research/programs/113126/results/3002022356. Last accessed on 2024-01-17.

EPRICA (2014). Range and Applicability of Heat Rate Improvements. https://www.epri.com/research/products/000000003002003457. Last accessed on 2024-01-17.

Fabrizio, K. R., N. L. Rose, and C. 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.

Fan, Y. and C. Yang (2025). Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety. *Journal of Political Economy 133*(5), 1401–1459.

Farrell, J. and C. Shapiro (1990). Horizontal Mergers: An Equilibrium Analysis. *The American Economic Review*, 107–126.

Farrell, J. and C. Shapiro (2010). Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition. *The BE Journal of Theoretical Economics 10*(1).

FERC (2012). Analysis of Horizontal Market Power Under the Federal Power Act: Order Reaffirming Commission Policy and Terminating Proceeding. 138 FERC 61,109. *Federal Energy Regulatory Commission*. Docket No. RM11-14-000.

Fitzgerald, P. and J. Gelorme (2015). Thermal Cycling, Air In-Leakage, Impact On Heat Rate, Emissions and Cost of Compliance. In *2015 Heat Rate Improvement Conference*. Electric Power Research Institute.

Focarelli, D. and F. Panetta (2003). Are Mergers Beneficial to Consumers? Evidence from the Market for Bank Deposits. *American Economic Review 93*(4), 1152–1172.

Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review 98*(1), 394–425.

Geurts, K. and J. Van Biesebroeck (2019). Employment Growth Following Takeovers. *The RAND Journal of Economics 50*(4), 916–950.

Gibbons, R. and R. Henderson (2012). What Do Managers Do?: Exploring Persistent Performance Differences among Seemingly Similar Enterprises. In *The Handbook of Organizational Economics*, pp. 680–731. Princeton University Press.

GMI (2024). US and Canada Energy Coverage Letter. https://www.spglobal.com/marketintelligence/en/documents/mi_nonfin_387622_energy-cv-us-and-canada-energy-coverage_letter-final.pdf. Last accessed on 2024-06-30.

Goodman-Bacon, A. (2021). Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics 225*(2), 254–277.

Gowrisankaran, G., A. Langer, and W. Zhang (2025). Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards. *Journal of Political Economy 133*(5), 1460–1502.

Grieco, P., J. Pinkse, and M. Slade (2018). Brewed in North America: Mergers, Marginal Costs, and Efficiency. *International Journal of Industrial Organization 59*, 24–65.

Haldi, J. and D. Whitcomb (1967). Economies of Scale in Industrial Plants. *Journal of Political Economy 75*(4, Part 1), 373–385.

Harrison, T. D. (2011). Do Mergers Really Reduce Costs? Evidence from Hospitals. *Economic Inquiry 49*(4), 1054–1069.

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

Hempling, S. (2018). Inconsistent with the Public Interest: FERC's Three Decades of Deference to Electricity Consolidation. *Energy Law Journal 39*, 233.

Hortaçsu, A., F. Luco, S. L. Puller, and D. Zhu (2019). Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets. *American Economic Review 109*(12), 4302–4342.

Hortaçsu, A. and S. L. Puller (2008). Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Spot Market. *The RAND Journal of Economics 39*(1), 86–114.

Hortaçsu, A. and C. Syverson (2007). Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices. *Journal of Political Economy 115*(2), 250–301.

Ishii, J. (2006). From Investor-Owned Utility to Independent Power Producer. *The Energy Journal 27*(3), 65–90.

Jeziorski, P. (2014). Estimation of Cost Efficiencies from Mergers: Application to US Radio. *The RAND Journal of Economics 45*(4), 816–846.

Jha, A. (2020). Dynamic Regulatory Distortions: Coal Procurement at US Power Plants. *Available at SSRN 3330740*.

Jha, A. and G. Leslie (2023). Start-up Costs and Market Power: Lessons from the Renewable Energy Transition. *Available at SSRN 3603627*.

Joskow, P. L., D. R. Bohi, and F. M. Gollop (1989). Regulatory Failure, Regulatory Reform, and Structural Change in the Electrical Power Industry. *Brookings Papers on Economic Activity. Microeconomics 1989*, 125–208.

Jovanovic, B. and P. L. Rousseau (2002). The Q-Theory of Mergers. *American Economic Review 92*(2), 198–204.

Kalnins, A., L. Froeb, and S. Tschantz (2017). Can Mergers Increase Output? Evidence from the Lodging Industry. *The RAND Journal of Economics 48*(1), 178–202.

Knittel, C. 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, R. B. (2017). Ready-to-Mix: Horizontal Mergers, Prices, and Productivity. *US Census Bureau Center for Economic Studies Paper No. CES-WP-17-38*.

Kwoka, J. (2014). *Mergers, Merger Control, and Remedies: A Retrospective Analysis of US Policy*. MIT Press.

Kwoka, J. and M. Pollitt (2010). Do Mergers Improve Efficiency? Evidence from Restructuring the US Electric Power Sector. *International Journal of Industrial Organization 28*(6), 645–656.

La Forgia, A. (2024). The Impact of Management on Clinical Performance: Evidence from Physician Practice Management Companies. *Management Science 69*(8), 4646–4667.

Lewis, M. S. and K. E. Pflum (2017). Hospital Systems and Bargaining Power: Evidence from Out-of-Market Acquisitions. *The RAND Journal of Economics 48*(3), 579–610.

Luco, F. and G. Marshall (2020). The Competitive Impact of Vertical Integration by Multiproduct Firms. *American Economic Review 110*(7), 2041–64.

Macchiavello, R. and A. Morjaria (2022). Acquisitions, Management, and Efficiency in Rwanda's Coffee Industry. *National Bureau of Economic Research No. 30230*.

MacKay, A. and I. Mercadal (2024). Do Markets Reduce Prices? Evidence from the US Electricity Sector. *Available at SSRN 3793305*.

Malmendier, U., E. Moretti, and F. S. Peters (2018). Winning by Losing: Evidence on the Long-Run Effects of Mergers. *The Review of Financial Studies 31*(8), 3212–3264.

Marschak, J. and W. H. Andrews (1944). Random Simultaneous Equations and the Theory of Production. *Econometrica 12*(3/4), 143–205.

Midrigan, V. and D. Y. Xu (2014). Finance and Misallocation: Evidence from plant-level Data. *American Economic Review 104*(2), 422–458.

Miller, N. H. and M. C. Weinberg (2017). Understanding the Price Effects of the MillerCoors Joint Venture. *Econometrica 85*(6), 1763–1791.

Morris, J. R. and D. Oska (2008). The Likely Effect of the Proposed Exelon-PSEG Merger on Wholesale Electricity Prices. *The Electricity Journal 21*(1), 45–54.

Moshary, S. and C. Slattery (2024). Consolidation and Political Influence in the Auto Retail Industry. *Available at SSRN 4832068*.

Nguyen, S. V. and M. 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.

Niefer, M. J. (2012). Explaining the Divide Between DOJ and FERC on Electric Power Merger Policy. *Energy LJ 33*(2), 505.

Nocke, V. and M. D. Whinston (2022). Concentration Thresholds for Horizontal Mergers. *American Economic Review 112*(6), 1915–48.

PacifiCorp Energy (2016). Fossil Fuel Heat Rate Improvement Plan. https://pscdocs.utah.gov/electric/16docs/1603519/275489PacifiCorpFossilFuelHeatRateImprovPlan4-29-2016.pdf. Last Last accessed on 2024-01-15.

Prager, E. and M. Schmitt (2021). Employer Consolidation and Wages: Evidence from Hospitals. *American Economic Review 111*(2), 397–427.

Rhodes-Kropf, M. and D. T. Robinson (2008). The Market for Mergers and the Boundaries of the Firm. *The Journal of Finance 63*(3), 1169–1211.

Rosaia, N. (2020). Competing Platforms and Transport Equilibrium: Evidence from New York City. *Working Paper*.

Röller, L.-H., J. Stennek, and F. Verboven (2006). Efficiency Gains from Mergers. In F. Ilzkovitz and R. Meiklejohn (Eds.), *European Merger Control: Do We Need an Efficiency Defence?*, Chapter 3. Edward Elgar for the European Commission.

Sargent & Lundy, LLC (2009). *Coal-Fired Power Plant Heat Rate Reductions*. Perrin Quarles Associates, Inc.

Schmitt, M. (2017). Do Hospital Mergers Reduce Costs? *Journal of Health Economics 52*, 74–94.

Setzler, B. (2022). *DiDforBigData: A Big Data Implementation of Difference-in-Differences Estimation with Staggered Treatment*. R package version 0.1.

Staudt, J. E. and J. Macedonia (2014). Evaluation of Heat Rates of Coal Fired Electric Power Boilers. In *Power Plant Pollutant Control "MEGA" Symposium*, pp. 19–21.

Syverson, C. (2011). What Determines Productivity? *Journal of Economic Literature 49*(2), 326–365.

Weber, R. A. and C. F. Camerer (2003). Cultural Conflict and Merger Failure: An Experimental Approach. *Management Science 49*(4), 400–415.

Weinberg, M. (2008). The Price Effects of Horizontal Mergers. *Journal of Competition Law and Economics 4*(2), 433–447.

Whinston, M. D. (2007). Antitrust Policy Toward Horizontal Mergers.

Williamson, O. (1968). Economies as an Antitrust Defense: The Welfare Tradeoffs. *American Economic Review*, 18–36.

Wolak, F. and S. McRae (2008). Merger Analysis in Restructured Electricity Supply Industries: The Proposed PSEG and Exelon Merger (2006). *The Antitrust Revolution*.

Do Mergers and Acquisitions Improve Efficiency? Evidence from Power Plants

Mert Demirer and Ömer Karaduman

Online Appendix

A Data Appendix

This section provides the details of the data sources used in the paper.

A.1 Unit-Level Data

We use EIA Forms 860 and 923, EPA's Continuous Emissions Monitoring Systems (CEMS), GMI, and Velocity Suite to construct a dataset for generator characteristics and production. The EIA forms and CEMS data sources are public, whereas GMI and Velocity Suite are private data providers for energy markets. The EIA Forms cover the universe of generators in the US, whereas the CEMS data includes generators with a capacity above 25 MW that are subject to a set of environmental regulations. GMI and Velocity are data providers that compile unit- and plant-level information from various resources, including EIA, EPA, FERC, and other proprietary sources. We merge these datasets based on generator names and plant identifiers (ORISPL code). The merged dataset comprises monthly panel data that includes information on plants and generators. This information provides regulation status, technology type, installation year, fuel type, coal type, boiler type, boiler model, boiler manufacturer, capacity, fuel cost, prime mover category, dispatch type, whether a unit is connected to the grid, internal generator, whether the unit is marginal or infra-marginal, and whether the unit can switch fuel. We provide more details about some of the variables below.

Generation Under EPA regulations, most fossil fuel power plants are required to make continual compliance determinations for environmental regulations. For this purpose, the EPA collects boiler-level hourly production and emissions data (heat input, gross electricity generation, emissions) from power plants and makes these data publicly available. The coverage of these data corresponds to roughly 96% of US fossil fuel-powered generation in 2018 (EPA, 2018a). While these data are available starting in 1995, they are primarily incomplete before 2000. For this reason, we restrict the study period from January 2000 to March 2023. With these restrictions, the final dataset includes all US fossil fuel generators that comply with the CEMS program, except those in Alaska and Hawaii. This procedure results in an hourly unit-level dataset on generation, fuel input, and heat rate between January 2000 and March 2023. We aggregate these data to weekly levels in some of the analyses employed in the paper.

The heat rate is calculated by dividing the total heat input by the total electricity output at the analysis frequency (hourly, daily, or weekly). If there are significant changes in production within the hour, the heat rate could be very high or very low. This sometimes generates noise in heat rate estimates, especially at small production levels. To account for this, we winsorize the weekly heat rates above 16 or below 6 MMBtu per MWh. This winsorization affects 2.47% of observations.[56]

We match unit-level generation data from CEMS to unit-level data from the abovementioned data sources. While most units are easily matched using the unit name, some do not match as the EPA uses boilers as units, whereas the EIA uses generator names. For those cases, we rely on the EPA's Power Sector Data Crosswalk on the EPA's website.[57] This crosswalk does not include units that retired before 2020. We manually match those retired and other unmatched units based on capacity, installation year, and retirement year information.

Emissions Together with generation and heat input, CEMS also provide hourly emissions data for pollutants SO_2, NO_x, and CO_2. Using these variables, we calculate emission intensity as total emission in each category divided by total generation at the weekly level. To eliminate outliers due to measurement errors, we winsorize these variables at the 1st and 99th percentiles.

Enviromental Programs CEMS provides information on which environmental programs units are subject to. These programs include the Acid Rain Program; Cross-State Air Pollution NO_x Annual Program; Cross-State Air Pollution NO_x Ozone Season Group 1 Program; Cross-State Air Pollution NO_x Ozone Season Group 2 Program; Cross-State Air Pollution NO_x Ozone Season Group 3 Program; Cross-State Air Pollution SO_2 Annual Group 1 Program; Cross-State Air Pollution SO_2 Annual Group 2 Program; Mercury and Air Toxics Standards; New Hampshire NO_x Program; NSPS Greenhouse Gas Rule; Regional Greenhouse Gas Initiative; SIP Call NO_x Budget Trading Program; and Texas SO_2 Trading Program.

Environmental Control Equipment CEMS provides data on environmental control equipment used in boilers for SO2, NOx, and particulate matter (PM) reduction. This includes

[56] In some very rare cases, generation is reported in different units, for example, in KWh instead of MWh. These cases are easy to detect because they lead to extremely high or extremely low and physically implausible heat rates. We correct these cases before further processing the data.
[57] https://www.epa.gov/airmarkets/power-sector-data-crosswalk.

the installation date and type of each piece of equipment. A generator may use multiple pieces of equipment for a particular pollutant. From this dataset, we create control variables that indicate whether a unit has at least one scrubber for each pollutant type.

Capacity Estimation EPA data do not provide capacity information. We infer yearly capacity from generation using the following algorithm. Each year, we keep generators that operate cumulatively for more than 120 hours (5 days). Then, we obtain the annual hourly generation distribution and use the 99th percentile of the observed hourly generation conditional on operating in the CEMS data every year. This algorithm yields a generator capacity that is stable over time for most units. If a unit generates for less than 120 hours, we do not use this algorithm for capacity estimates due to the small sample size. For those units, we backfill capacity information from previous years. To check the accuracy of this algorithm, we run it for the units that have a perfect match in the EPA and EIA data, for which we have the true capacity information from the EIA. We find that the capacity generated from the EPA data aligns with that provided by the EIA.

A.2 Plant-Level Data

We use Velocity Suite and GMI to construct data for plant-level characteristics. From these data sources, we obtain information on location, ISO, FERC region, regulation status, and other important plant-level information. We also obtain data on non-fuel inputs from Velocity Suite, such as capital expenditures, number of personnel, and non-fuel costs. Velocity Suite compiles these data from two sources. The first dataset is the annual FERC Form 1, a comprehensive financial and operating report submitted for electric rate regulation. The second dataset is the Rural Utility Service (RUS) Form "Financial and Operating Report Electric Power Supply". This form is only mandatory for major electric utilities as defined by FERC, so the coverage for these variables is lower than the coverage of other variables.

A.3 Personnel Data

Each power plant subject to at least one EPA program must submit a representative contact to the EPA. This representative information is useful for the EPA, as potential problems need to be addressed quickly, and responsible parties should be accountable. These data include the representative's name, start and end dates, and contact information and are available through the EPA's Envirofacts Data Service API.[58] We use these data on plant representatives from the EPA between 2000 and 2020 to construct personnel data.

[58] https://www.epa.gov/enviro/envirofacts-data-service-api.

However, this database 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% of representatives on LinkedIn. The match rate improves over time, reaching 80–90% in later years. LinkedIn profiles provide a history of job titles, employment, and education. The job titles suggest that about 78% of representatives are plant managers, and the rest are engineers or regulatory compliance managers. Considering that most of these representatives are plant managers, we treat the representative personnel as the plant managers in this study.

This procedure results in monthly plant-level panel data on plant managers. If the managers are successfully matched to LinkedIn profiles, we also know their start and end dates of employment and education history.

A.4 Divestiture Data

Significant deregulation reshaped the power generation industry after the 1990s. To differentiate deregulation-driven divestitures from typical acquisitions, we construct a list of power plants that were subject to divestiture-related ownership changes after 2000. Because no comprehensive public dataset of forced divestitures exists during our sample period, we use multiple data sources. These include Velocity Suite, GMI, EIA, and replication packages of previous academic research, specifically Cicala (2015) and Abito et al. (2024).

We start with data from Cicala (2015), who provides a list of plant divestitures between 1990 and 2009. We also use data from Abito et al. (2024), who provided a list of ownership changes from the EIA Electric Power Monthly reports until 2009, documenting plant transactions, including divestitures and other ownership changes within deregulated markets. Finally, we compile regulatory status information from EIA-860 forms (2006–2020) and identify regulatory status changes, classifying transitions from regulated (RE) to not regulated (NR) as deregulation and from not regulated (NR) to regulated (RE) as re-regulation (which is rare).

Next, we match this data to our sample and manually validate whether the ownership changes correspond to regulatory status changes, using a ±15-month window around the acquisition date recorded in our dataset. Lastly, we examine regulatory status changes recorded in Velocity Suite and GMI datasets not previously identified. We again verify these changes within a ±15-month period.

This approach results in 615 generator divestiture events from 2000 to 2023, which are excluded from our primary acquisition analysis.

A.5 Ownership and Acquisition Data

Every power plant acquisition should be notified to the corresponding state or federal agency. For this reason, the power generation industry has comprehensive information on power plant acquisitions. We construct ownership change data using two separate ownership and transaction datasets from GMI. We augment this dataset using company press releases and newspaper articles about these acquisitions. GMI was previously called SNL Financial and has been used by many researchers to study electricity markets (Davis and Hausman, 2016; Borenstein and Bushnell, 2022; Jha, 2020).

GMI gathers professional and ownership data across various industries, including the energy sector, using multiple sources. For the energy industry in the US, GMI leverages regulatory filings from agencies like the Securities and Exchange Commission (SEC), as well as specific electricity industry-related filings from the FERC, Rural Utilities Service (RUS), EIA, and State-Regulated Utilities (GMI, 2024). Additionally, GMI collects a wide range of data from news aggregators, company websites, press releases, industry reports, interviews, and corporate announcements.[59]

GMI ownership data come in the form of generator-owner-share. The ownership information for each generator share is characterized by the name of the owner company, its percent share of equity in the generating unit, and the owner's ultimate parent company. If a generator's ownership changes over time due to an acquisition, GMI records this by updating the power plant shares as an event with an event ID (transaction ID) and an end date. The status of each share is recorded as either "Current" or "Sold". "Current" shares do not have an end date as they indicate the current ownership, while "Sold" shares do have an end date indicating the end of a past ownership. There are also "Pending" ownership shares, but these represent transactions that have yet to be completed as of March 2023, so these observations are disregarded.

From the raw generator-share data, we construct a monthly panel that records information about the companies that own each generating unit for the duration of our study period. We rely on the dates listed with each ownership share to determine when a generating unit should enter the panel and when ownership changes occur. These data record the companies that own each generating unit, the percent shares attributable to each owner, and each owner's ultimate parent company. If an ownership group is active for less than an entire month, meaning a power plant is acquired after the first of the month and resold before the end of that same month, then we exclude the ownership group involved in this intra-month change from the panel. Intra-month ownership changes account for less

[59] https://www.spglobal.com/marketintelligence/en/solutions/sp-capital-iq-pro.

than one percent of the generator-share data. GMI backfills any company name change, so firm name changes are not reflected as ownership changes. Moreover, GMI maintains a consistent company identifier for owners throughout the panel, so we do not need to rely on company names. To summarize, this procedure results in a month-generator panel dataset 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 acquisitions data. This dataset provides detailed information for every transaction, such as buyers, sellers, transaction type, and deal value. This dataset includes a transaction ID and transaction description where one can see acquired assets, acquisition motives, and other important information. We merge the transaction data to the ownership panel using transaction IDs available in both datasets. The merged data give us a complete picture of ownership changes, including new and previous owners and important transaction characteristics.

In the GMI transaction data, we observe that approximately 21.8% of transactions do not include a deal description. For these transactions, we manually search for information about the companies involved to understand the nature of the transactions. Most of these transactions were cases where the power plant changed ownership between two subsidiaries of the same parent company, either due to corporate restructuring or forced divestitures (for example, from the utility subsidiary to the independent power producer subsidiary, see (Ishii, 2006)). We also notice that a small fraction of these acquisitions are false due to company name changes. For this reason, we exclude the acquisitions with no description and their corresponding ownership changes from our estimating sample.

Finally, in some cases, we observe another ownership or share change occurring shortly after an acquisition. After manually reviewing these cases, we noticed that these are typically follow-up ownership changes related to the first transaction. Therefore, when we detect multiple acquisitions within a 3-month period, we treat them as a single acquisition in our sample. A total of 182 unit acquisitions fall into this category.

A.6 Firm-Level Data

Even though the ownership and transaction data provide buyer and seller names and identifiers, they do not provide information on firms, such as their industry and asset size. To obtain this information, we used another data portal owned by GMI called S&P Capital IQ Pro.[60] S&P Capital IQ Pro and GMI use the same company identifier if the firm is classified as a utility. For other firms, we manually searched for company names

[60] https://www.capitaliq.com/.

in the platforms to create a crosswalk between company identifiers. We could match all company names except for a few companies that went bankrupt or were investment funds. We merge the S&P Capital IQ Pro database with our ownership panel using these company identifiers. We obtained information about firms, such as their industry and publicly listed status.

A.7 Maintenance and Outage Data

The Generating Availability Data System (GADS), managed by the NERC, is a database and reporting system that collects and analyzes data on the performance and reliability of power plants. The collected data helps utilities and other stakeholders analyze performance trends, develop equipment reliability and availability benchmarks, and make informed decisions about plant operations and maintenance. The GADS database is divided into Events, Performance, and Unit datasets containing information on unit maintenance and disruptions at the hourly level, monthly unit generation, and time-invariant unit characteristics, respectively. The intersection of these datasets yields a panel ranging from 2013 to 2021 for 6,914 units that experienced any event.

To our knowledge, our paper is the first in the economics literature to use the GADS data at the generator event level. For this reason, we provide a detailed description of data construction below.

A.7.1 GADS Data Description

The primary focus of the GADS database is the Events dataset, which is aggregated at the event level and describes the duration of disruptions and other issues experienced by generators. These events can be broadly categorized as outages, which indicate a complete disruption of production; derates, which are associated with periods of production lower than expected capacity; non-curtailing events, which do not affect the productive capabilities of units; and inactive periods, during which units are not producing for some reason other than those associated with outages. Depending on the urgency, outages and derates are further categorized as forced, planned, or maintenance events. Forced events must be addressed immediately or near-immediately, whereas planned and maintenance events are disruptions that have been anticipated over a longer period of time; planned events typically coincide with planned inspections and are thus scheduled months in advance, whereas maintenance events are less emergent than forced events but require attention before the next planned event.

The GADS data describe unit generation in addition to the events. These data, called Performance data, are aggregated at the monthly level and report generation in terms of hours, along with descriptive information such as fuel and unit type. NERC also provides time-invariant unit characteristics in the Unit dataset; of particular interest in these data are the unit's geographic location and nameplate capacity.

A.7.2 Processing GADS Data

We use the raw GADS data to construct an events panel unique at the unit-hour level. The foundation of this panel is the Events data, though the performance and unit data supplement the Events data with unit characteristics and production information. The Unit data provide capacity and geographic information of a unit, and the Performance data provide fuel type information, which is taken as the most recently reported fuel type for a given unit, as well as monthly production hours. The processing of these data is minimal; the Performance data contain some duplicate observations that are dropped, but cleaning efforts are otherwise focused on the Events data.

Similar to the Performance data, the Events data include some apparent duplicate observations. The data documentation suggests that the data should be unique at the unit-event level, where a combination of descriptors and date range describes an event. These event descriptors include event type, contribution code, cause code, and amplification code, where contribution code indicates whether an event was the primary cause of disruption, and cause and amplification codes each provide more detail describing the event (such as a particular piece of equipment malfunctioning). We assign our own event identifier based on the combination of descriptors detailed above and drop duplicate events based on this definition. This cleaning step drops 1,380 observations, which account for 0.03% of the raw sample.

The raw Events data are split into yearly files, so individual events that span multiple years must be concatenated manually. The raw data include a flag indicating whether an observation corresponds to an event that continues into the next year or is a continuation of an event from a previous year. We concatenate events across years, using this flag to distinguish between events that span calendar years instead of those that start (end) on the first (last) date of a given year. It follows that an event continuing into the next year should match a corresponding event such that the end and start dates are the same. For a given pair of adjacent years, we concatenate events when appropriate, matching them based on the event descriptors described above (i.e., unit owner, unit, event type, contribution code, cause code, and amplification code) as well as coinciding start and end dates. We repeat this exercise once more to account for events that may span multiple years, matching 75% of events flagged as spanning multiple years.

The Events data also include partially or completely overlapping events that are otherwise identical; in other words, there are events that are identical across descriptors that start at the same time but end at different times, and vice versa. These events are always derating or non-curtailing events, and likely correspond to different periods of work or, alternatively, different periods of capacity restriction. Given our focus on the timing and nature of events and not the extent of work or capacity restrictions, we drop any event that is completely subsumed by another event that is identical in terms of descriptors. One important distinction is partially overlapping events, which we keep in the data as distinct disruptions. This excludes 1,424 superfluous events.

In sum, this process drops 0.7% of the raw sample and yields a dataset defined at the event level, in which a unit can experience multiple concurrent events. Observations are identified by Unit-Owner-Event Type-Cause-Contribution-Amplification-Start Date combinations. From this dataset, we generate a restricted Events dataset that excludes units that are not located in the contiguous United States or produce for less than 100 hours over the sample period. Additionally, hydro and nuclear units are excluded from this restricted set, which accounts for 71.2% of units and 69.3% of reported events in the full Events dataset.

Taking the cleaned Events data as input, we construct a balanced hourly events panel for all units that includes event ID and descriptors for up to five concurrent events. Units seldom experience more than five events simultaneously; these events account for 0.05% of all unit-hours or 0.09% of unit-hours conditional on at least one observed event, so this constraint has minimal qualitative impact.

A.7.3 Matching GADS Units to CEMS Data

The hourly events panel facilitates a granular comparison between GADS event incidence and CEMS production. Additionally, the GADS performance data allow a direct comparison between GADS production and CEMS production data, albeit at the monthly level, due to the limitations of GADS granularity. As a first step, we attempt to match units across datasets by correlating monthly production hours; this approach is supplemented with an hourly comparison, in which we calculate the probability of production conditional on events. This process, described in further detail below, matches 3,988 GADS units to CEMS units.

Monthly Algorithm The monthly algorithm attempts to match units across datasets by correlating monthly production hours. The GADS performance data report monthly production in terms of generation hours; to ease comparison, the hourly CEMS production data are aggregated to the monthly level by calculating the number of hours in a given month during which a unit produces any output. Units are grouped into "buckets" based on state and a broad categorization of fuel type, distinguishing among coal, gas, and other fuels. Within each bucket, production is correlated for all unit pairs across datasets over the months during which both units are available for production. Correlations are calculated with a variety of measures to account for outlier sensitivity; the key measures include Spearman, Kendall, and Pearson coefficients, though we run additional correlations for robustness that winsorize the production hour distributions and, separately, focus on months where both units are actively producing. We calculate the average correlation coefficient across measures and produce a scatter plot for each potential match. We manually reviewed the scatter plots to determine true matches. Each potential match is given a score of 1 through 6, based on the following definitions: definite unit match, definite plant match, probable match, multiple potential matches, no match, and no need to match on account of low production. To supplement the correlation calculation and distinguish among multiple strong candidates, we also compare unit characteristics, such as retirement status or capacity. In sum, this algorithm matches 3,671 (3,469) GADS units to CEMS units.

This process is repeated on the subset of unmatched GADS units with buckets defined by the state, allowing for some flexibility with fuel type when comparing units across datasets. For this additional iteration, we restrict attention to unmatched CEMS units as well as CEMS units that are not matched with utmost confidence; in other words, we exclude CEMS units matched to GADS units for which we are reasonably certain that the unit match is precise. This step generates an additional 304 GADS unit matches.

Hourly Algorithm The hourly algorithm matches units across datasets by calculating the probability of production, conditional on an observed event. We evaluate matches based on the rate at which unit production accords with event occurrence. To underpin this logic, we take the monthly matches as given and plot the distribution of production probabilities conditional on various event types. Focusing in particular on units that are matched with confidence, these distributions show that units experiencing an outage event overwhelmingly do not produce; likewise, a unit experiencing an isolated derate event (i.e., a derating without any other concurrent events) is very likely to produce. These findings square with the GADS documentation describing each of these events. Intuitively, a unit that experiences an outage cannot produce, whereas a unit experiencing a derating event without any other extenuating circumstances (i.e., simultaneous outage or reserve shutdown) should operate at reduced capacity. Approximately 66.4% of units experience an isolated derating event, though 99% of units experience an outage event at some point; the overwhelming majority of units that experience an isolated derating event also experience an outage at some point during the sample period.

Given these patterns, we devise scores to rate the extent to which a unit's production coincides with expectations, given an outage or isolated derating event. We calculate four different scores: (i) a derate score, which is the probability of production, conditional on an isolated derating event; (ii) an outage score, which is the probability of no production, conditional on any outage; (iii) an average score, which is a simple average of the derate and outage scores[61]; and (iv) a composite score: the probability of production conditional on an isolated derating or the probability of no production conditional on any outage.

These scores are calculated over the intersecting periods of GADS event times and CEMS production times. As such, we do not calculate conditional probabilities during times prior to a unit's entry or following a unit's retirement. This is meant to reduce the number of false positives that may arise from an inactive unit perfectly overlapping with an outage event. However, this approach introduces an additional source of false positives in that production and event times may overlap minimally and thus achieve an erroneously high score. To account for this possibility, we also calculate the share of event time during which a unit is available for production and scale the average and composite scores by this share. Doing so minimizes the potential for false positives by scaling down the scores of matches that barely overlap.

As in the initial monthly algorithm, GADS units are compared to CEMS units based on the state and broad fuel grouping. To provide additional focus, comparisons are restricted to those units that operate over similar time periods; this operational overlap is calculated as the share of months during which both units produce any amount of output. We only attempt to match those units whose production shares are within 2.5 percentage points.

We calculate the scores discussed above for all unit combinations within these constraints. We rely on the monthly matches as a benchmark to identify additional matches because these have been manually reviewed and verified. For each score, we calculate the share of GADS units for which the hourly and monthly algorithms generate the same match; the score that yields the highest rate of concurrence across algorithms is considered the optimal score.

Concurrence rates tend to vary considerably across fuel types, so we select optimal scores for each fuel grouping: the optimal score for coal, gas, and other units is composite, average scaled by overlap, and composite scaled by overlap, respectively. For each unit, we calculate the difference in optimal scores between the top matches based on the assumption that the score of a true match will far exceed the next best option. We plot the distribution of these differences, breaking them out by fuel grouping, as well as whether the match concurs with the monthly algorithm match. Focusing on the GADS units that are matched

[61] If a unit does not experience an outage (derate), then the average score is equal to the derate (outage) score.

to different CEMS units (i.e., the hourly algorithm does not correctly identify the unit match given the monthly results), we use this distribution to identify a threshold above which false positives are very unlikely based on the right tail of the distribution. We apply this threshold to GADS units that were not matched to any unit in the monthly algorithm in order to identify additional matches that are unlikely to be false positives. Taking these additional matches, we manually review correlation matrices similar to those generated in the monthly algorithm to weed out erroneous matches, applying match scores based on the scheme outlined in the monthly algorithm. This procedure yields an additional 13 matches to CEMS units.

Match Results This iterative matching process yields 3,994 matched GADS units in total; matches to CEMS units account for 81.1% of GADS units, and matches make up 92.8% of CEMS capacity. The entire process matches approximately 90% of GADS coal units and 87.5% of GADS gas units to CEMS units, while less than 50% of other units are matched. Though we were not able to match every GADS unit, these matches do account for the vast majority of GADS units as well as the vast majority of CEMS capacity. The bulk of unmatched units fall into the "other" fuel category; likewise, the capacity of these units tends to be towards the lower extremes or negligible, suggesting that the most relevant and significant units have been matched.

B Estimation Details

This section provides the details of various estimation procedures employed in the main text.

B.1 Estimation of Residual Productivity

This section explains how we estimate the annual residual log-productivity reported in Figure 2. Our goal is to account for the observable factors that can affect generator productivity and document large heterogeneity in residual generator productivity over time and across firms.

We estimate regressions with a rich set of observables and fixed effects to obtain residual productivity. In particular, in the first step, we use weekly heat rate data aggregated from hourly data and regress the logarithm of the inverse heat rate on time-varying observed plant characteristics and unit-year indicators. These time-varying variables include week fixed effects, state-month fixed effects, regulation status, total load, the number of idle hours, the standard deviation of heat rate, and the number of times the production increases by more than 2% and 5% of the production capacity in that week. By controlling for these factors, we account for the potential effects of production profiles on efficiency. In the second step, we take the estimated unit-year fixed effects and regress them on time-invariant unit characteristics that include capacity, fuel type, boiler manufacturer, and generator model.[62] The second regression accounts for productivity differences explained by observable generator characteristics. We plot the estimated residuals from this second regression in Figure 2. The time-varying observables in the first-step regression explain 45% of the variation in weekly heat rate, and the time-invariant observables explain 42% of the remaining variation in the second step.

B.2 Heat Rate Curve Estimation

We estimate the generator-specific heat rate curves using hourly data before and after the acquisition by controlling for productivity level (percent of capacity) and ramp rate. We define ramp rate as the change in production compared to the previous hour, relative to capacity.

We use the sample of acquired generators in estimating Equation (3) for the treated group. Then, we take the production profile of these generators for the three years preceding the acquisition and years 2 and 3 following the acquisition. We exclude data within

[62]Generator model and characteristics are missing for about 20% of generators. For these, we include a missing dummy variable.

the first year post-acquisition because efficiency improvements take time to materialize, and we aim to measure the long-term effects of the acquisition. We remove generators from the sample if a generator is inactive more than 80% of the time, either during the pre-period or post-period. The results are robust to this restriction, but they tend to be unstable because the estimates for rarely active generators could be extremely noisy for some generation levels in the heat rate curve. We non-parametrically estimate the heat rate curves with this sample using a local polynomial regression. In particular, we use the `loess()` function in R's `stats` package with the default tuning parameters for bandwidth selection. We separately applied the local polynomial regression for each generator pre- and post-acquisition to estimate their heat rate curve.

To construct the control group, we match each acquired generator to a never-acquired generator. For the matching procedure, we follow what is described in Section C.4 except that we match each generator to only one rather than three. After constructing the control sample, we estimate pre- and post-acquisition heat rate curves as if these control generators are acquired at the same time as the matched acquired generators.

We estimate the confidence band for the difference between pre- and post-acquisition heat rate curves for the treated generators using a bootstrap procedure. We resample the treated generators with replacement and estimate the heat rate curve for the sample. We repeat this 500 times and report the 2.5 and 97.5 percentiles of the bootstrap distribution.

B.3 Calculation of CO_2 Emission Reductions Due to Acquisitions

To quantify the efficiency gains from acquisitions in terms of changes in CO_2 emissions for each unit, we limit our analysis to the Subsidiary/Parent Company changes from the first acquisitions. We assume that after the acquisition, generators produce the same amount of electricity as they would have if not acquired. Additionally, we assume a uniform industry-wide efficiency increase of 0.3% per year, applied consistently across months and plants.

Using the CEMS dataset, we analyze monthly CO_2 emissions at the unit level. We identify the month of the first acquisition for each unit and calculate CO_2 emission intensity for every month. Post-acquisition, we adjust this intensity to account for non-acquisition-related gains due to the industry-wide efficiency increase. Then, we aggregate the total generation and CO_2 emissions before the acquisition and compare them to the total generation and implied CO_2 emissions after the acquisition to determine CO_2 intensity changes. Assuming no change in production post-acquisition, we calculate the hypothetical total emissions if the unit had maintained its pre-acquisition CO_2 emission intensity. The total CO_2 emission savings are then determined by the difference between this hypothetical scenario and the actual post-acquisition emissions.

With this set of assumptions, the total cumulative decline in CO_2 emissions between 2000 and 2023 is roughly 360 million tons. This corresponds to the emissions reduction from replacing 800 TWh of electricity generated from natural-gas-fired plants with renewables, assuming CO_2 emissions are roughly 0.4 tons per MWh for gas-fired power plants (EIA, 2024a). With the assumption of a 30% utilization rate for wind power plants (EIA, 2024b), this is roughly equal to 13 GW of capacity investment in wind power plants from 2000 to 2023.

B.4 Details on Heterogeneity Analysis

In this subsection, we describe the constructions of variables used in the heterogeneity analysis.

- **Plant Age > Median:** An indicator variable that equals 1 if the age of the acquired generator is greater than the median. We consider all generators in our main specification to calculate the median age.

- **Unit Capacity > Median:** An indicator variable that equals 1 if the capacity of the acquired generator is larger than the median. To calculate the median capacity, we consider all generators in our main specification and find the median capacity.

- **Serial Acquirer:** An indicator variable that equals 1 if the total capacity acquired between 2000 and 2023 is larger than the median of the total capacity acquired by firms during the same period.

- **Acquirer Size > Median:** An indicator variable that equals 1 if the total capacity of the acquirer pre-transaction is larger than the median capacity of firms involved in a transaction between 2000 and 2023.

- **Cross-Market:** An indicator variable that equals 1 if the capacity owned by the acquirer in the markets involved in the acquisition is 0 before the acquisition. We define a market as a power control area.

B.5 Calculation of Fuel Cost Share in Operating Expenses

To compare fuel costs for gas and coal power plants to total variable costs, we use cost data from Velocity Suite. This dataset, derived from FERC Form 1 (pages 320-323), provides annual electric operating and maintenance expenses for investor-owned utilities. We focus on total fuel expenses and total power production non-fuel operations and maintenance (O&M) expenses for our analysis.

Since the non-fuel O&M expenses cover all fuel types, we first exclude utilities with hydro and nuclear power production. As the report does not include renewable production expenses separately, we focus on the years 2000-2012 to avoid including non-fuel O&M expenses related to renewables. This approach yields 1,900 utility-year observations. We then calculate the ratio of fuel expenses to the sum of fuel and non-fuel O&M expenses. For the 2000-2012 period, this ratio is 79%. Extending the calculation to the 2000-2022 sample results in a slightly lower ratio of 76%.

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 spans 23 years, many generators have been acquired multiple times. Approximately half of the 2,048 units that have ever been acquired experienced multiple ownership changes during the study period. In our main specification, we consider only the first acquisition of each generator because, with multiple acquisitions, the post-acquisition period of the first overlaps with subsequent acquisitions. For those generators, it is unclear how to conduct a proper event study. In this section, we investigate the robustness of our results to this sample restriction by estimating event studies that include all acquisitions.

The first robustness check includes all acquisitions except those within 36 months of each other. We exclude these acquisitions because the post- and pre-acquisition periods overlap. Using this sample, we estimate Equation (2) with some key differences. For each event, we include post-treatment indicator variables for 36 months following the acquisition and pre-treatment indicator variables for 36 months before the acquisition. The treatment variables are set to zero for 36 months after an acquisition and 36 months before the next acquisition. Therefore, we assume treated plants follow the same trends as the control group between the two acquisitions.

C.2 Data Frequency

We estimated our main specification using weekly data, where efficiency is defined as total electricity output divided by total heat input for that week. We chose weekly frequency because it reduces the computational complexity and decreases noise from aggregating hourly data. In this section, we analyze whether our results are robust to changes in data frequency by considering daily frequency.

Estimation with daily data follows the same steps as the estimation with weekly data. We aggregate fuel input and electricity output to a daily level and define daily efficiency as total daily electricity output divided by total daily fuel input. The treatment variables are monthly indicator variables for each month, 36 months before and 36 months after acquisition. We estimate the same specification as in Equation (2), but include the day of the week as an additional control variable. We also tried to estimate the hourly specification but could not do so when including control variables, as the number of observations reached close to 1 billion, and memory requirements exceeded 500 GB.

C.3 Staggered Difference-in-Differences

Our main specification includes weekly heat rate data, but the treatment coefficients are included at the monthly level to increase precision. Since staggered treatment effect estimation requires data frequency to match treatment frequency, we aggregate our data to the monthly level by taking the average weekly heat rates in a given month and estimating the staggered difference-in-differences at the monthly level. We use the never-treated group as the control group. With this sample, we estimate the Callaway and Sant'Anna (2021) method.

Due to the computational complexity that comes with a large number of treated units and a long panel, we were not able to use standard staggered difference-in-differences packages.[63] Instead, we use the R package DiDforBigData (Setzler, 2022), which provides a big-data-friendly and memory-efficient difference-in-differences estimator in the staggered treatment contexts.

C.4 Matching Difference-in-Differences

We match each of our acquired units to the five nearest neighbors from the pool of control units that have never been acquired during our sample period. For each treated unit, we first find the never-treated active units during the acquisition time with the same fuel and technology type but in a different ISO (to prevent spillovers). This never-treated sample constitutes the pool of candidate control units for that unit. Then, we find the nearest neighbor units with respect to 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 control units to be matched to multiple acquired plants. Using these nearest neighbors, we calculate the unit-specific treatment effect as follows:

$$\hat{\Delta} Y_{it} = Y_{it}(1) - \hat{Y}_{it}(1), \tag{8}$$

where $\hat{Y}_{it}(1)$ is the average heat rate of the control units that are matched to i and scaled such that the average outcome of the control one month before the acquisition is the same as the outcome of the treated unit. By indexing the levels to a baseline period, we obtain a unit-specific "difference-in-differences" estimate. We take the average of the unit-specific treatment effects to obtain the final estimates reported in Figure C.4.

To construct the confidence intervals, we employ a bootstrap procedure, where we

[63] For example, the R package provided by the authors (Callaway and Sant'Anna, 2020).

resample without replacement of the treated generators and follow the same matching procedure described above. We repeat this procedure 100 times and obtain a distribution of efficiency gains from the bootstrap samples. We take the 2.5 and 97.5 percentiles of the bootstrap distribution to construct the confidence intervals.

C.5 Observation Weights

In our regressions, we weighted units equally. A natural alternative to this is to weight each generator by its capacity, which would be robust to a potential concern that all efficiency gains come from small units. Moreover, capacity-weighted estimates would be more informative about the total production affected by efficiency gains. To investigate this, we estimate Equations (2) and (3) by weighting units by their capacity in that year. The results from this estimation are reported in Table OA-13, Table OA-15, and Figure OA-23. We find that the efficiency effect is slightly larger when we weigh units by capacity, which is consistent with the findings reported in Table 4 that the efficiency effect is larger for larger units.

C.6 Placebo Test

We use minority acquisitions as a placebo test to control for potential unobservable factors that might influence both acquisitions and changes in efficiency. If such unobservables are present, they would likely affect minority acquisitions as well. Our analysis includes the 663 minority acquisitions from our dataset, characterized by the majority owner remaining the same post-acquisition. Based on the mechanisms of efficiency gains outlined in our study, we would not expect any impact on power plant efficiency from these minority acquisitions. We estimate the same specification as in Equation (2) by treating these minority acquisitions as events. The results, reported in Table OA-13, confirm our expectation that we do not see any significant change in power plant efficiency after these events.

D Additional Figures

Figure OA-1: A Slide from Investor Presentation About Efficiency Claims

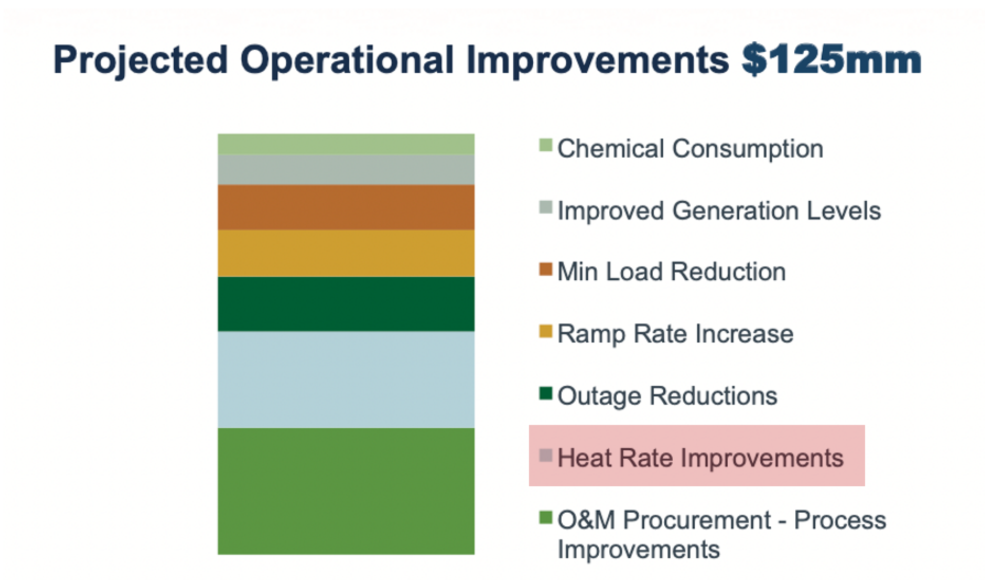

Note: This figure is from a slide deck presented in the conference call of the acquisition of Dynegy by Vistra Energy in 2018.

Figure OA-2: Market Share of Acquired Asset across Transactions

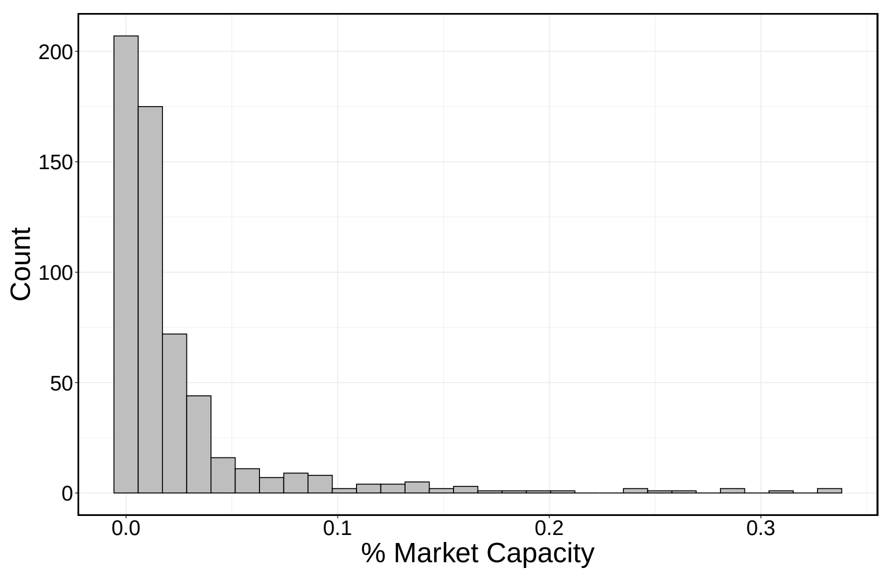

Notes: Distribution of total capacity acquired as a share of total fossil-fuel firm capacity in each market (defined as a power control area). Each observation represents a transaction–market pair.

Figure OA-3: Illustration of Ownership Change Types

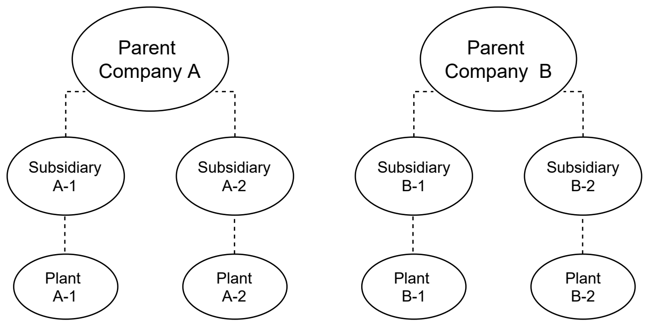

(a) Ownership Structure Before Acquisition

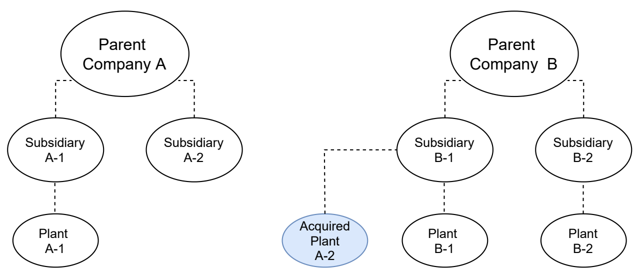

(b) Asset Acquisitions

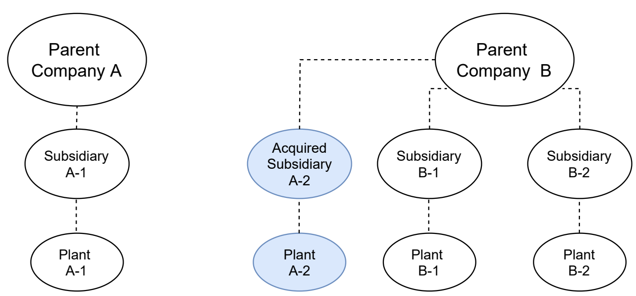

(c) Subsidiary Acquisitions

Note: This figure demonstrates different types of acquisitions. Panel (a) is the corporate structure of companies before the acquisition. Panels (b) and (c) show the corporate structure after the acquisition for asset sales and subsidiary acquisitions, respectively.

Figure OA-4: Slides and Case Studies of Heat Rate Improvements

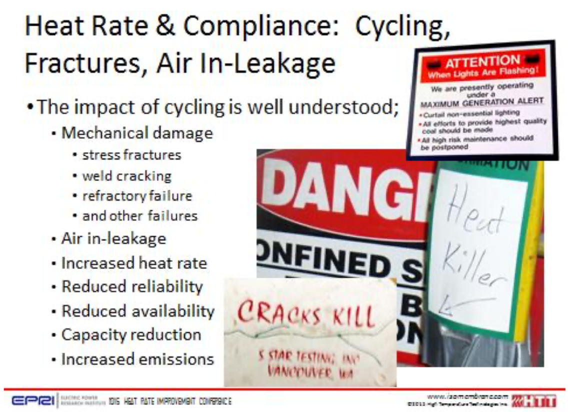

(a) Case Study 1

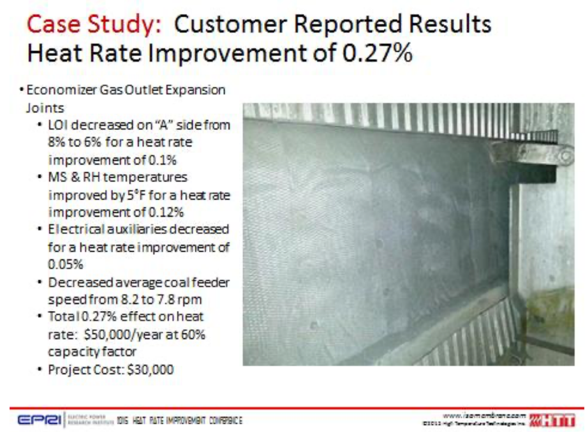

(b) Case Study 2

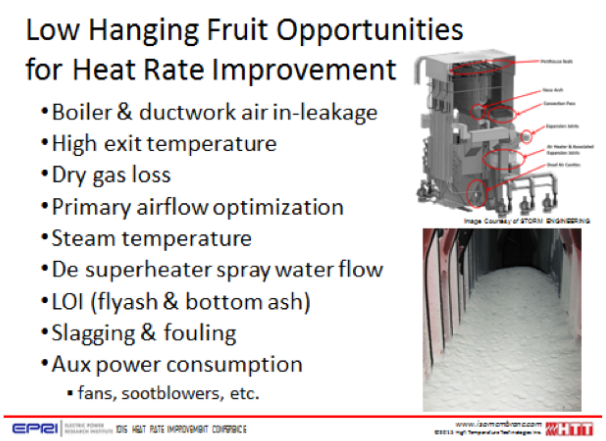

(c) Case Study 3

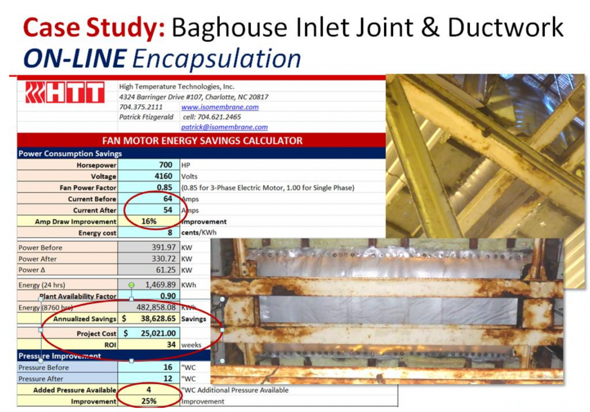

(d) Case Study 4

Note: These pictures demonstrate some methods that were implemented in power plants to improve heat rate. Source: Fitzgerald and Gelorme (2015).

Figure OA-5: Change in US Power Generation Market Concentration by Capacity

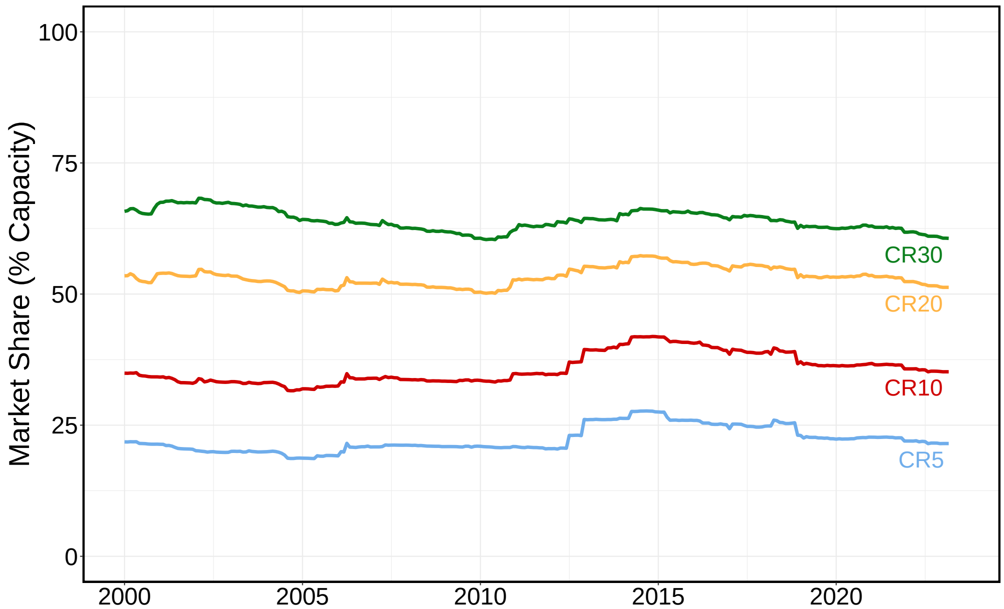

Note: This figure shows the change of national concentration ratios in the overall US fossil fuel power plant market between 2000 and 2023. For every concentration ratio, we calculate the total fossil fuel capacity of the largest corresponding number of firms in the US and divide that by the total fossil fuel capacity in the US.

Figure OA-6: Distribution of Generator Capacity

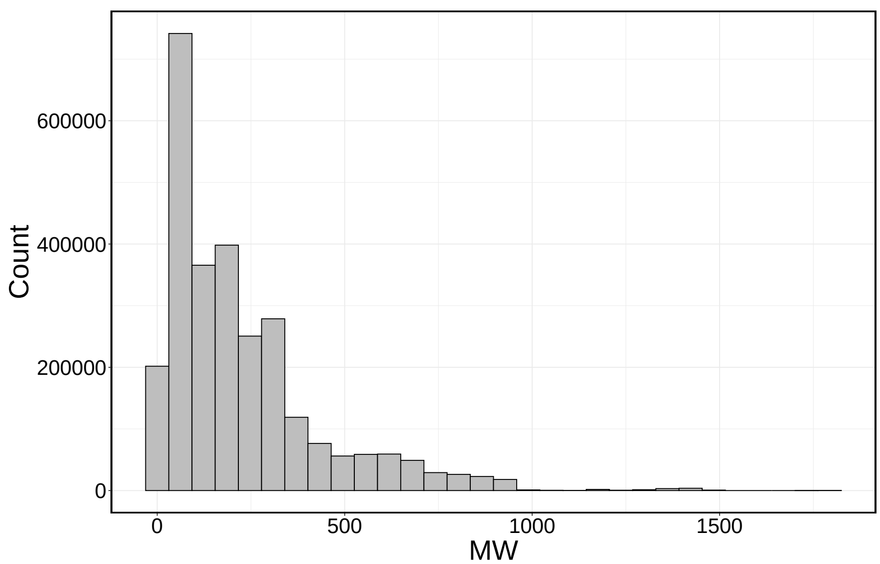

Notes: This distribution shows generator capacity (in MW). Each observation represents a generator-week, so the distribution is weighted by the number of weeks each generator appears in the sample.

Figure OA-7: Firms with Largest Capacity Increase, 2010–2023

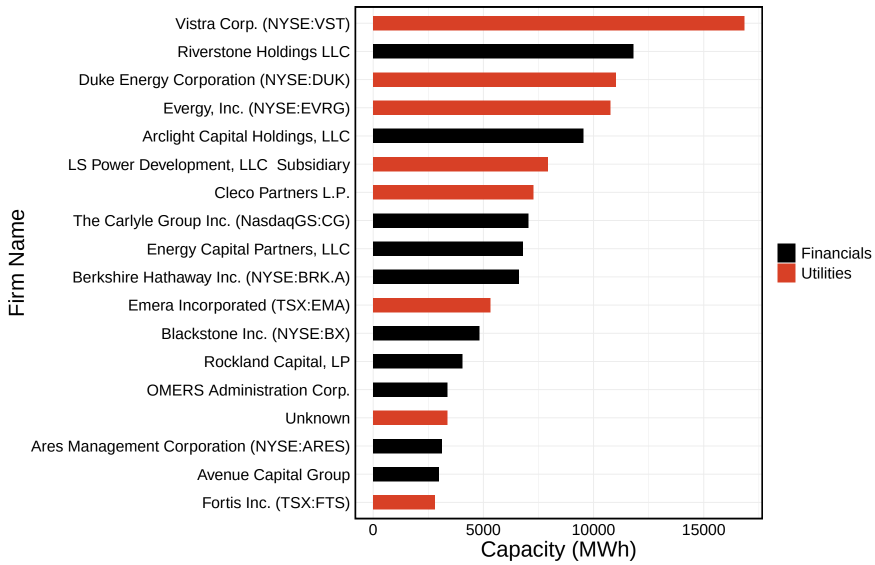

Note: This figure shows firms with the largest capacity increase in fossil fuel generation capacity in the US between 2010 and 2020.

Figure OA-8: Firms with Largest Capacity Decrease, 2010–2023

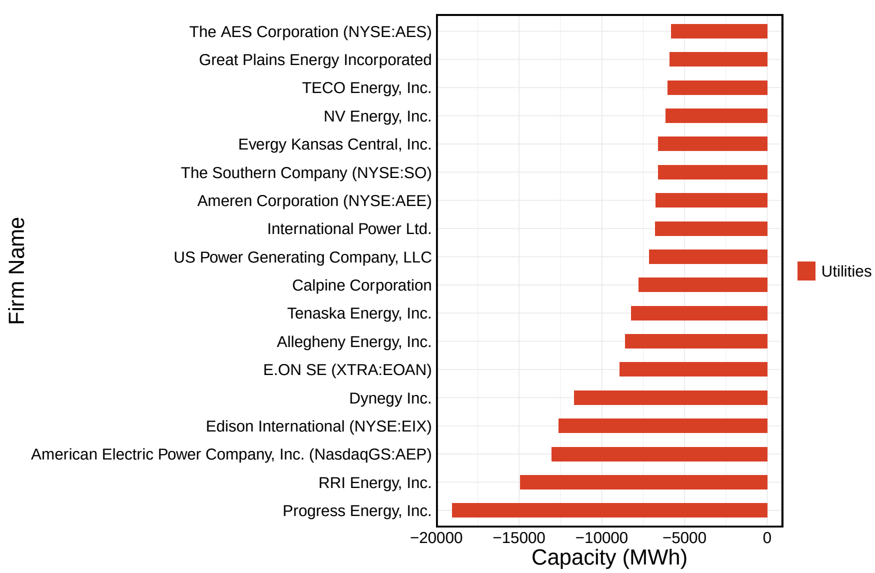

Note: This figure shows firms with the largest capacity decrease in fossil fuel generation capacity in the US between 2010 and 2020.

Figure OA-9: Geographic Distribution of Ownership Changes by Capacity

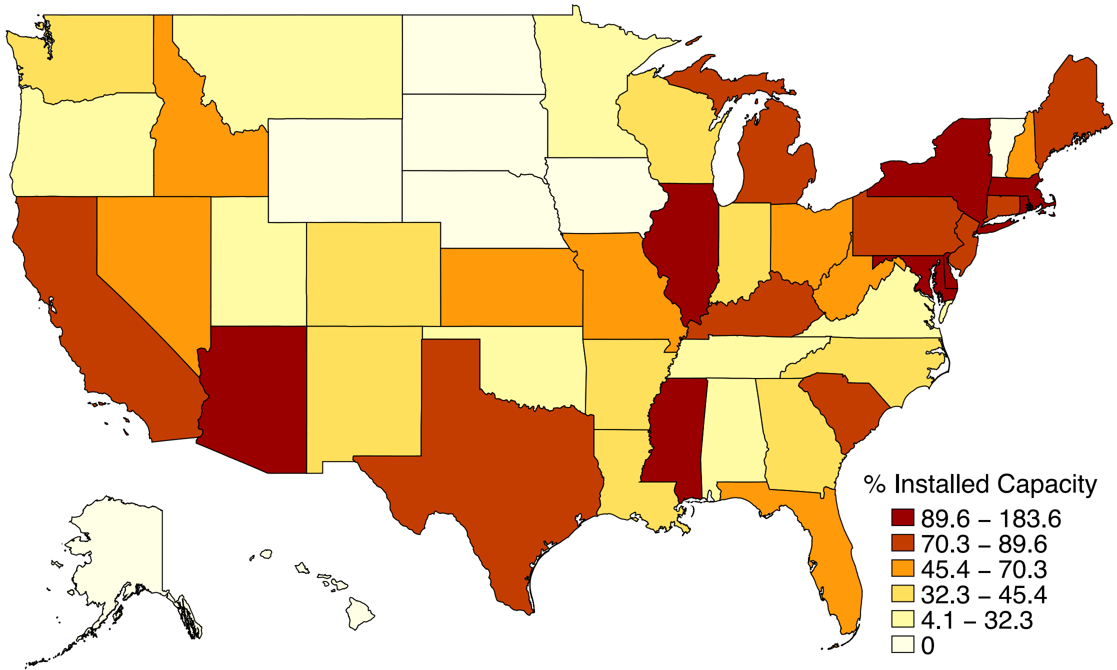

Note: Geographical distribution of power plant acquisitions by capacity between January 2000 - March 2023.

Figure OA-10: The Effects of Manager Changes without Acquisitions on Efficiency

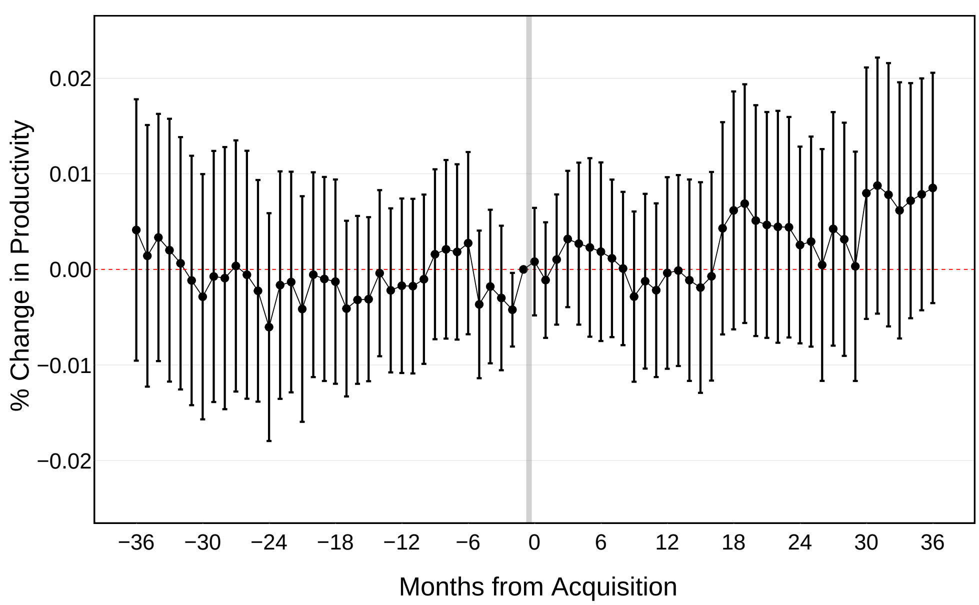

Note: This figure shows the effects of manager change on efficiency estimated using the specification given in Equation (3). In particular, we treat a unit if the manager of that unit changes and there is no acquisition in the three months preceding and following the manager change. Error bars show 95% confidence intervals.

Figure OA-11: Average Within-Plant Annual Productivity Change

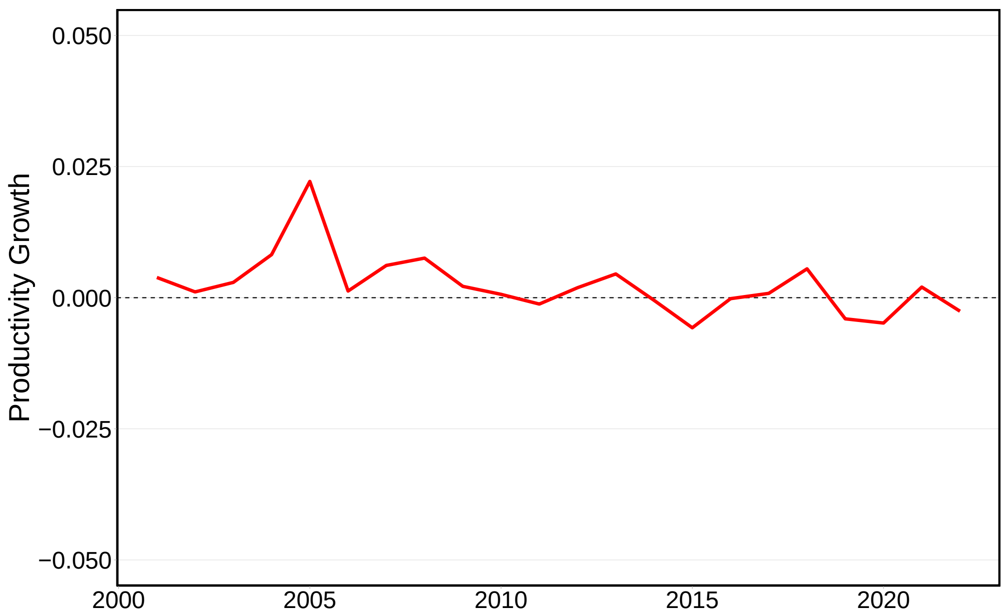

Note: This figure shows average year-to-year within-plant productivity growth for the plants that were not involved in an acquisition.

Figure OA-12: Confidence Band for the Difference in Heat Rate Curves

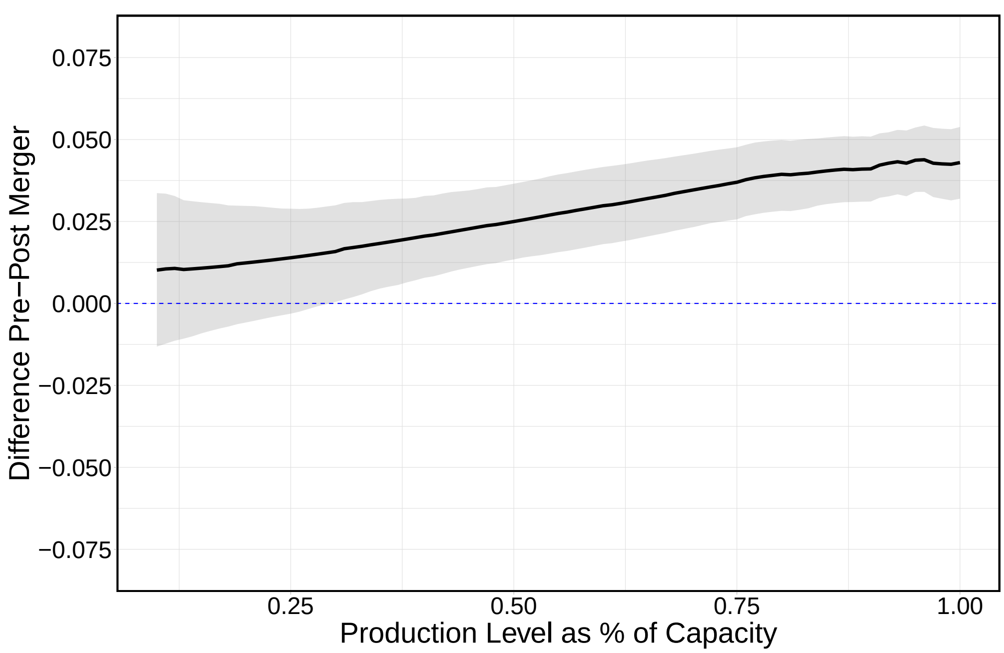

Note: This figure shows the 95% confidence interval for the difference between acquired firms' heat rate curves pre- and post-acquisition, as reported in Figure 7(a). The estimates are reported from 200 bootstrap replications.

Figure OA-13: Effects of Acquisitions on Generator Performance Measures (Event Studies)

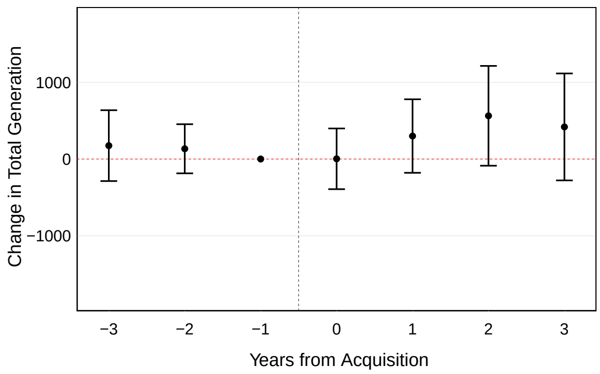

(a) Total Generation

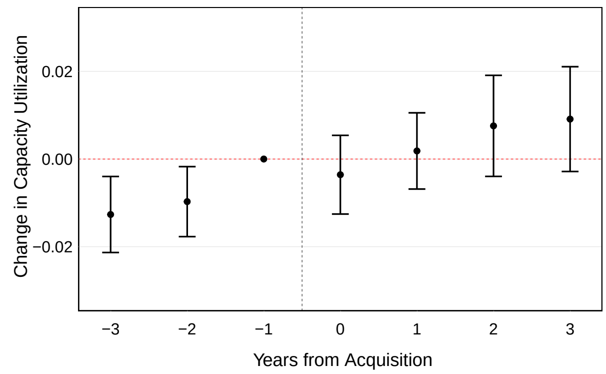

(b) Capacity Utilization

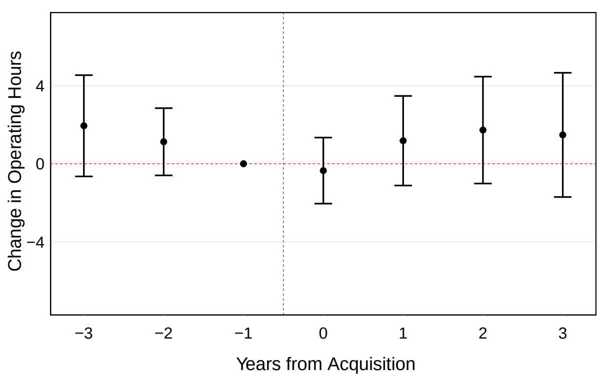

(c) Operating Hours

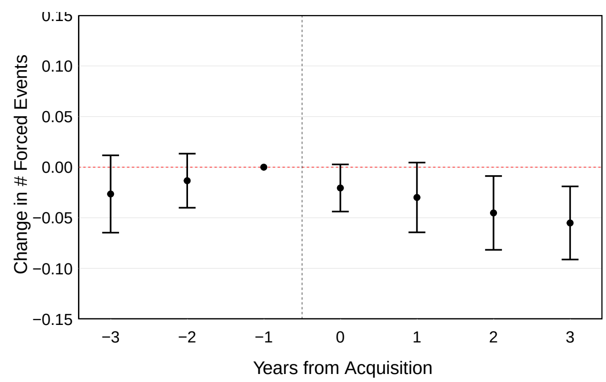

(d) Forced Outages/Derates

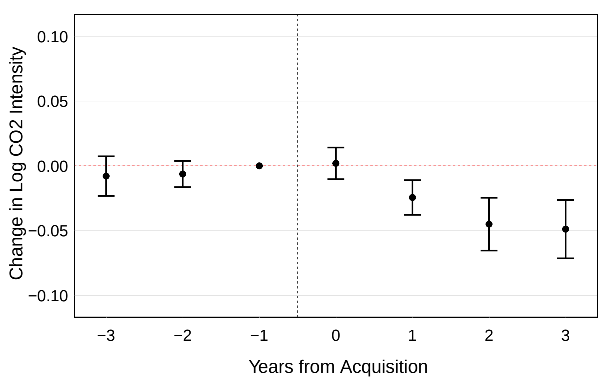

(e) Log CO_2 Intensity

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-3, 3)$ from estimating Equation (2) with yearly treatment indicators to improve precision. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. Unit of observation is generator-week. Error bars show 95% confidence intervals. Standard errors are clustered at the acquisition level.

Figure OA-14: Change in Emissions After Acquisitions

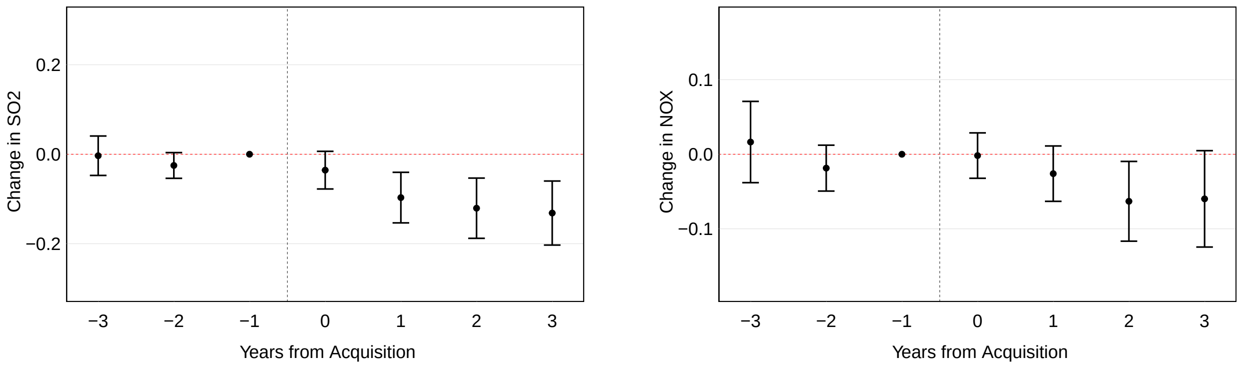

(a) SO_2 **(b)** NO_x

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-3, 3)$ from estimating Equation (2) with yearly treatment indicators to improve precision. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. Unit of observation is generator-week. Error bars show 95% confidence intervals. Standard errors are clustered at the acquisition level.

Figure OA-15: Results on Dynamic Efficiency Mechanism (Event Studies)

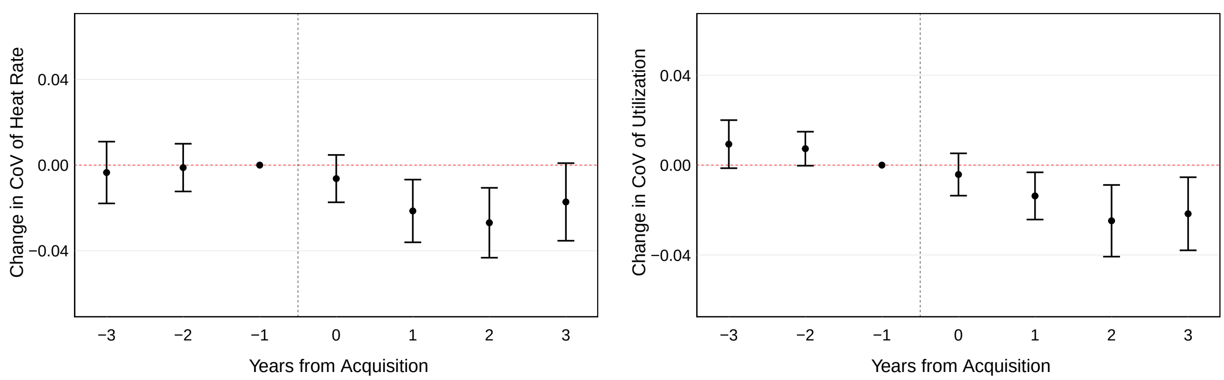

(a) CoV of Heat Rate **(b)** CoV of Utilization

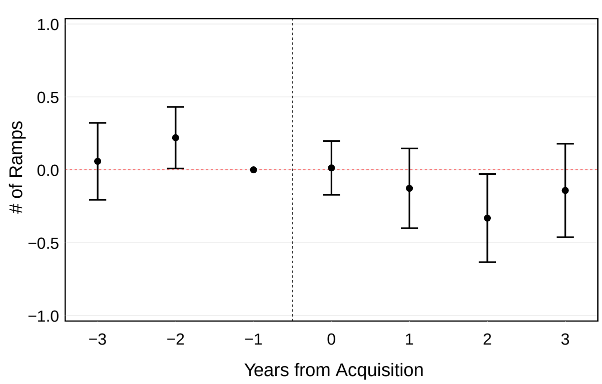

(c) Number of Ramps

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-3, 3)$ from a regression of the CoV of heat rate and utilization and number of ramps on treatment dummies using Equation (2), using yearly treatment indicators to improve precision. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. Error bars show 95% confidence intervals.

Figure OA-16: Dynamic Effects of Acquisitions on Productivity (5-Year)

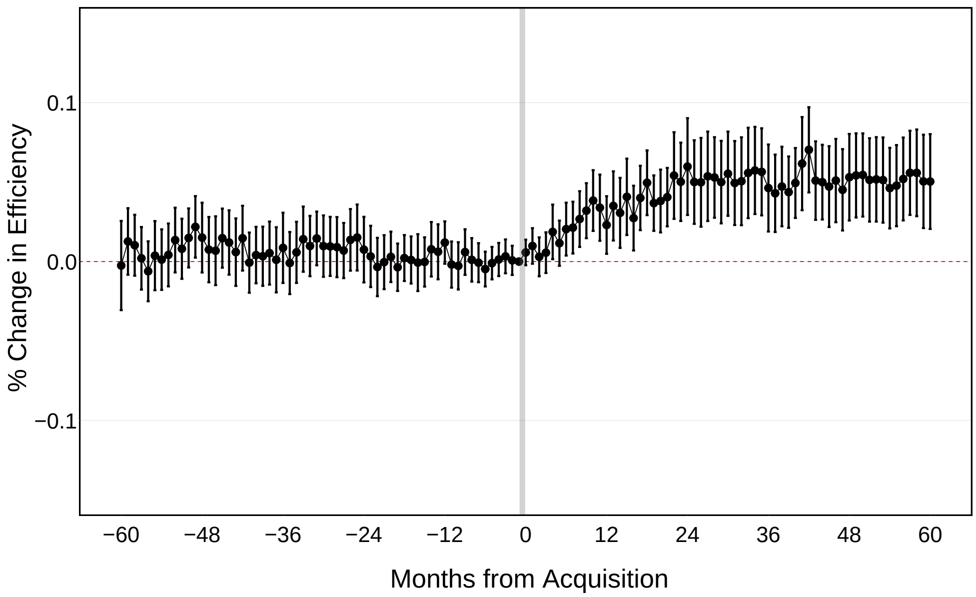

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-60, 60)$ from Equation (3) along with 95% confidence intervals. The dependent variable is the logarithm of the inverse weekly heat rate. The unit of observation is generator-week. Standard errors are clustered at the acquisition level.

Figure OA-17: Effects of Acquisitions on Generator Performance Measures (5-Years)

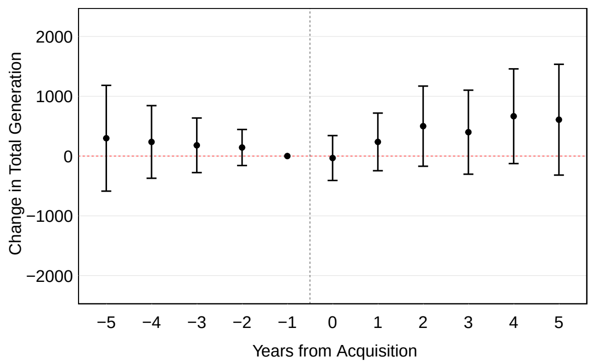

(a) Total Generation

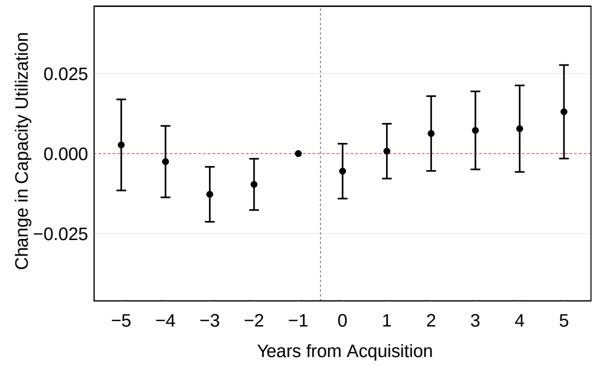

(b) Capacity Utilization

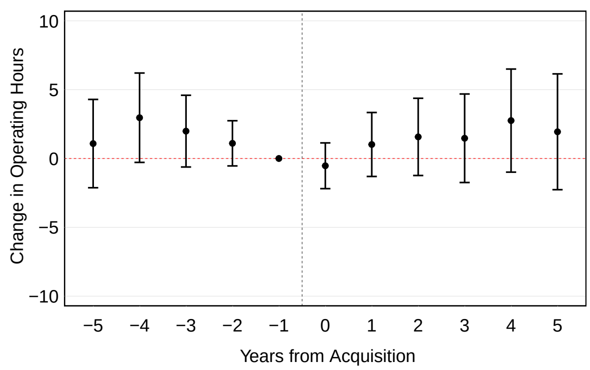

(c) Operating Hours

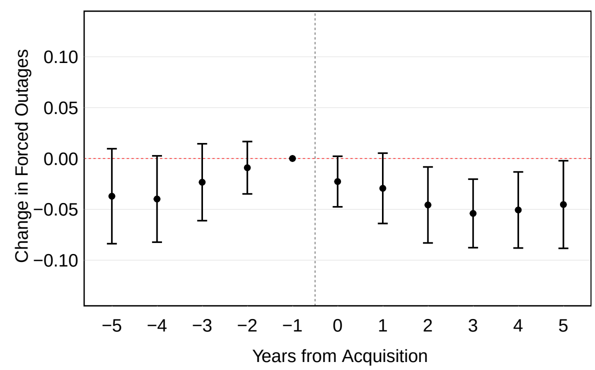

(d) Forced Outages

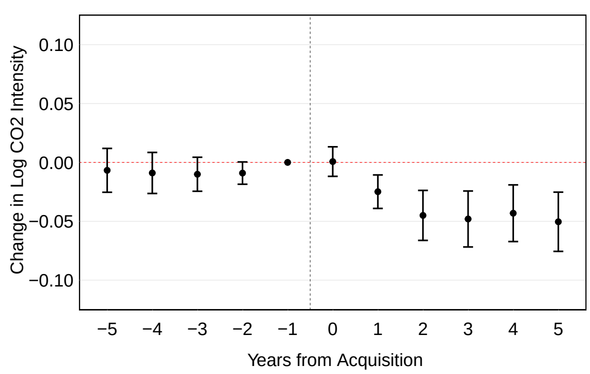

(e) Log CO_2 Intensity

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-5, 5)$ from estimating Equation (2) with yearly treatment indicators to improve precision. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. Unit of observation is generator-week. Error bars show 95% confidence intervals. Standard errors are clustered at the acquisition level.

E Additional Tables

Table OA-1: Largest 25 Acquisitions by Fossil Fuel Power Plant Capacity

Acquirer	Target	Year	Cap. (MW)	# of units
Vistra Energy Corp.	Dynegy Inc.	2018	27198	99
NRG Energy, Inc.	GenOn Energy, Inc.	2012	26174	139
Volt Parent, Lp	Calpine Corporation	2018	22991	127
RRI Energy, Inc.	Mirant Corporation	2000	22748	140
Duke Energy Corporation	Progress Energy, Inc.	2012	19048	134
Duke Energy Corporation	Cinergy Corp.	2006	14923	70
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
Vistra Corp.	TXU Corp.	2007	11116	45
Exelon Corporation	Constellation Energy Group	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
NextEra Energy, Inc.	Engie SA	2017	8604	39
Dynegy Inc.	Duke Energy Corporation	2015	8387	26
Reliant Resources, Inc.	Orion Power Holdings, Inc.	2002	8247	85
AES Corporation	DPL Inc.	2006	7879	33
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

Note: Largest 25 acquisitions in the fossil fuel power generation industry between January 2000 and March 2023. The columns indicate the year the transaction occurred, total production capacity involved in the transaction, and the total number of units that changed ownership.

Table OA-2: Generator Acquisitions by State

State	Number of Units	Number of Unit Acquisitions	Ratio	Regulation Status in 2022
Rhode Island	11	20	1.82	Deregulated
Massachusetts	87	134	1.54	Deregulated
Illinois	237	339	1.43	Deregulated
New York	323	451	1.40	Deregulated
Pennsylvania	218	235	1.08	Deregulated
West Virginia	55	59	1.07	Regulated
New Hampshire	16	17	1.06	Deregulated
Maine	16	16	1.00	Deregulated
New Jersey	197	180	0.91	Deregulated
Nevada	77	67	0.87	Regulated
Ohio	167	139	0.83	Deregulated
Maryland	78	62	0.79	Deregulated
Connecticut	75	59	0.79	Deregulated
Arizona	102	80	0.78	Regulated
Kentucky	117	87	0.74	Regulated
Washington	23	15	0.65	Regulated
California	319	205	0.64	Deregulated
Arkansas	50	32	0.64	Regulated
Georgia	179	111	0.62	Mixed
Mississippi	108	65	0.60	Regulated
Missouri	122	72	0.59	Regulated
Texas	530	301	0.57	Deregulated
Kansas	64	36	0.56	Regulated
South Carolina	88	49	0.56	Regulated
Michigan	141	76	0.54	Mixed
Colorado	82	43	0.52	Regulated
Indiana	150	76	0.51	Regulated
Florida	399	191	0.48	Regulated
Utah	32	15	0.47	Regulated
Delaware	39	17	0.44	Deregulated
New Mexico	39	16	0.41	Regulated
Virginia	115	45	0.39	Mixed
Louisiana	109	41	0.38	Regulated
Oklahoma	100	32	0.32	Regulated
Montana	16	5	0.31	Mixed
North Carolina	162	41	0.25	Regulated
Wisconsin	130	32	0.25	Regulated
Alabama	109	26	0.24	Regulated
Tennessee	95	18	0.19	Regulated
North Dakota	21	2	0.10	Regulated
Oregon	13	1	0.08	Mixed
Wyoming	25	1	0.04	Regulated
Minnesota	82	3	0.04	Regulated

Notes: This table reports four statistics for each state: (i) the number of unique generators, (ii) the number of generator acquisitions, (iii) share of acquisitions relative to total number of generators, and (iv) the state's generation-market regulatory status. The regulatory status is based on the authors' own classification from multiple resources. "Mixed" indicates states that have partial regulation.

Table OA-3: Effects of Acquisitions on Generator Productivity (Subsidiary Change Acquisitions)

	Subsidiary and Parent Changes (1)	Subsidiary and Parent Changes (2)	Subsidiary and Parent Changes (3)	Subsidiary and Parent Changes (4)
	Dependent Variable: Log of Efficiency			
Pre-acquisition (1 Year)	0.004	0.003	-0.001	-0.001
	(0.007)	(0.006)	(0.005)	(0.005)
Post-acquisition (1 Year)	0.03	0.026	0.016	0.016
	(0.008)	(0.007)	(0.006)	(0.006)
Post-acquisition (2 Years)	0.056	0.054	0.039	0.039
	(0.012)	(0.01)	(0.009)	(0.009)
Post-acquisition (3 Years)	0.069	0.071	0.05	0.05
	(0.015)	(0.014)	(0.012)	(0.012)
Ambient Temp. & Humidity	X	X	X	X
Unit & Week FE	X	X	X	X
State by Month FE		X	X	X
Unit Characteristic by Month FE			X	X
Scrubber & Enviro. Prog. FE				X
R^2	0.712	0.735	0.762	0.763
# of Observations	1.494M	1.494M	1.494M	1.494M
# of Never-Treated Units	2311	2311	2311	2311
# of Treated Units	1089	1089	1089	1089

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3$, and δ_4 from estimating Equation (2) only for acquisitions that have both subsidiary and parent ownership changes. Unit characteristic fixed effects include installation year, fuel, technology, and unit capacity bins. The dependent variable is the logarithm of the inverse heat rate. Standard errors are clustered at the acquisition level.

Table OA-4: Effects of Small Acquisitions on Heat Rate (< 10% Market Capacity)

	All Acquisitions Types (1)	All Acquisitions Types (2)	All Acquisitions Types (3)	All Acquisitions Types (4)	Subsidiary and Parent Changes (5)	Only Parent Changes (6)
	Dependent Variable: Log of Efficiency					
Pre-acquisition (1 Year)	0.003	0.001	-0.004	-0.004	-0.002	-0.003
	(0.004)	(0.003)	(0.004)	(0.004)	(0.005)	(0.004)
Post-acquisition (1 Year)	0.021	0.019	0.007	0.007	0.015	-0.004
	(0.006)	(0.006)	(0.005)	(0.005)	(0.007)	(0.006)
Post-acquisition (2 Years)	0.039	0.040	0.022	0.022	0.039	0.007
	(0.009)	(0.008)	(0.008)	(0.008)	(0.010)	(0.008)
Post-acquisition (3 Years)	0.045	0.044	0.024	0.024	0.049	0.001
	(0.012)	(0.011)	(0.010)	(0.010)	(0.012)	(0.009)
Ambient Temp. & Humidity	X	X	X	X	X	X
Unit & Week FE	X	X	X	X	X	X
State by Month FE		X	X	X	X	X
Unit Characteristic by Month FE			X	X	X	X
Scrubber & Enviro. Prog. FE				X	X	X
R^2	0.71	0.73	0.756	0.756	0.765	0.768
# of Observations	1.711M	1.711M	1.711M	1.711M	1.476M	1.438M
# of Controls	2311	2311	2311	2311	2311	2311
# of Treated Units	1590	1590	1590	1590	1030	689

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3$, and δ_4 from estimating Equation (2) only for acquisitions that account for less than 10% of market capacity. Unit characteristic fixed effects include installation year, fuel, technology, and unit capacity bins. The dependent variable is the logarithm of the inverse heat rate. Standard errors are clustered at the acquisition level.

Table OA-5: Heterogeneous Effects of Acquisitions on Productivity (Full Results)

Interaction Var. (Z)	Capacity >Median (1)	Age >Median (2)	Serial Acquirers (3)	Firm Size >Median (4)	Cross-Market Acquisitions (5)
	Dependent Variable: Log of Efficiency				
Pre-acquisition (1 Year)	-0.001	-0.001	-0.002	-0.001	-0.001
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Post-acquisition (1 Year)	0.002	0.016	0.009	0.011	0.015
	(0.009)	(0.01)	(0.008)	(0.008)	(0.009)
Post-acquisition (2 Years)	0.018	0.038	0.014	0.019	0.049
	(0.01)	(0.013)	(0.009)	(0.01)	(0.013)
Post-acquisition (3 Years)	0.029	0.054	0.026	0.034	0.068
	(0.013)	(0.017)	(0.013)	(0.013)	(0.016)
Post-acquisition (1 Year) × Z	0.023	-0.001	0.014	0.012	0.002
	(0.011)	(0.012)	(0.012)	(0.012)	(0.012)
Post-acquisition (2 Years) × Z	0.035	0.004	0.059	0.049	-0.021
	(0.015)	(0.016)	(0.016)	(0.017)	(0.015)
Post-acquisition (3 Years) × Z	0.034	-0.011	0.058	0.041	-0.039
	(0.018)	(0.02)	(0.02)	(0.02)	(0.019)
Ambient Temp. & Humidity	X	X	X	X	X
Unit & Week FE	X	X	X	X	X
Unit Characteristic by Month FE	X	X	X	X	X
Scrubber & Enviro. Prog. FE	X	X	X	X	X
R^2	0.763	0.763	0.763	0.763	0.763
# of Observations	1.494M	1.494M	1.494M	1.494M	1.494M
# of Units	2311	2311	2311	2311	2311
# of Acquisitions	1089	1089	1089	1089	1089

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3$, and δ_4 from estimating Equation (4). Each column reports results from a different regression by varying the interaction variable, Z. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. Unit of observation is generator-week and the dependent variable is the logarithm of the inverse weekly heat rate. Standard errors are clustered at the acquisition level. Details about the heterogeneity variables are provided in Appendix B.4.

Table OA-6: Heterogeneous Effects of Acquisitions on Productivity (Market and State Characteristics)

Interaction Var. (Z)	In ISO (1)	In a High Acquisition Activity State (2)	In a Deregulated State (3)
	Dependent Variable: Log of Efficiency		
Post-acquisition (1 Year) × Z	0.020	0.01	0.014
	(0.014)	(0.013)	(0.013)
Post-acquisition (2 Years) × Z	0.025	0.003	0.044
	(0.018)	(0.017)	(0.015)
Post-acquisition (3 Years) × Z	0.045	0.018	0.052
	(0.020)	(0.023)	(0.018)
Ambient Temp. & Humidity	X	X	X
Unit & Week FE	X	X	X
Unit Characteristic by Month FE	X	X	X
Scrubber & Enviro. Prog. FE	X	X	X
R^2	0.763	0.763	0.763
# of Observations	1.494M	1.494M	1.494M
# of Units	2311	2311	2311
# of Acquisitions	1089	1089	1089

Note: This table presents the coefficient estimates of $\bar{\delta}_2, \bar{\delta}_3$, and $\bar{\delta}_4$ from estimating Equation (4). Each column reports results from a different regression by varying the interaction variable, Z. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. The unit of observation is generator-week, and the dependent variable is the logarithm of the inverse weekly heat rate. Standard errors are clustered at the acquisition level. High-acquisition activity is determined by first calculating the share of acquired units relative to the total number of units in each state. Then, the states that are higher than the median of this measure are classified as high-acquisition activity states. Regulation status is based on the authors' own classification derived from multiple sources; see Table OA-2. States with partial regulation are counted as deregulated for the purposes of this regression.

Table OA-7: Effects of Acquisitions with or without Manager Changes

	Acquisitions w/ Manager Changes (1)	Acquisitions w/o Manager Changes (2)
Dependent Variable: Log of Efficiency		
Pre-acquisition (1 Year)	-0.008	0.013
	(0.006)	(0.01)
Post-acquisition (1 Year)	0.011	0.025
	(0.008)	(0.013)
Post-acquisition (2 Years)	0.038	0.039
	(0.011)	(0.017)
Post-acquisition (3 Years)	0.052	0.057
	(0.014)	(0.022)
Ambient Temp. & Humidity	X	X
Unit & Week FE	X	X
Unit Characteristic by Month FE	X	X
Scrubber & Enviro. Prog. FE	X	X
R^2	0.777	0.777
# of Observations	1.34M	1.254M
# of Never-Treated Units	2311	2311
# of Treated Units	691	331

Note: This table presents the coefficient estimates of δ_1-δ_4 and $\bar{\delta}_1$-$\bar{\delta}_4$ from estimating Equation (2). Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. The unit of observation is generator-week, and the dependent variable is the logarithm of the inverse weekly heat rate. Standard errors are clustered at the acquisition level. Columns (1-2) split these acquisitions into two subsamples. In column (2), we include acquisitions with accompanying manager changes; in column (3), we include acquisitions with no manager changes. For Columns (2-3), the sample period for the regression is between 2000 and 2020 due to the availability of manager data.

Table OA-8: Poisson Regressions

Dep. Var.	Total Load (MWh) (1)	Mean Load (%) (2)	Operating Time (hrs) (3)
Pre-acquisition (1 Year)	-0.019	0.004	-0.031
	(0.02)	(0.007)	(0.021)
Post-acquisition (1 Year)	0.001	0.009	-0.009
	(0.029)	(0.007)	(0.026)
Post-acquisition (2 Years)	0.046	0.020	0.007
	(0.039)	(0.008)	(0.032)
Post-acquisition (3 Years)	0.041	0.022	-0.009
	(0.04)	(0.009)	(0.034)
Ambient Temp. & Humidity	X	X	X
Unit & Week FE	X	X	X
Unit Characteristic by Month FE	X	X	X
Scrubber & Enviro. Prog. FE	X	X	X
# of Observations	2.592M	1.494M	2.597M
# of Controls	2311	2311	2311
# of Treated Units	1089	1089	1089

Notes: This table presents the percent effects of acquisitions on different outcome measures from estimating Poisson regressions. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. The unit of observation is generator-week.

Table OA-9: Effects of Acquisitions on Emissions Intensity

Dep. Var.	Log NOx Intensity (1)	Log SO_2 Intensity (2)
Pre-acquisition (1 Year)	-0.013	0.01
	(0.015)	(0.016)
Post-acquisition (1 Year)	-0.045	0.003
	(0.022)	(0.019)
Post-acquisition (2 Years)	-0.066	-0.044
	(0.027)	(0.026)
Post-acquisition (3 Years)	-0.058	-0.059
	(0.03)	(0.033)
Ambient Temp. & Humidity	X	X
Unit & Week FE	X	X
Unit Characteristic by Month FE	X	X
Scrubber & Enviro. Prog. FE	X	X
R^2	0.873	0.949
# of Observations	1.493M	1.441M
# of Controls	2310	2216
# of Treated Units	1089	985

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3$, and δ_4 from estimating Equation (2). Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. The unit of observation is generator-week. Standard errors are clustered at the acquisition level.

Table OA-10: Heterogeneous Effects of Acquisitions on Productivity (Single Regression)

Interaction Var. (Z)		Capacity >Median	Age >Median	Serial Acquirers	Firm Size >Median	Cross-Market Acquisitions
		(1)	(2)	(3)	(4)	(5)
		Panel A: Heterogeneous Effects on Efficiency				
Post-acquisition (1 Year) × Z		0.019	–0.001	0.011	–0.002	–0.001
		(0.012)	(0.013)	(0.030)	(0.028)	(0.012)
Post-acquisition (2 Years) × Z		0.021	–0.011	0.068	–0.015	–0.027
		(0.015)	(0.017)	(0.034)	(0.034)	(0.015)
Post-acquisition (3 Years) × Z		0.019	–0.035	0.101	–0.048	–0.051
		(0.020)	(0.022)	(0.040)	(0.038)	(0.019)
		Panel B: Effect on Most Commonly Observed Types				
N	Post-acq (3 Years)					
104	0.075	X		X	X	X
	(0.021)					
96	0.003					X
	(0.016)					
90	0.073	X				
	(0.026)					
88	0.073		X	X	X	
	(0.024)					
87	0.092	X	X	X	X	
	(0.021)					
87	0 .020		X			
	(0.015)					
83	0.022	X				X
	(0.020)					
68	0.022		X	X	X	X
	(0.023)					
67	0.054					
	(0.023)					
53	0.108			X	X	
	(0.034)					

Notes: Panel A presents results from the estimation of heterogeneous effects in a single regression where all interaction variables are included. Panel B lists the 10 most commonly observed acquisition types based on combinations of the observable characteristics used in the heterogeneity analysis (denoted as Z). For example, the first type is an acquisition of a unit with above-median unit capacity and below-median unit age by a serial, above-median size acquirer in a cross-market acquisition. The column "N" reports the number of unit acquisitions in each category, and the "Post-acq (3 Years)" column shows the sum of the corresponding coefficients, including the constant term, which corresponds to the omitted category.)

Table OA-11: Long-Run Effects of Acquisitions on Heat Rate (5-Years)

	All Acquisitions Types (1)	All Acquisitions Types (2)	All Acquisitions Types (3)	All Acquisitions Types (4)	Subsidiary and Parent Changes (5)	Only Parent Changes (6)
	Dependent Variable: Log of Efficiency					
Pre-acquisition (1 Year)	-0.006	-0.004	-0.007	-0.007	-0.005	-0.014
	(0.007)	(0.007)	(0.005)	(0.005)	(0.005)	(0.006)
Post-acquisition (1 Year)	0.010	0.012	0.003	0.003	0.013	-0.017
	(0.008)	(0.008)	(0.006)	(0.006)	(0.007)	(0.008)
Post-acquisition (2 Years)	0.027	0.031	0.017	0.017	0.035	-0.010
	(0.009)	(0.009)	(0.007)	(0.007)	(0.010)	(0.008)
Post-acquisition (3 Years)	0.030	0.032	0.016	0.016	0.046	-0.016
	(0.012)	(0.011)	(0.009)	(0.009)	(0.012)	(0.009)
Post-acquisition (4 Years)	0.034	0.032	0.013	0.013	0.045	-0.014
	(0.012)	(0.011)	(0.009)	(0.009)	(0.012)	(0.009)
Post-acquisition (5 Years)	0.039	0.041	0.018	0.018	0.045	-0.009
	(0.012)	(0.012)	(0.009)	(0.009)	(0.012)	(0.010)
Unit & Week FE	X	X	X	X	X	X
State by Month FE		X	X	X	X	X
Unit Characteristic by Month FE			X	X	X	X
Scrubber & Enviro. Prog. FE				X	X	X
R^2	0.706	0.725	0.752	0.753	0.761	0.765
# of Observations	1.94M	1.94M	1.94M	1.94M	1.547M	1.632M
# of Controls	2311	2311	2311	2311	2311	2311
# of Treated Units	2046	2046	2046	2046	1089	1142

Note: This table presents the coefficient estimates from estimating Equation (2) with 5-year post acquisition coefficients. Columns (4-6) present our baseline specification, where we allow for time trends to vary flexibly by unit characteristic and include weather, scrubber, and environmental program controls. Unit characteristic fixed effects include installation year, fuel, technology, and unit capacity bins. The dependent variable is the logarithm of the inverse heat rate. Standard errors are clustered at the acquisition level. Table OA-3 presents the same analysis results but for the subsample of acquisitions with both subsidiary and parent company changes.

Table OA-12: Heterogeneous Effects of Acquisitions on Manager Turnover

Independent Variable:	Capacity > Median (1)	Age > Median (2)	Serial Acquirers (3)	Firm Size > Median (4)	Cross-Market Acquisitions (5)
Manager Turnover	-0.034	-0.069	0.020	0.075	-0.026
	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)

Notes: Regression of an indicator variable for whether a manager is replaced following an acquisition on observable unit/firm/acquisition characteristics. Each observation represents an acquired generator.

F Robustness Checks Results

Table OA-13: Effects of Acquisitions on Generator Productivity (Robustness)

	After 2010 (1)	Weighted Regressions (2)	Net Generation (3)	All Acquisitions (4)	Matching (5)	Minority (Placebo) (6)
	Dependent Variable: Log of Efficiency					
Pre-acquisition (1 Year)	-0.002	-0.001	-0.004	-0.001	-0.002	-0.003
	(0.005)	(0.005)	(0.005)	(0.005)	(0.002)	(0.005)
Post-acquisition (1 Year)	0.002	0.016	0.009	0.013	0.021	-0.003
	(0.007)	(0.006)	(0.006)	(0.006)	(0.004)	(0.007)
Post-acquisition (2 Years)	0.021	0.039	0.028	0.027	0.036	-0.013
	(0.011)	(0.009)	(0.008)	(0.008)	(0.006)	(0.009)
Post-acquisition (3 Years)	0.037	0.05	0.033	0.033	0.047	-0.011
	(0.015)	(0.012)	(0.01)	(0.009)	(0.007)	(0.009)
Ambient Temp. & Humidity	X	X	X	X	X	
Unit & Week FE	X	X	X	X	X	
Unit Characteristic by Month FE	X	X	X	X	X	
Scrubber & Enviro. Prog. FE	X	X	X	X	X	
R^2	0.769	0.763	0.667	0.77	-	0.783
# of Observations	1.387M	1.494M	1.491M	1.769M	-	1.407M
# of Never-Treated Units	2311	2311	2308	2311	1089	2311
# of Treated Units	529	1089	1089	1541	1089	663

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3$, and δ_4 from estimating Equation (2) with various robustness checks, discussed in detail in Section C. Column (1) excludes acquisitions prior to 2010 that may have resulted from deregulation. Column (2) estimates the effect of acquisitions on efficiency measured by total output net of any energy input. Column (3) includes event-study estimates for units acquired multiple times during the sample period. Column (4) reports regression estimates weighed by yearly unit capacity. Column (5) includes results from the matching estimation, in which treated units are matched to controls in different ISOs based on fuel type, age, and capacity. Finally, Column (6) is a placebo specification that considers changes in minority ownership as a treatment event. Unit characteristic fixed effects include state, installation year, fuel type, technology type, and unit capacity bins. Standard errors are clustered at the acquisition level.

Table OA-14: Effects of Acquisitions on Generator Productivity (First and Subsequent Acquisitions)

	All Acquisitions Types (1)	All Acquisitions Types (2)	All Acquisitions Types (3)	All Acquisitions Types (4)	Subsidiary or Parent Changes (5)	Only Parent Changes (6)
	Dependent Variable: Log of Efficiency					
Pre-acquisition (1 Year)	0	0	-0.002	-0.002	-0.001	-0.015
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)
Post-acquisition (1 Year)	0.01	0.01	0.006	0.005	0.013	-0.015
	(0.004)	(0.004)	(0.004)	(0.004)	(0.006)	(0.005)
Post-acquisition (2 Years)	0.019	0.02	0.012	0.012	0.027	-0.01
	(0.005)	(0.005)	(0.005)	(0.005)	(0.008)	(0.005)
Post-acquisition (3 Years)	0.022	0.022	0.011	0.011	0.033	-0.016
	(0.006)	(0.006)	(0.006)	(0.006)	(0.009)	(0.006)
Ambient Temp. & Humidity	X	X	X	X	X	X
Unit & Week FE	X	X	X	X	X	X
State by Month FE		X	X	X	X	X
Unit Characteristic by Month FE			X	X	X	X
Scrubber & Enviro. Prog. FE				X	X	X
R^2	0.713	0.731	0.76	0.761	0.77	0.767
# of Observations	2.335M	2.335M	2.335M	2.335M	1.769M	1.795M
# of Never-Treated Units	2311	2311	2311	2311	2311	2311
# of Treated Units	3515	3515	3515	3515	1541	1449

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3,$ and δ_4 from estimating Equation (2) using all acquisitions. Unit characteristic fixed effects include installation year, fuel type, technology type, and unit capacity bins. The acquisition sample is described in Section C.1. Standard errors are clustered at the acquisition level.

Table OA-15: Effects of Acquisitions on Generator Productivity (Weighted Regressions)

	All Acquisitions Types (1)	All Acquisitions Types (2)	All Acquisitions Types (3)	All Acquisitions Types (4)	Subsidiary or Parent Changes (5)	Only Parent Changes (6)
	Dependent Variable: Log of Efficiency					
Pre-acquisition (1 Year)	-0.003	-0.006	-0.008	-0.008	-0.001	-0.014
	(0.005)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)
Post-acquisition (1 Year)	0.013	0.009	0.001	0.001	0.016	-0.016
	(0.007)	(0.006)	(0.005)	(0.005)	(0.006)	(0.006)
Post-acquisition (2 Years)	0.03	0.028	0.015	0.015	0.039	-0.009
	(0.009)	(0.008)	(0.007)	(0.007)	(0.009)	(0.007)
Post-acquisition (3 Years)	0.034	0.031	0.015	0.015	0.05	-0.014
	(0.011)	(0.011)	(0.009)	(0.009)	(0.012)	(0.008)
Ambient Temp. & Humidity	X	X	X	X	X	X
Unit & Week FE	X	X	X	X	X	X
State by Month FE		X	X	X	X	X
Unit Characteristic by Month FE			X	X	X	X
Scrubber & Enviro. Prog. FE				X	X	X
R^2	0.707	0.726	0.752	0.753	0.763	0.764
# of Observations	1.838M	1.838M	1.838M	1.838M	1.494M	1.575M
# of Never-Treated Units	2311	2311	2311	2311	2311	2311
# of Treated Units	2046	2046	2046	2046	1089	1142

Note: This table presents the coefficient estimates of δ_1, δ_2, δ_3, andδ_4 from estimating Equation (2) by weighting observations by capacity as described in Section C.5. Unit characteristic fixed effects include installation year, fuel, technology, and unit capacity bins. Standard errors are clustered at the acquisition level.

Table OA-16: Effects of Acquisitions on Generator Productivity (Acquisitions After 2010)

	All Acquisitions Types (1)	All Acquisitions Types (2)	All Acquisitions Types (3)	All Acquisitions Types (4)	Subsidiary or Parent Changes (5)	Only Parent Changes (6)
	Dependent Variable: Log of Efficiency					
Pre-acquisition (1 Year)	0.002	-0.001	-0.005	-0.006	-0.002	-0.008
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)
Post-acquisition (1 Year)	0.006	0.004	-0.006	-0.006	0.002	-0.014
	(0.007)	(0.007)	(0.006)	(0.006)	(0.007)	(0.007)
Post-acquisition (2 Years)	0.015	0.017	0.004	0.004	0.021	-0.01
	(0.01)	(0.009)	(0.008)	(0.008)	(0.011)	(0.007)
Post-acquisition (3 Years)	0.021	0.023	0.01	0.01	0.037	-0.013
	(0.012)	(0.013)	(0.01)	(0.01)	(0.015)	(0.009)
Ambient Temp. & Humidity	X	X	X	X	X	X
Unit & Week FE	X	X	X	X	X	X
State by Month FE		X	X	X	X	X
Unit Characteristic by Month FE			X	X	X	X
Scrubber & Enviro. Prog. FE				X	X	X
R^2	0.72	0.74	0.764	0.764	0.769	0.768
# of Observations	1.657M	1.657M	1.657M	1.657M	1.387M	1.507M
# of Never-Treated Units	2311	2311	2311	2311	2311	2311
# of Treated Units	1170	1170	1170	1170	529	819

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3, \text{and} \delta_4$ from estimating Equation (2) for acquisitions occurring after 2010. Unit characteristic fixed effects include installation year, fuel type, technology type, and unit capacity bins. Standard errors are clustered at the acquisition level.

Table OA-17: Effects of Acquisitions on Generator Productivity (Net Generation)

	All Acquisitions Types (1)	All Acquisitions Types (2)	All Acquisitions Types (3)	All Acquisitions Types (4)	Subsidiary or Parent Changes (5)	Only Parent Changes (6)
	Dependent Variable: Log of Efficiency					
Pre-acquisition (1 Year)	-0.004	-0.005	-0.008	-0.008	-0.004	-0.009
	(0.005)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
Post-acquisition (1 Year)	0.003	0.004	-0.005	-0.005	0.009	-0.016
	(0.008)	(0.006)	(0.005)	(0.005)	(0.006)	(0.008)
Post-acquisition (2 Years)	0.018	0.022	0.008	0.008	0.028	-0.007
	(0.009)	(0.007)	(0.006)	(0.006)	(0.008)	(0.007)
Post-acquisition (3 Years)	0.023	0.024	0.008	0.008	0.033	-0.009
	(0.01)	(0.009)	(0.007)	(0.007)	(0.01)	(0.008)
Ambient Temp. & Humidity	X	X	X	X	X	X
Unit & Week FE	X	X	X	X	X	X
State by Month FE		X	X	X	X	X
Unit Characteristic by Month FE			X	X	X	X
Scrubber & Enviro. Prog. FE				X	X	X
$R2$	0.599	0.621	0.65	0.651	0.667	0.661
# of Observations	1.834M	1.834M	1.834M	1.834M	1.491M	1.572M
# of Never-Treated Units	2308	2308	2308	2308	2308	2308
# of Treated Units	2046	2046	2046	2046	1089	1142

Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3, \text{and} \delta_4$ from estimating Equation (2) using net generation to calculate productivity. Unit characteristic fixed effects include installation year, fuel type, technology type, and unit capacity bins. Standard errors are clustered at the acquisition level.

Figure OA-18: Dynamic Effects of Acquisitions on Productivity (Daily)

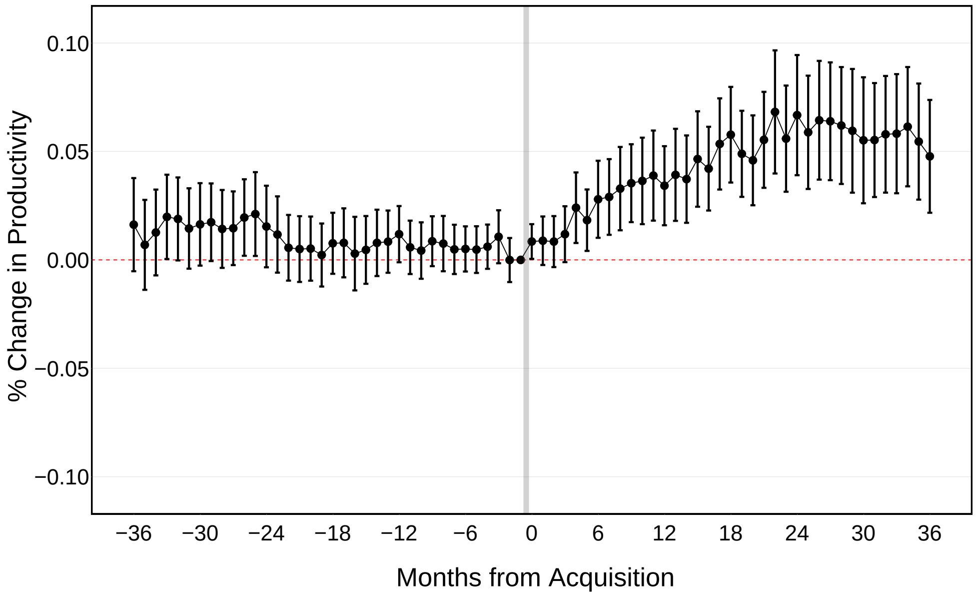

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors. The dependent variable is the logarithm of the inverse daily heat rate, as discussed in Section 7. Error bars show 95% confidence intervals. Standard errors are clustered at the acquisition level.

Figure OA-19: Dynamic Effects of Acquisitions on Productivity (Net Generation)

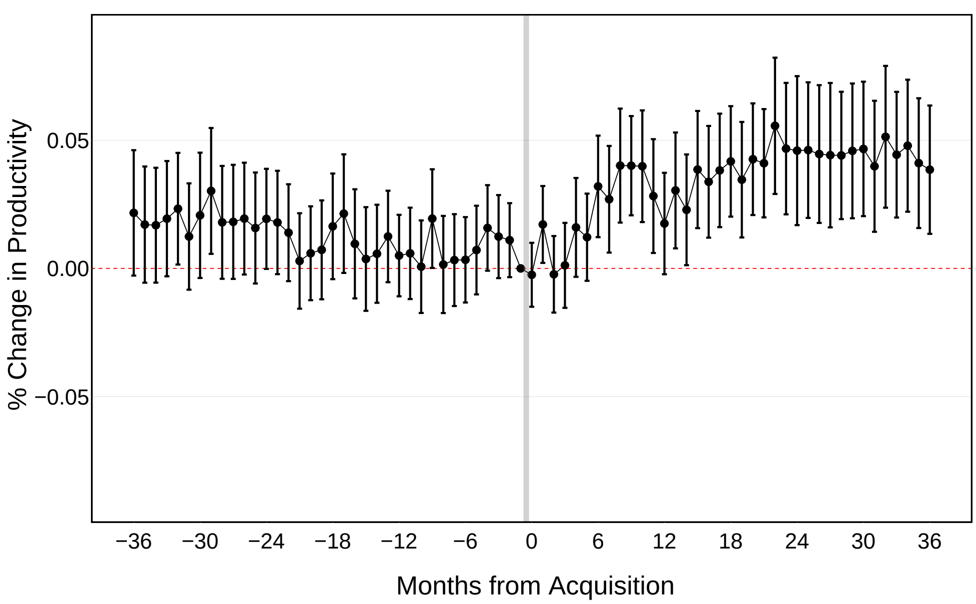

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors. The dependent variable is the logarithm of the inverse weekly heat rate, calculated using net generation as opposed to gross generation as discussed in Section 7. Error bars show 95% confidence intervals. Standard errors are clustered at the acquisition level.

Figure OA-20: Dynamic Effects of Acquisitions on Productivity (All Acquisitions)

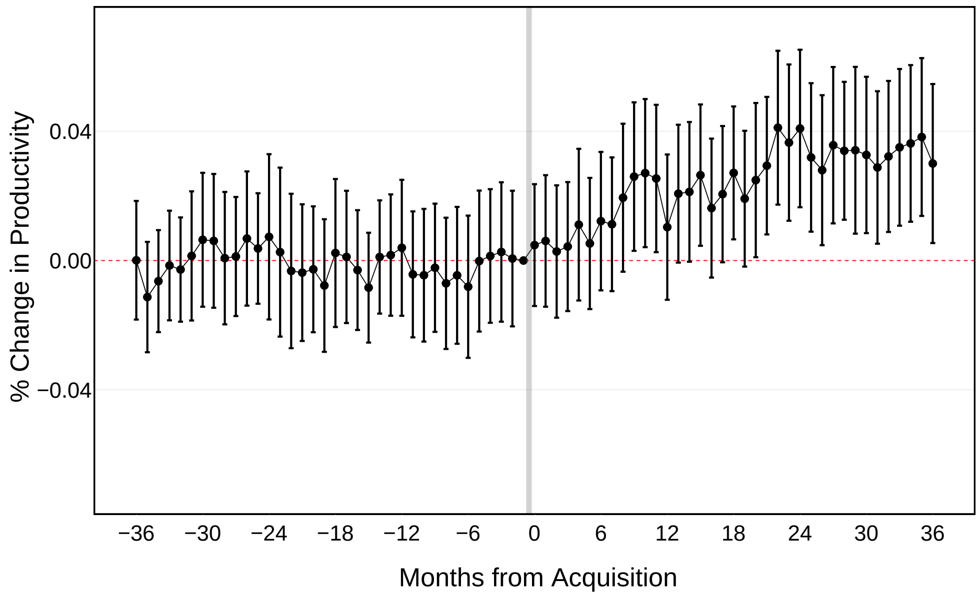

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors, using all acquisitions. The acquisition sample is described in Section C.1. The dependent variable is the logarithm of the inverse weekly heat rate. Error bars show 95% confidence intervals. Standard errors are clustered at the acquisition level.

Figure OA-21: Dynamic Effects of Acquisitions on Productivity (Staggered Difference-in-Differences)

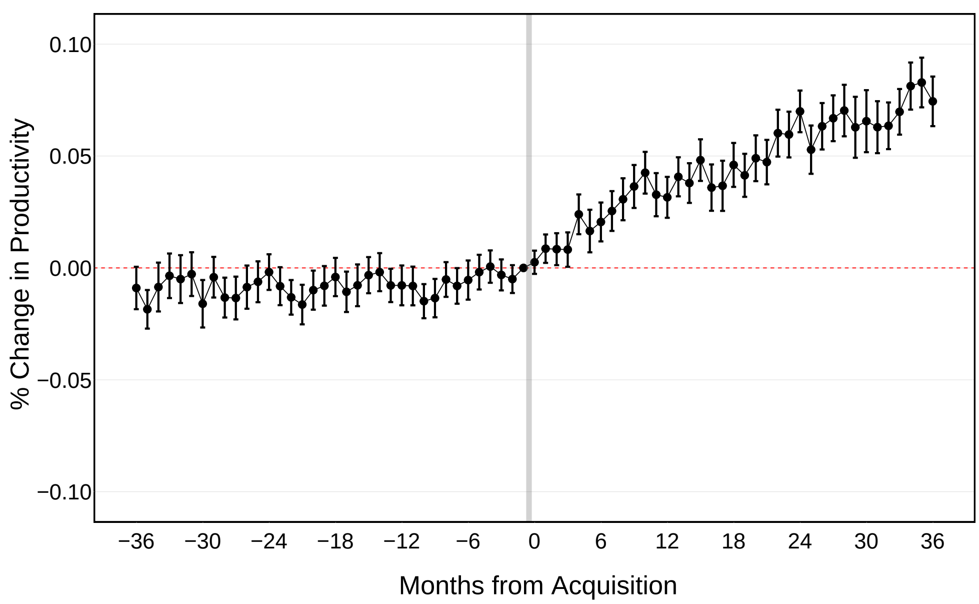

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors, using the method of Callaway and Sant'Anna (2021). The details are provided in Section C.3. The dependent variable is the logarithm of the inverse weekly heat rate. Standard errors are clustered at the acquisition level. This specification does not include unit characteristics and time trends due to computational complexity. Error bars show 95% confidence intervals.

Figure OA-22: Dynamic Effects of Acquisitions on Productivity (Matching Estimator)

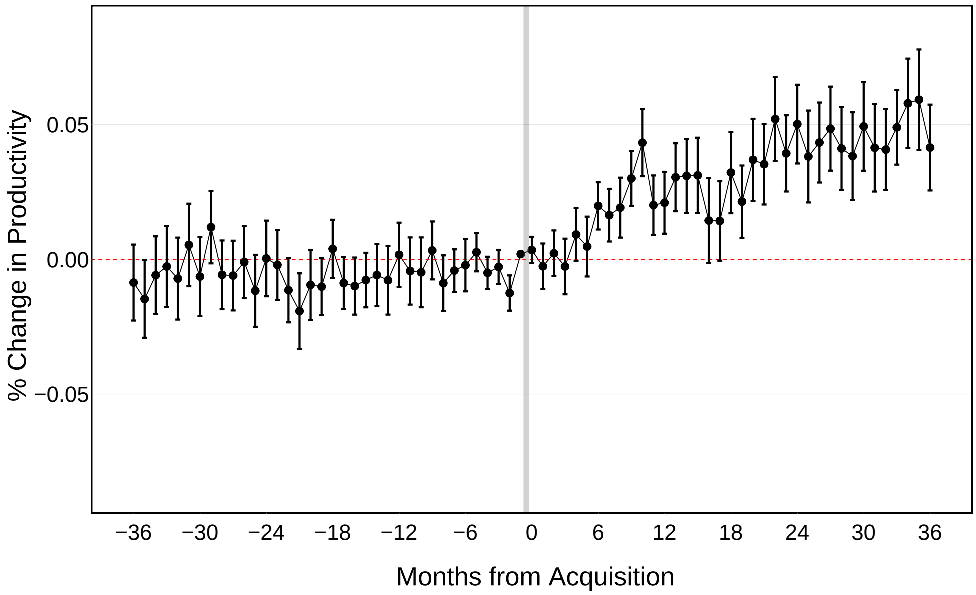

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors, using the matching method described in Section C.4. The dependent variable is the logarithm of the inverse weekly heat rate. Error bars show 95% confidence intervals.

Figure OA-23: Dynamic Effects of Acquisitions on Productivity (Weighted By Capacity)

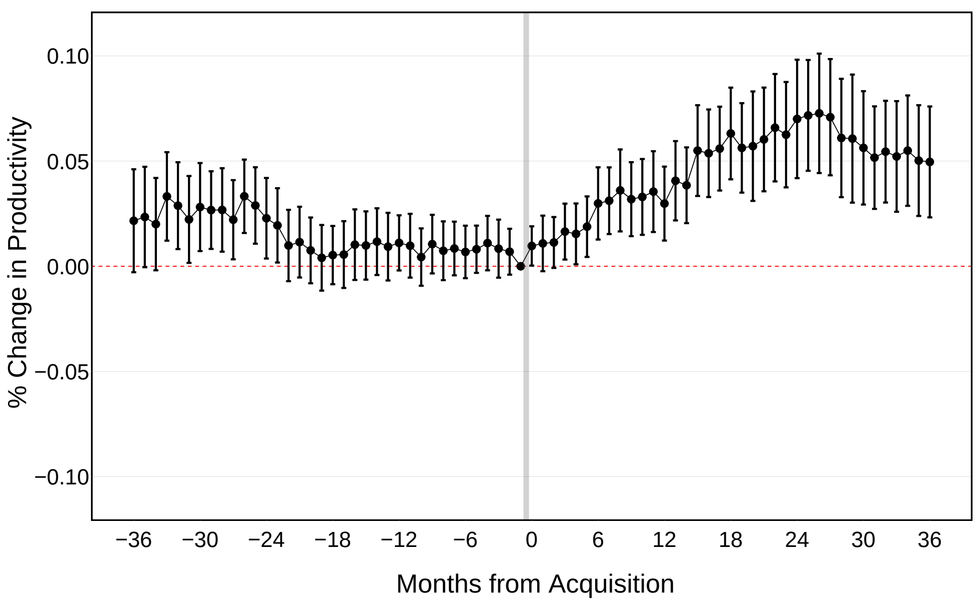

Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors, weighting observations by capacity as described in Section C.5. The dependent variable is the logarithm of the inverse weekly heat rate. Error bars show 95% confidence intervals. Standard errors are clustered at the acquisition level.