

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 heat rates above 16 or below 6 MMBtu per MWh. This winsorization affects less than 3% of all observations.

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.⁴⁸ 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.

Environmental 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 SO₂ Annual Group 1 Program; Cross-State Air Pollution SO₂ 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 SO₂ Trading Program.

Environmental Control Equipment CEMS provides data on environmental control equipment used in boilers for SO₂, NO_x, and particulate matter (PM) reduction. This includes 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 two weeks. Then, we obtain the annual hourly generation distribution and use the 99th percentile of the observed hourly generation

⁴⁸<https://www.epa.gov/airmarkets/power-sector-data-crosswalk>.

conditional on operating in the CEMS data every year. This algorithm yields generator capacity that is stable over time for most units. If a unit generates for less than two weeks over the course of a year, this algorithm does not yield any capacity estimates. 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. 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.⁴⁹ We use these data on plant representatives from the EPA between 2000 and 2020 to construct personnel data. 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.

⁴⁹<https://www.epa.gov/enviro/envirofacts-data-service-api>

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. We compiled a dataset focusing on divestiture-related acquisitions after 2000 to differentiate between mergers and deregulation-driven divestitures from various resources.

Our initial step involves data from [Cicala \(2022\)](#) and [Abito et al. \(2024\)](#), which provide a list of divested power plants from 2000 to 2002 from the EIA and other resources. For data after 2002, similar to [Cicala \(2022\)](#) and [Abito et al. \(2024\)](#), we primarily rely on the EIA Electric Power Monthly reports, which include a section listing divestitures, titled “Electric Utility Plants Sold/Transferred and Reclassified as Non-utility Plants.” These reports are available only until 2012. After 2012, we use the regulation status information of the owner in the GMI and Velocity Suite and consider the status change from “regulated” to “deregulated” as a change in the regulation status.

The next step in our methodology is to identify the ownership changes due to divestitures, as the exact dates of ownership changes and changes in the regulation status might not match in different datasets. For this, we use a 12-month window. If a regulation status change falls within a 12-month window of the acquisition, we consider that acquisition a divestiture. This procedure results in 615 acquisition events between 2000 and 2023 that we remove from the acquisition sample.

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, 2018](#); [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.⁵⁰

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 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

⁵⁰<https://www.spglobal.com/marketintelligence/en/solutions/sp-capital-iq-pro>.

previous owners and important transaction characteristics.

In the GMI transaction data, we observe that approximately 20% 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.

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.⁵¹ 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 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.

⁵¹<https://www.capitaliq.com/>.

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 scores⁵²; 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

⁵²If a unit does not experience an outage (derate), then the average score is equal to the derate (outage) score.

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 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 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.⁵³ 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 three years preceding the acquisition and 1-3 years following the acquisition. We exclude data within the first year

⁵³Generator model and characteristics are missing for about 20% of generators. For these, we include a missing dummy variable.

post-acquisition because efficiency improvements take time to materialize, and we aim to measure the long-term effects of 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 re-sample the treated generators with replacement and estimate the heat rate curve for the sample. We repeat this 200 times and report the 2.5 and 97.5 percentiles of the bootstrap distribution.

B.3 Calculation of CO₂ Emission Reductions Due to Acquisitions

To quantify the efficiency gains from acquisitions in terms of changes in CO₂ 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 CO₂ emissions at the unit level. We identify the month of the first acquisition for each unit and calculate CO₂ 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 CO₂ emissions before the acquisition and compare them to the total generation and implied CO₂ emissions after the acquisition to determine CO₂ intensity changes. Assuming no change in production post-acquisition, we calculate the hypothetical total emissions if the unit had maintained its pre-acquisition CO₂ emission intensity. The total CO₂ 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 CO₂ 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 CO₂ 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 500GB.

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 SantAnna \(2021\)](#) 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.⁵⁴ 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 three 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:

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

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

⁵⁴For 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-5, Table OA-7, and Figure OA-18. 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.

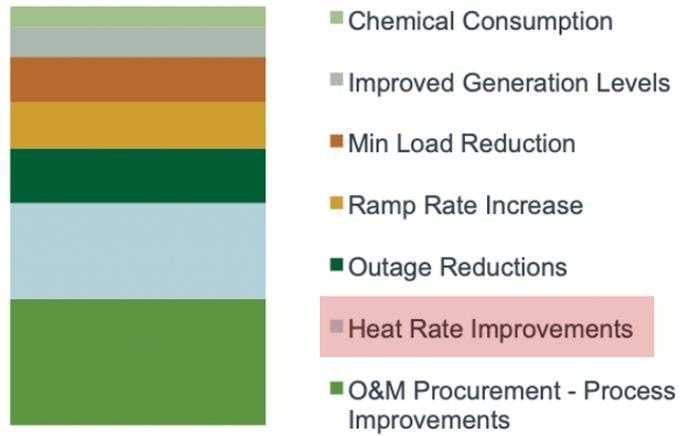
C.6 Placebo Test

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-5, confirm our expectation that we do not see any significant change in power plant efficiency after these events.

D Additional Figures

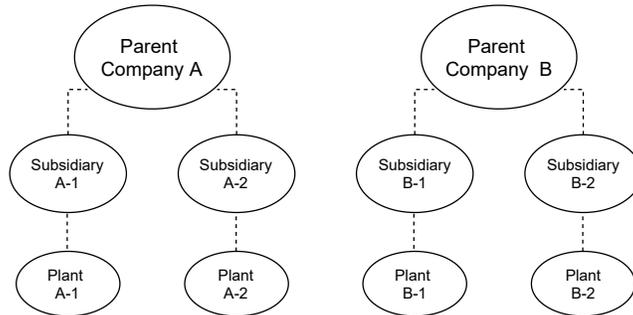
Figure OA-1: A Slide from Investor Presentation About Efficiency Claims

Projected Operational Improvements **\$125mm**

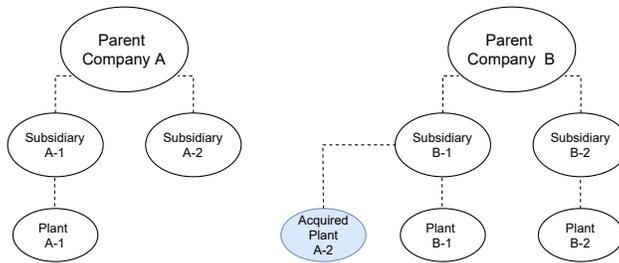


Note: This figure is from a slide deck presented in the conference call of the acquisition of Dynegy by Vistra Energy in 2018.

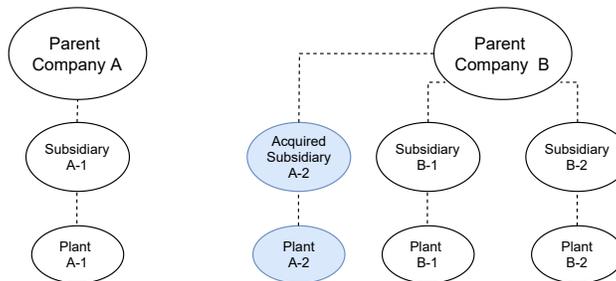
Figure OA-2: Illustration of Ownership Change Types



(a) Ownership Structure Before Acquisition



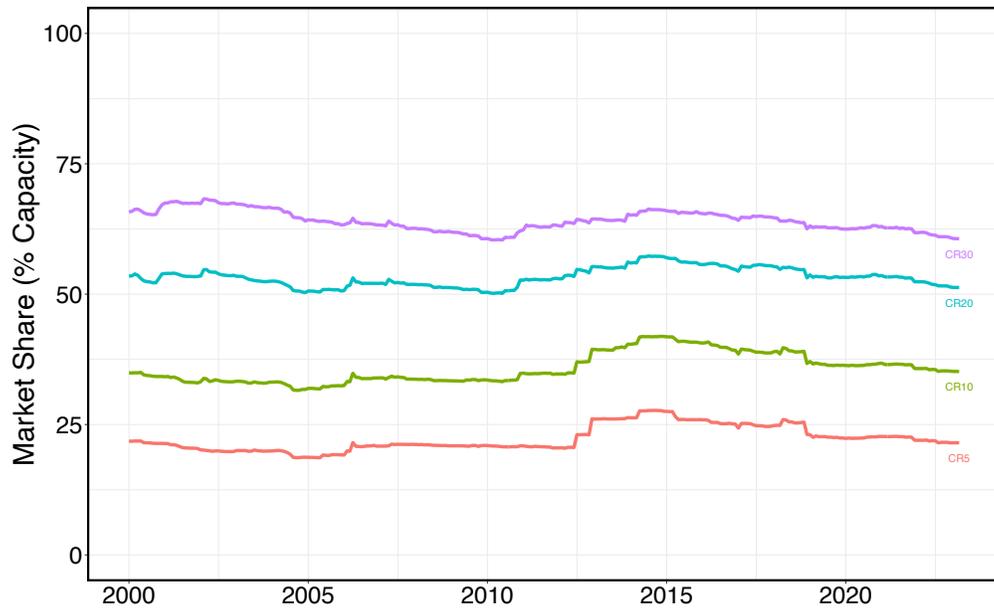
(b) Asset Acquisitions



(c) Subsidiary Acquisitions

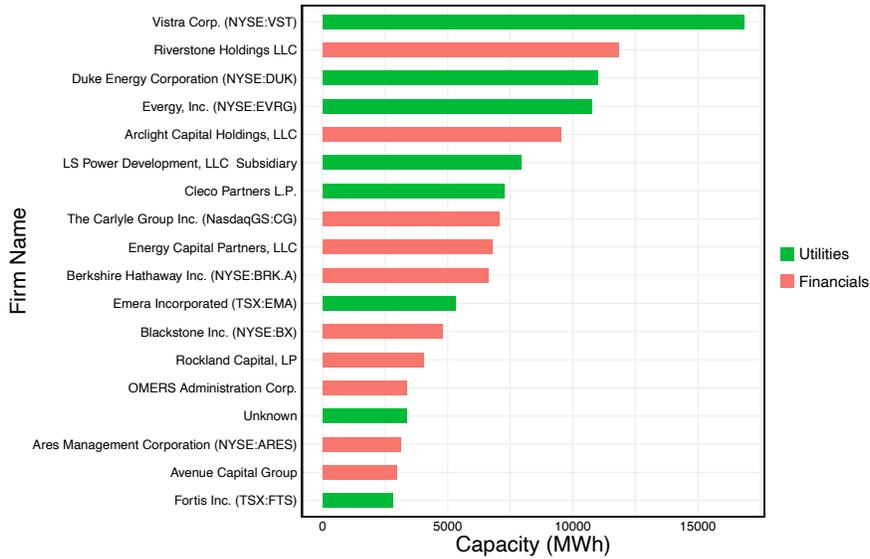
Note: This figure demonstrates different types of acquisitions. Panel (a) is the corporate structure of companies before the acquisition. Panels (b) and (c) show the corporate structure after the acquisition for asset sales and subsidiary acquisitions, respectively.

Figure OA-4: Change in US Power Generation Market Concentration by Capacity



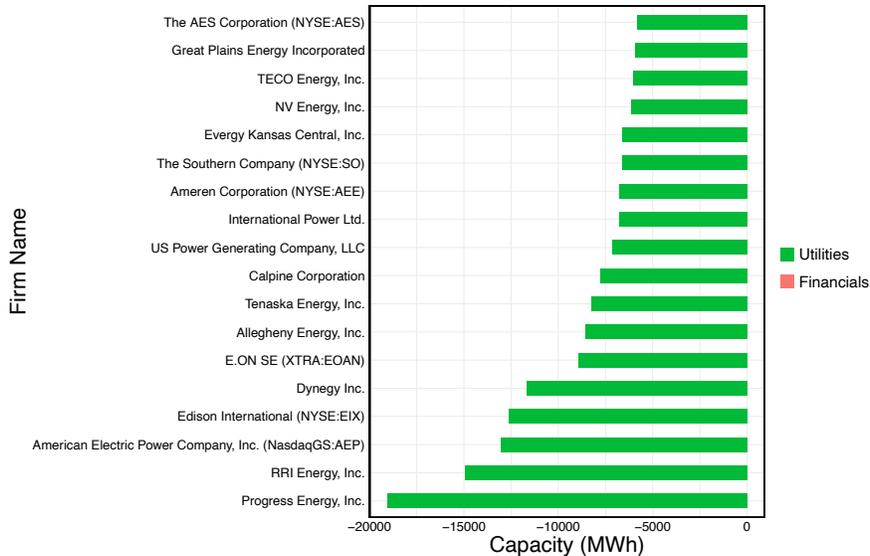
Note: This figure shows the change of national concentration ratios in the overall US fossil fuel power plant market between 2000 and 2023. For every concentration ratio, we calculate the total fossil fuel capacity of the largest corresponding number of firms in the US and divide that by the total fossil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

Figure OA-5: Firms with Largest Capacity Increase, 2010–2023



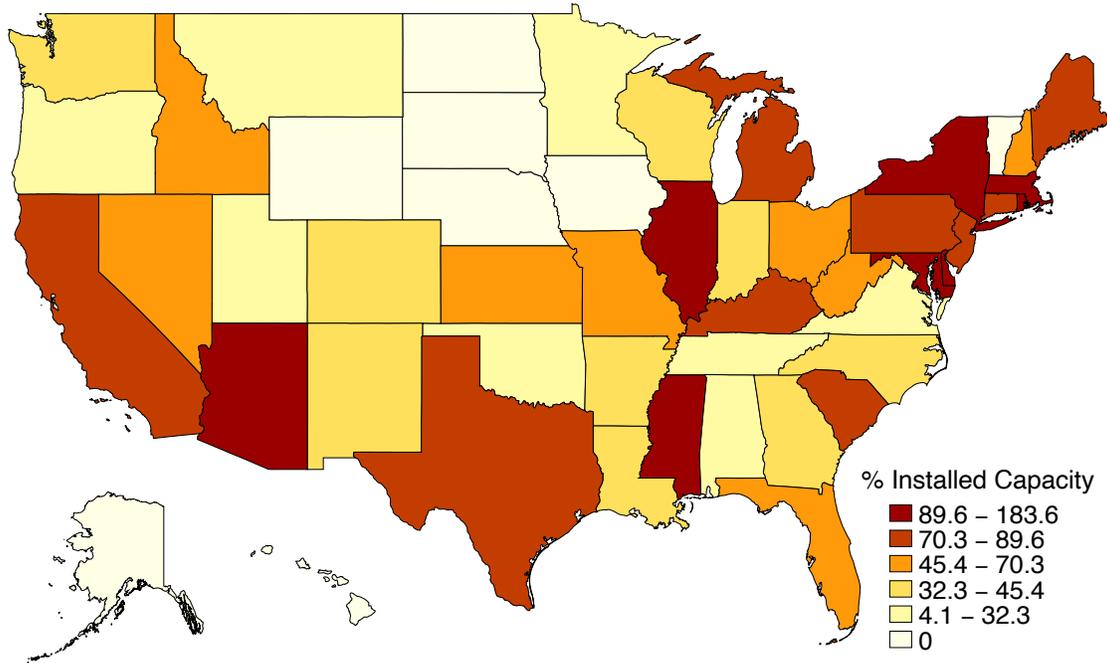
Note: This figure shows firms with the largest capacity increase in fossil fuel generation capacity in the US between 2010 and 2020.

Figure OA-6: Firms with Largest Capacity Decrease, 2010–2023



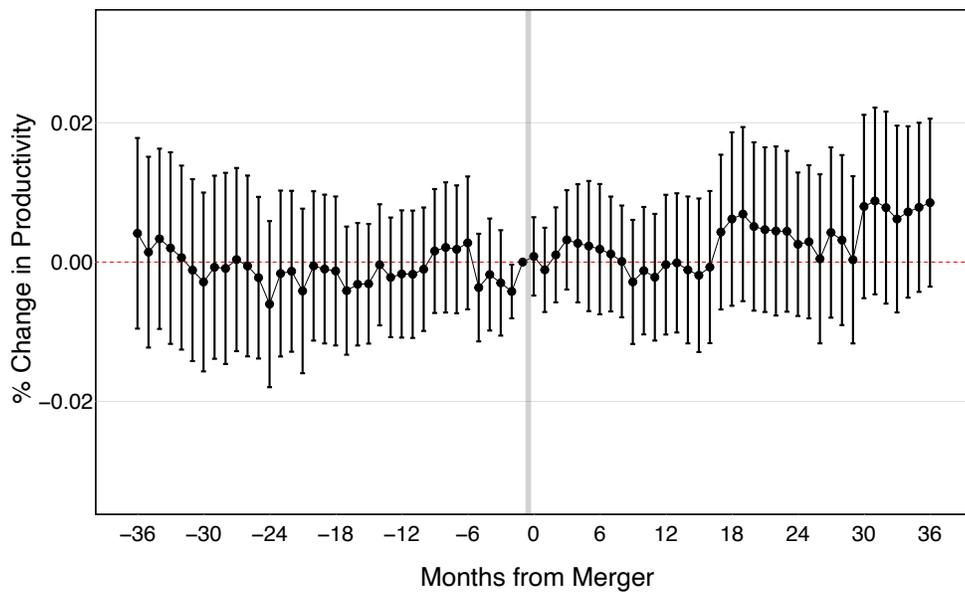
Note: This figure shows firms with the largest capacity decrease in fossil fuel generation capacity in the US between 2010 and 2020.

Figure OA-7: Geographic Distribution of Ownership Changes by Capacity



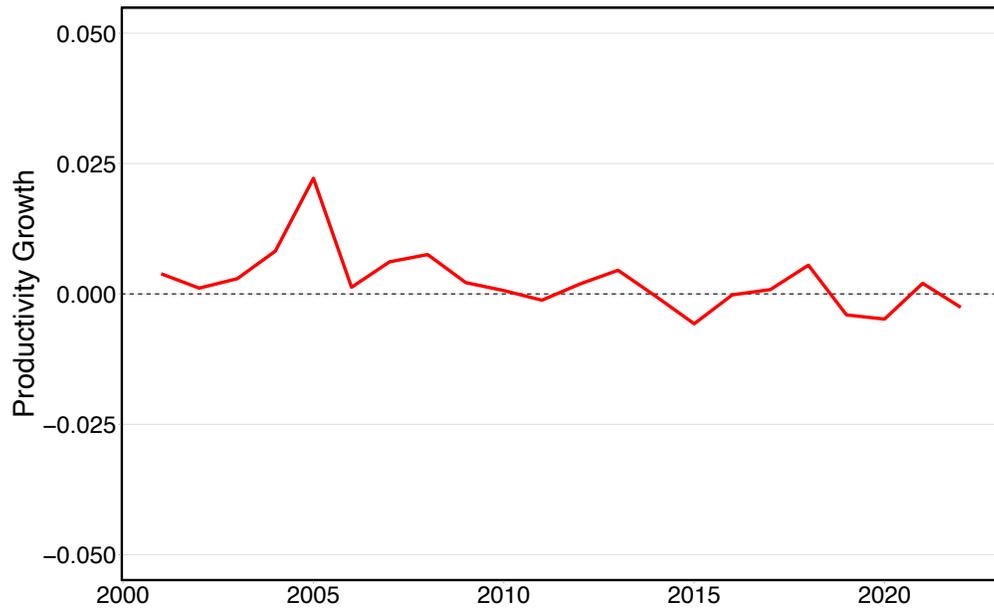
Note: Geographical distribution of power plant acquisitions by capacity between January 2000 - March 2023.

Figure OA-8: The Effects of Manager Changes without Acquisitions on Efficiency



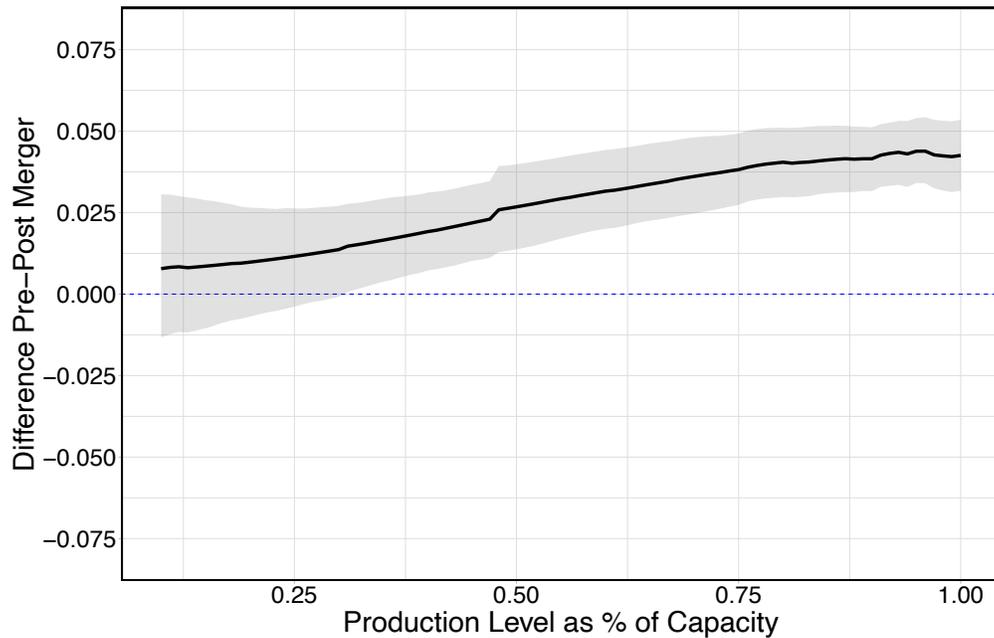
Note: This figure shows the effects of manager change on efficiency estimated using the specification given in Equation (3). In particular, we treat a unit if the manager of that unit changes and there is no acquisition in the three months preceding and following the manager change. Error bars show 95%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

Figure OA-9: Average Within-Plant Annual Productivity Change



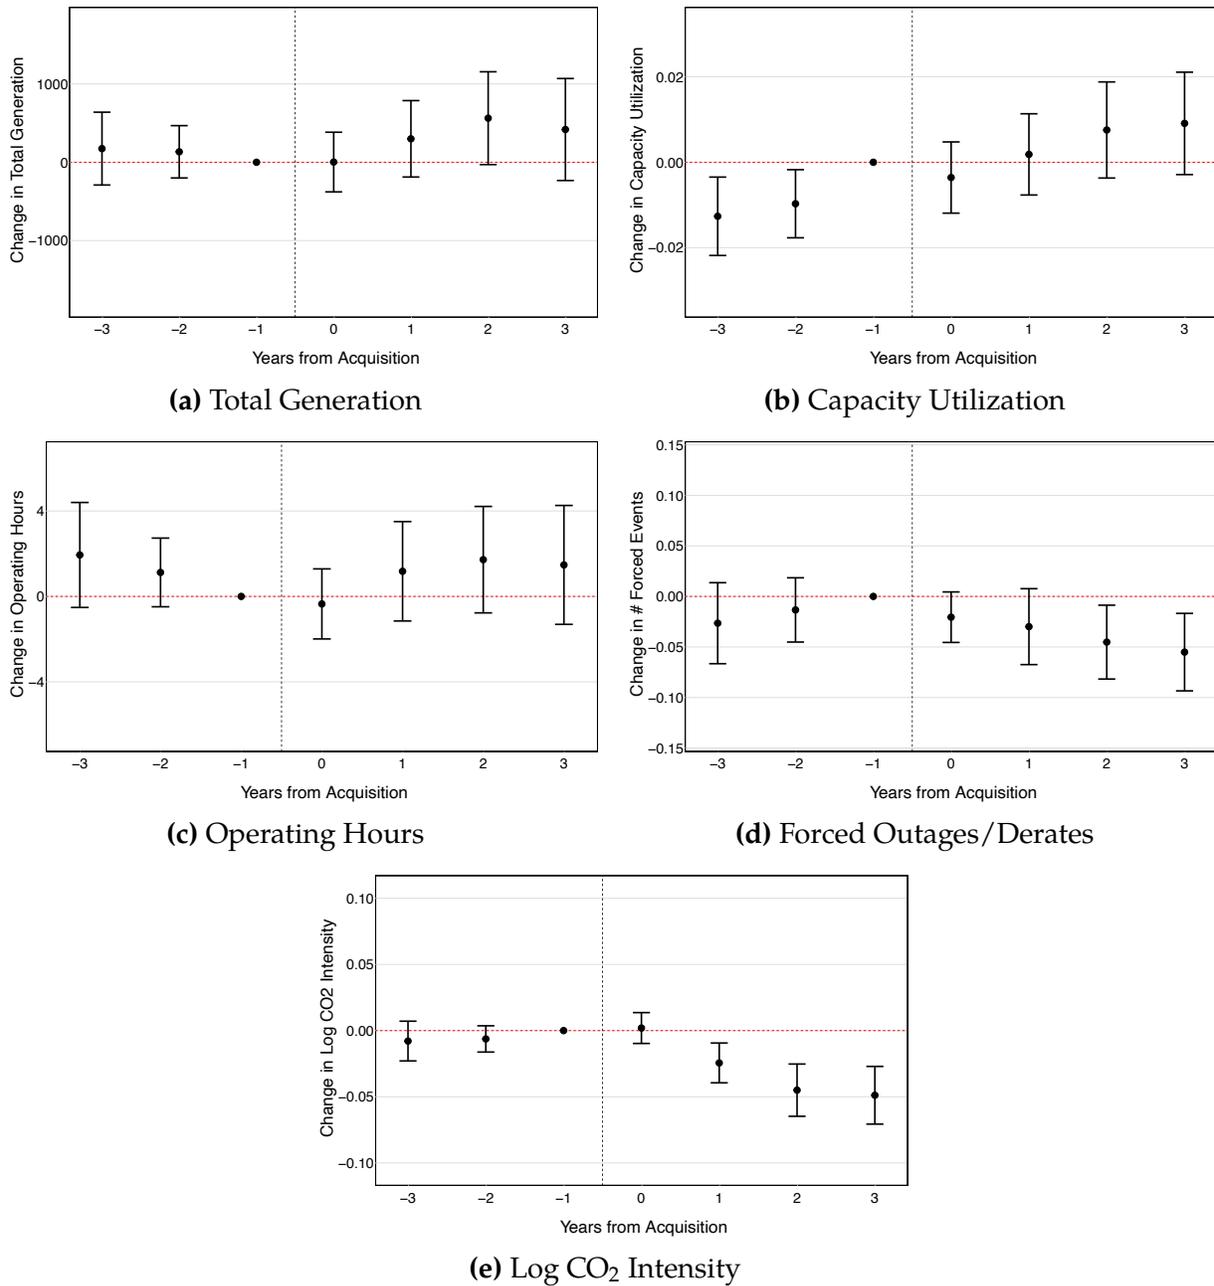
Note: This figure shows average year-to-year within-plant productivity growth for the plants that were not involved in an acquisition.

Figure OA-10: Confidence Band for the Difference in Heat Rate Curves



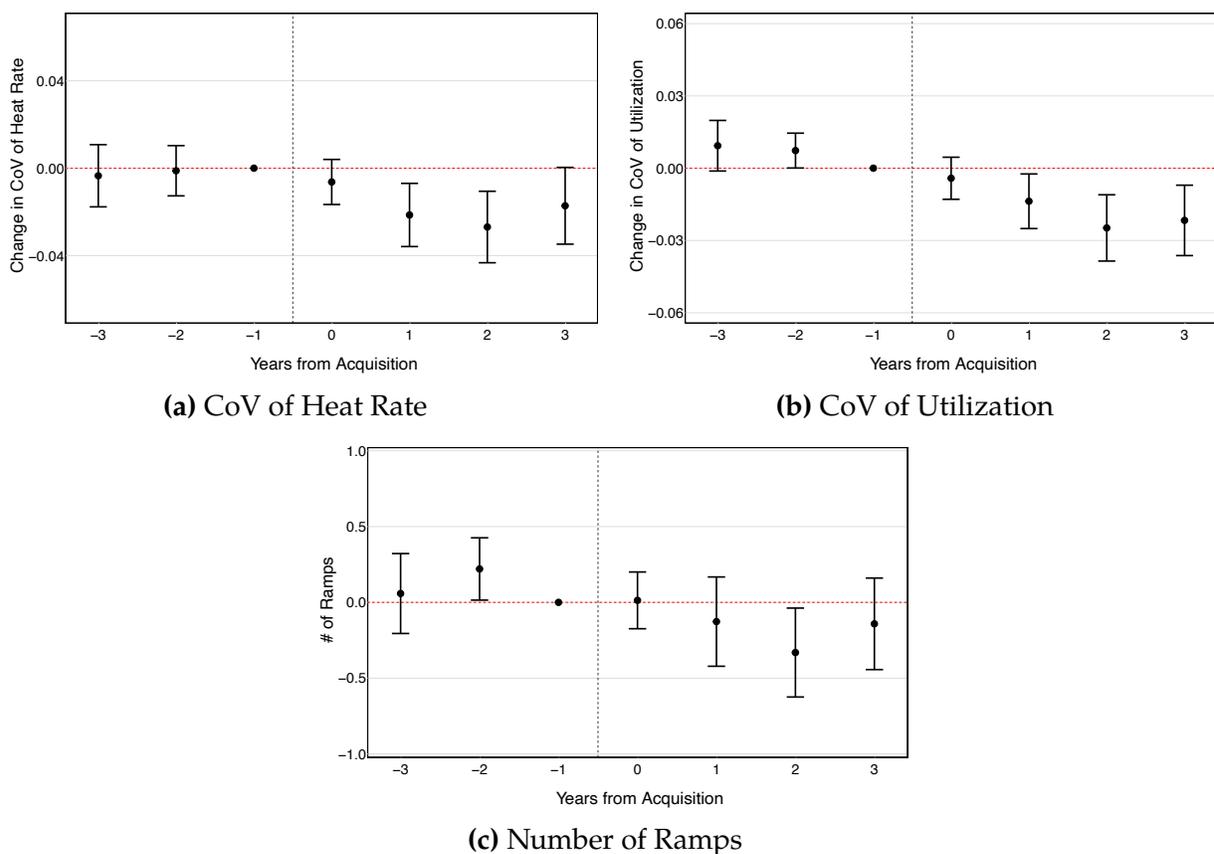
Note: This figure shows the 95% confidence interval for the difference between acquired firms' heat rate curves pre- and post-acquisition, as reported in Figure 7(a). The estimates are reported from 200 bootstrap replications.

Figure OA-11: Effects of Acquisitions on Generator Performance Measures (Event Studies)



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 plant level

Figure OA-12: Results on Dynamic Efficiency Mechanism (Event Studies)



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.

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: 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)
<i>Dependent Variable: Log of Efficiency</i>				
Pre-acquisition (1 Year)	0.004 (0.005)	0.003 (0.005)	-0.001 (0.005)	-0.001 (0.005)
Post-acquisition (1 Year)	0.03 (0.008)	0.026 (0.008)	0.016 (0.007)	0.016 (0.007)
Post-acquisition (2 Years)	0.056 (0.011)	0.054 (0.01)	0.039 (0.009)	0.039 (0.009)
Post-acquisition (3 Years)	0.069 (0.013)	0.071 (0.013)	0.05 (0.011)	0.05 (0.011)
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 plant level.

Table OA-3: 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)
<i>Dependent Variable: Log of Efficiency</i>					
Pre-acquisition (1 Year)	-0.001 (0.005)	-0.001 (0.005)	-0.002 (0.005)	-0.001 (0.005)	-0.001 (0.005)
Post-acquisition (1 Year)	0.002 (0.009)	0.016 (0.01)	0.009 (0.008)	0.011 (0.008)	0.015 (0.009)
Post-acquisition (2 Years)	0.018 (0.01)	0.038 (0.013)	0.014 (0.009)	0.019 (0.01)	0.049 (0.013)
Post-acquisition (3 Years)	0.029 (0.013)	0.054 (0.017)	0.026 (0.013)	0.034 (0.013)	0.068 (0.016)
Post-acquisition (1 Year) × Z	0.023 (0.011)	-0.001 (0.012)	0.014 (0.012)	0.012 (0.012)	0.002 (0.012)
Post-acquisition (2 Years) × Z	0.035 (0.015)	0.004 (0.016)	0.059 (0.016)	0.049 (0.017)	-0.021 (0.015)
Post-acquisition (3 Years) × Z	0.034 (0.018)	-0.011 (0.02)	0.058 (0.02)	0.041 (0.02)	-0.039 (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 plant level. Details about the heterogeneity variables are provided in Appendix B.4.

Table OA-4: Effects of Acquisitions with or without Manager Changes

	Acquisitions w/ Manager Changes (1)	Acquisitions w/o Manager Changes (2)
<i>Dependent Variable: Log of Efficiency</i>		
Pre-acquisition (1 Year)	-0.008 (0.005)	0.013 (0.009)
Post-acquisition (1 Year)	0.011 (0.008)	0.025 (0.012)
Post-acquisition (2 Years)	0.038 (0.01)	0.039 (0.016)
Post-acquisition (3 Years)	0.052 (0.013)	0.057 (0.021)
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 plant 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.

F Robustness Checks Results

Table OA-5: Effects of Acquisitions on Generator Productivity (Robustness)

	After 2010 (1)	Weighted Regressions (2)	Net Generation (3)	All Acquisitions (4)	Matching (5)	Minority (Placebo) (6)
<i>Dependent Variable: Log of Efficiency</i>						
Pre-acquisition (1 Year)	-0.002 (0.004)	-0.001 (0.005)	-0.004 (0.005)	-0.001 (0.005)	-0.003 (0.002)	-0.001 (0.006)
Post-acquisition (1 Year)	0.002 (0.006)	0.016 (0.007)	0.009 (0.006)	0.013 (0.006)	0.021 (0.004)	-0.007 (0.009)
Post-acquisition (2 Years)	0.021 (0.011)	0.039 (0.009)	0.028 (0.008)	0.027 (0.007)	0.036 (0.006)	-0.023 (0.013)
Post-acquisition (3 Years)	0.037 (0.014)	0.05 (0.011)	0.033 (0.01)	0.033 (0.009)	0.047 (0.007)	-0.013 (0.011)
Ambient Temp. & Humidity 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	1.769M	1.493M	-	1.407M
# of Never-Treated Units	2311	2311	2311	2311	-	2311
# of Treated Units	529	1089	1089	1541	-	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 plant level.

Table OA-6: 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)
<i>Dependent Variable: Log of Efficiency</i>						
Pre-acquisition (1 Year)	0 (0.004)	0 (0.004)	-0.002 (0.004)	-0.002 (0.004)	-0.001 (0.005)	-0.015 (0.005)
Post-acquisition (1 Year)	0.01 (0.004)	0.01 (0.004)	0.006 (0.004)	0.005 (0.004)	0.013 (0.006)	-0.015 (0.006)
Post-acquisition (2 Years)	0.019 (0.005)	0.02 (0.005)	0.012 (0.004)	0.012 (0.004)	0.027 (0.007)	-0.01 (0.006)
Post-acquisition (3 Years)	0.022 (0.006)	0.022 (0.005)	0.011 (0.005)	0.011 (0.005)	0.033 (0.009)	-0.016 (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	0.713	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 δ_1 , δ_2 , δ_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 plant level.

Table OA-7: 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)
<i>Dependent Variable: Log of Efficiency</i>						
Pre-acquisition (1 Year)	-0.003 (0.004)	-0.006 (0.004)	-0.008 (0.003)	-0.008 (0.003)	-0.001 (0.005)	-0.014 (0.004)
Post-acquisition (1 Year)	0.013 (0.006)	0.009 (0.006)	0.001 (0.005)	0.001 (0.005)	0.016 (0.007)	-0.016 (0.007)
Post-acquisition (2 Years)	0.03 (0.008)	0.028 (0.008)	0.015 (0.007)	0.015 (0.007)	0.039 (0.009)	-0.009 (0.008)
Post-acquisition (3 Years)	0.034 (0.009)	0.031 (0.01)	0.015 (0.009)	0.015 (0.009)	0.05 (0.011)	-0.014 (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	0.707	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) 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 plant level.

Table OA-8: 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)
<i>Dependent Variable: Log of Efficiency</i>						
Pre-acquisition (1 Year)	0.002 (0.003)	-0.001 (0.003)	-0.005 (0.003)	-0.006 (0.003)	-0.002 (0.004)	-0.008 (0.004)
Post-acquisition (1 Year)	0.006 (0.006)	0.004 (0.006)	-0.006 (0.006)	-0.006 (0.006)	0.002 (0.006)	-0.014 (0.007)
Post-acquisition (2 Years)	0.015 (0.009)	0.017 (0.009)	0.004 (0.008)	0.004 (0.008)	0.021 (0.011)	-0.01 (0.008)
Post-acquisition (3 Years)	0.021 (0.011)	0.023 (0.012)	0.01 (0.01)	0.01 (0.01)	0.037 (0.014)	-0.013 (0.011)
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

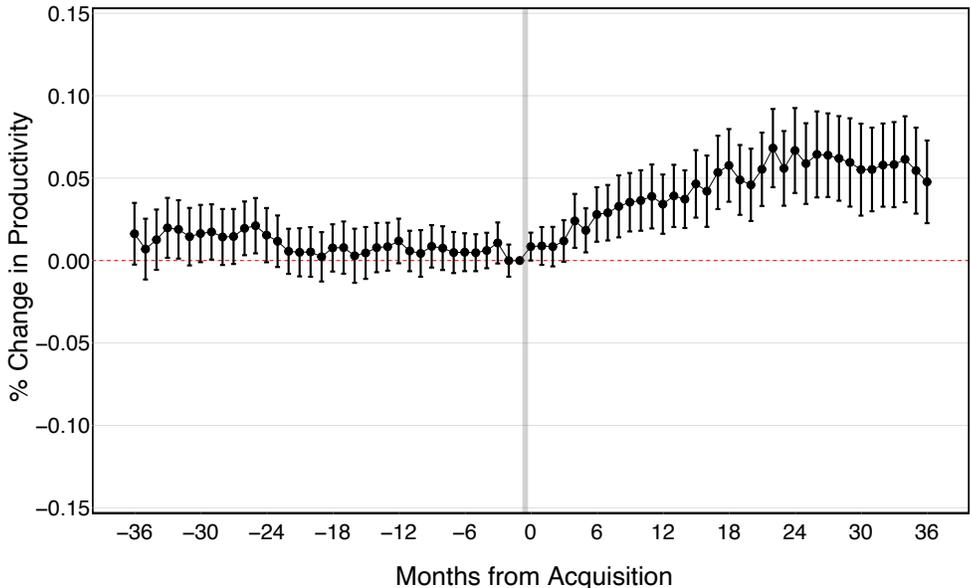
Note: This table presents the coefficient estimates of δ_1 , δ_2 , δ_3 , and δ_4 from estimating Equation (2) for acquisitions occurring after 2010. Unit characteristic fixed effects include installation year, fuel type, technology type, and unit capacity bins. Standard errors are clustered at the plant level.

Table OA-9: 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)
<i>Dependent Variable: Log of Efficiency</i>						
Pre-acquisition (1 Year)	-0.004 (0.004)	-0.005 (0.004)	-0.008 (0.004)	-0.008 (0.004)	-0.004 (0.005)	-0.009 (0.005)
Post-acquisition (1 Year)	0.003 (0.006)	0.004 (0.006)	-0.005 (0.005)	-0.005 (0.005)	0.009 (0.006)	-0.016 (0.007)
Post-acquisition (2 Years)	0.018 (0.008)	0.022 (0.007)	0.008 (0.006)	0.008 (0.006)	0.028 (0.008)	-0.007 (0.008)
Post-acquisition (3 Years)	0.023 (0.009)	0.024 (0.008)	0.008 (0.007)	0.008 (0.007)	0.033 (0.01)	-0.009 (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
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

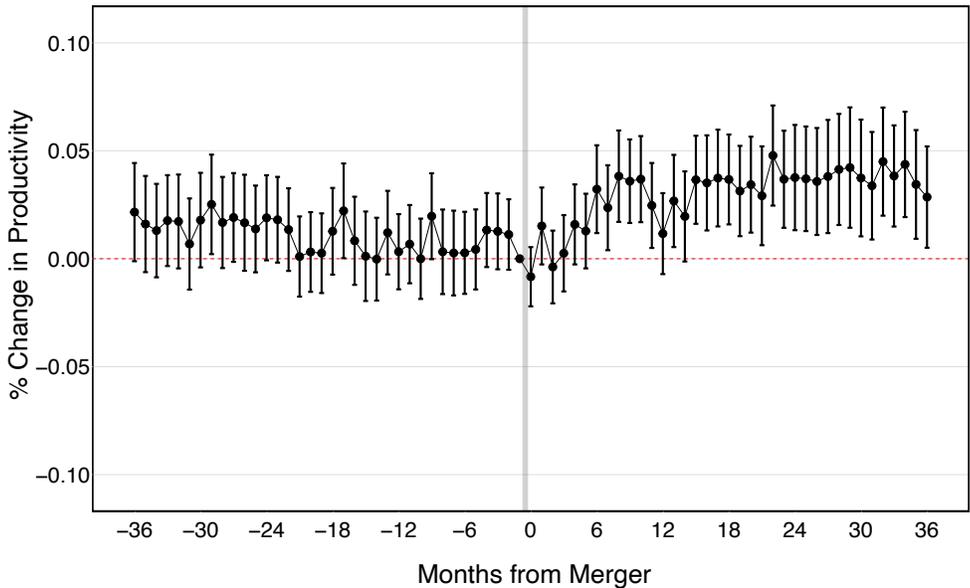
Note: This table presents the coefficient estimates of $\delta_1, \delta_2, \delta_3,$ and δ_4 from estimating Equation (2) using net generation to calculate productivity. Unit characteristic fixed effects include installation year, fuel type, technology type, and unit capacity bins. Standard errors are clustered at the plant level.

Figure OA-13: Dynamic Effects of Acquisitions on Productivity (Daily)



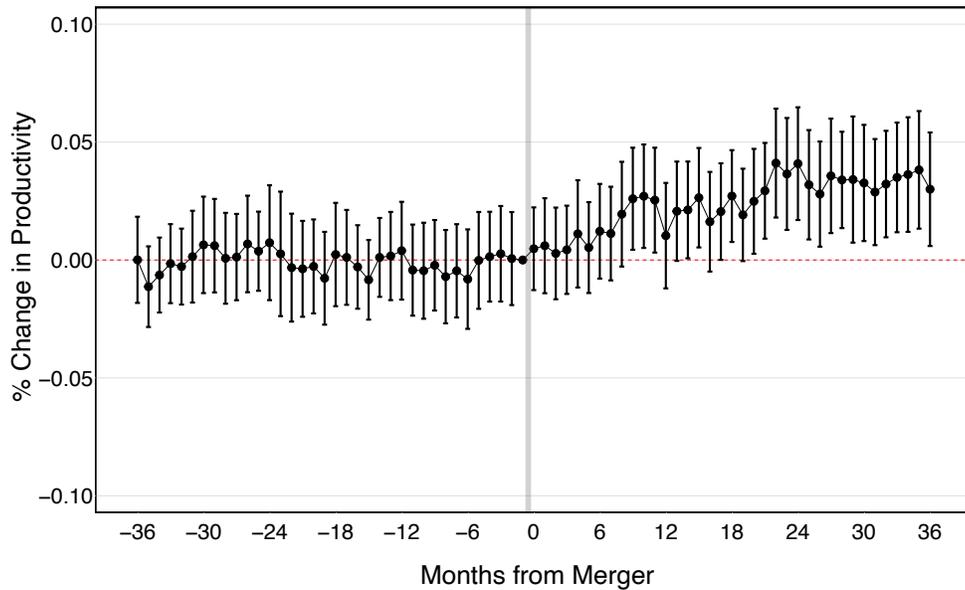
Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors. The dependent variable is the logarithm of the inverse daily heat rate, as discussed in Section 7. Error bars show 95% confidence intervals. Standard errors are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

Figure OA-14: Dynamic Effects of Acquisitions on Productivity (Net Generation)



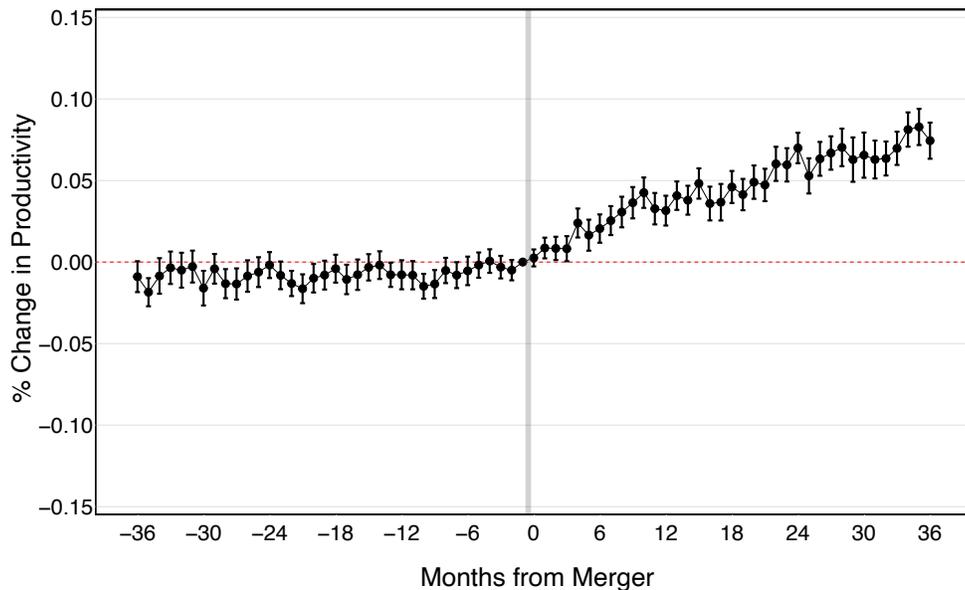
Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors. The dependent variable is the logarithm of the inverse weekly heat rate, calculated using net generation as opposed to gross generation as discussed in Section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

Figure OA-15: Dynamic Effects of Acquisitions on Productivity (All Acquisitions)



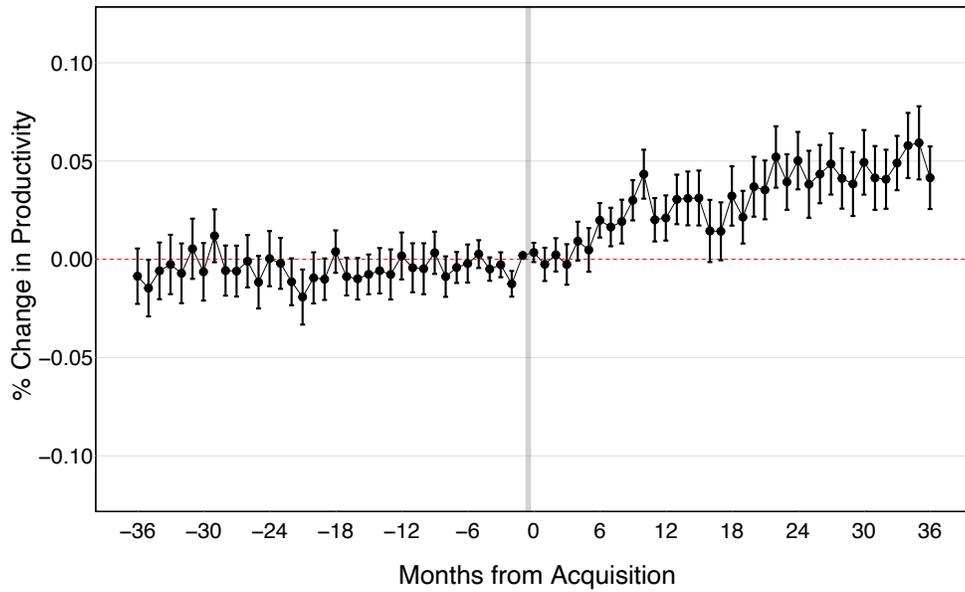
Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors, using all acquisitions. The acquisition sample is described in Section C.1. The dependent variable is the logarithm of the inverse weekly heat rate. Error bars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

Figure OA-16: Dynamic Effects of Acquisitions on Productivity (Staggered Difference-in-Differences)



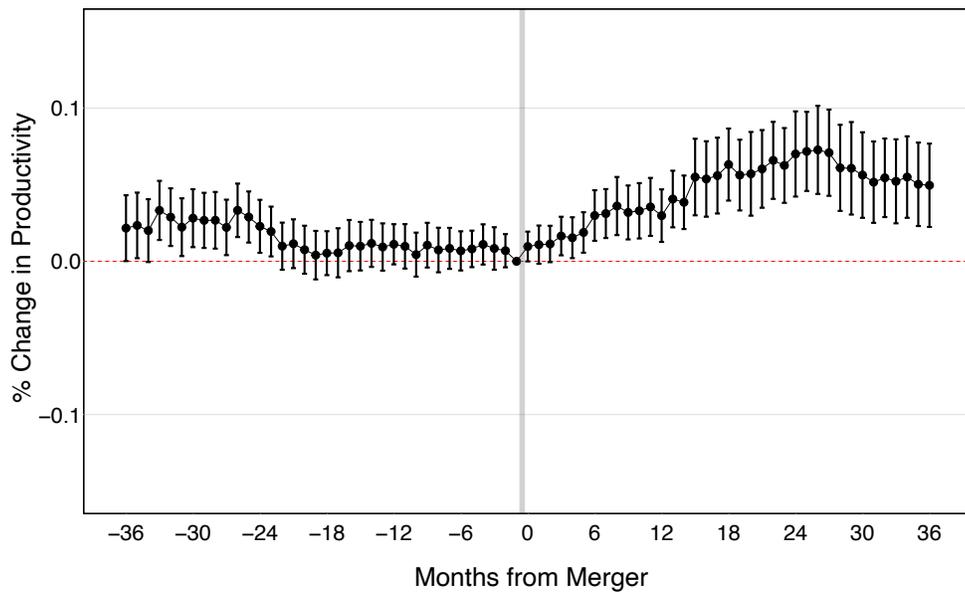
Note: This figure presents the coefficient estimates of $\hat{\delta}_s$ where $s \in (-36, 36)$ from Equation (3) along with standard errors, using the method of Callaway and Sant'Anna (2021). The details are provided in Section C.3. The dependent variable is the logarithm of the inverse weekly heat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

Figure OA-17: 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-18: 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 plant level.