

BATTERY STORAGE AND NATURAL GAS GENERATOR MARKET POWER

Maghfira Ramadhani*
Job Market Paper

October 27, 2026

Click [here](#) for most recent version.

Abstract

Battery storage is usually valued for integrating renewable energy, but it also competes directly with the natural gas generators that set prices and exercise market power during the evening, when electricity demand peaks. I study this competitive effect in Texas, exploiting the staggered expansion of storage capacity across counties. I find that storage lowers the evening markups of nearby gas generators. An additional 100 megawatt-hours of storage reduces gas operating profits by 6 percent, and at least 73 percent of this loss is market power rent. Gas generators keep setting prices, but at smaller margins. Storage also displaces evening gas output, especially from higher-cost generators. The resulting \$47 million annual transfer flows mainly to consumers, while the production cost savings from shifting energy across hours are small.

*Maghfira Ramadhani: School of Economics, Georgia Institute of Technology, maghfira.ramadhani@gatech.edu. I thank Dylan Brewer, Bobby Harris, Casey Wichman, Laura Taylor, Gaurav Doshi, Severin Borenstein, Matt Oliver, and audience at the Energy and Environmental Lunch at Georgia Tech, the Energy Camp 2025 (UC Berkeley) and Southeastern Workshop on Energy & Environmental Economics & Policy 2026 for their valuable feedback. This research was partially supported by the Georgia Tech EPICenter through its 2025 Summer Research Program.

1 Introduction

Battery storage is usually valued as a clean energy technology, but it also brings new competition into electricity markets. A battery storage operator (henceforth, storage) buys electricity (charges) when it is cheap and sells it back (discharges) when it is expensive. Economists have mostly studied how this helps integrate intermittent renewable energy: storage smooths renewable output, reduces curtailment, and relieves transmission congestion (Karaduman, 2023; Butters, Dorsey and Gowrisankaran, 2025; Andres-Cerezo and Fabra, 2025; Kirkpatrick, 2026; Dudek, 2026). This view overlooks that storage competes directly with natural gas generators during the evening, when electricity demand peaks. In these hours, gas generators are typically the marginal generators that set the price, and they exercise market power by withholding output to raise their markups (Borenstein, Bushnell and Wolak, 2002; Woerman, 2023). Storage entry displaces part of their output, which can lower production costs, and makes withholding less profitable, which lowers their markups and mainly transfers surplus to consumers. Economists have developed theoretical and structural models of how storage entry affects market outcomes, including in markets where firms hold market power (Andrés-Cerezo and Fabra, 2023; Karaduman, 2023; Butters, Dorsey and Gowrisankaran, 2025). Yet there is little ex post evidence on how incumbent gas generators change their output and markups once storage actually enters.

In this paper, I provide, to my knowledge, the first ex post evidence on how storage entry changes the output and realized markups of incumbent gas generators, and on how much of their lost profit comes from market power rent rather than a lower competitive price. I study the Texas electricity market, where more than 12 GWh of storage capacity has been deployed since 2021, a scale of market penetration substantially larger than those evaluated in prior studies.¹ The Texas grid also experiences frequent transmission congestion, which splits the grid into local markets where a few gas generators can exercise market power. My analysis proceeds in three stages. First, I estimate how storage entry changes the markups and operating profits of nearby gas generators. Second, I estimate how much gas output storage displaces, and which generators it displaces most. I then combine these estimates into a welfare accounting that values the surplus transferred through lower markups, production cost savings, and avoided emissions from displaced gas output.

To guide the empirical analysis, I develop a simple two-period model of a nodal market in which a marginal natural gas generator competes with a price-taking storage fleet in a uniform-price auction. The model follows Doshi (2025) and Woerman (2023), in which natural gas generators bid on the price and quantity of electricity. The model yields two channels through which storage affects the generator’s incentive to set its markup. Storage changes the *level* of demand left for the marginal generator, lowering residual demand when storage discharges and raising it when storage charges. Because storage discharges more when the price is higher, it also makes the *slope* of residual demand flatter, which raises the cost of withholding output for the marginal generator. The model predicts that storage lowers gas markups through both channels during discharge hours,

¹This capacity is roughly ten to one hundred times the storage analyzed in Butters, Dorsey and Gowrisankaran (2025), Karaduman (2023), and Kirkpatrick (2026).

while its effect during charge hours is ambiguous.

My empirical strategy exploits the staggered entry of storage capacity across Texas counties. I construct a generator-level panel that links each gas generator’s output, price offers, and marginal cost to storage entry within its county.² I define the local markets that congestion creates in each interval, using nodal price divergence across locations and the structure of the grid (Arkolakis and Walsh, 2024; Mercadal, 2022; Woerman, 2023). A storage unit in the same county as a gas generator shares its local market more than 99 percent of the time, so I approximate each generator’s exposure to storage competition at the county level. I compare the markups of generators in counties that storage has commercially operated in with generators in counties that storage has not yet reached, before and after entry, using counties’ installed storage capacity as a continuous treatment. Because the model’s predictions differ between charge and discharge hours, I estimate the effect of storage capacity separately for each hour of the day. Using data from the bid that generators submit, I also compute the slope of each gas generator’s residual demand from the offers of its rivals within the local market.

Storage lowers the markups of incumbent gas generators during the evening peak. An additional 100 MWh of local storage lowers the markup by \$2.71 per MWh at 7 PM and by \$1.50 per MWh on average from 5 PM to 9 PM, about 5 percent of the pre-entry evening markup. Markups do not change overnight or in the morning, when storage charges. The lower markups reduce daily operating profit by \$5.45 per MW-day, about 6 percent of the pre-entry average. The event study estimates show no pre-entry trends, and markups fall at every level of installed storage capacity. Consistent with the model’s prediction, the residual demand of nearby gas generators becomes flatter after storage enters, which raises their cost of withholding output. The flattening is most pronounced for combined cycle generators in the evening, when baseline market power is greater.

The loss in operating profit comes mainly from lower market power rent rather than a lower competitive price. A fall in profit alone does not necessarily mean that storage reduces market power, because storage could also lower the competitive price that gas generators would earn without markups. Following Borenstein, Bushnell and Wolak (2002), I decompose each generator’s operating profit into a competitive rent, earned at the price that cost-based dispatch would produce within its local market, and a market power rent above that price. I bound the competitive price with two benchmarks: the cost-based dispatch price and the cost of the generator that the system operator identifies as the price-setter. Market power rent accounts for between 73 percent and all of the lost profit under the two benchmarks. If storage lowered gas profits only by adding supply, it would push cheaper units along the merit order to the margin. Instead, gas generators continue setting the price after storage enters, but at a smaller margin above cost.

The theoretical model shows that the same two channels also affect the output of gas generators. In discharge hours, the level channel lowers gas output and the slope channel raises it, so output falls if storage serves more demand than the generators stop withholding. In charge hours, both channels

²I define the entry date of storage capacity as the first day the storage unit is awarded a positive quantity in the real-time market.

raise gas output, although the increase is small when storage charges from surplus renewable output rather than from gas. The slope channel also determines which generators lose output. A higher-cost generator must offer at a higher price, and on a flatter residual demand that higher price costs it more sales. Storage therefore displaces higher-cost gas generators more relative to the lower-cost gas generators.

I estimate the empirical analog of the generator’s optimal output rule using an instrumental variables approach. The optimal output of each generator is a function of local storage capacity, its marginal cost, and their interaction. The coefficient on storage capacity measures the effect on gas output at the typical cost of each technology, and the interaction measures how this effect varies with the generator’s marginal cost. Marginal cost is endogenous because realized heat rates — the amount of fuel needed to produce one MWh of electricity — move with output. I instrument marginal cost with the product of each generator’s pre-entry heat rate and the Henry Hub natural gas price. The optimality condition holds only when a generator runs between its operating limits, so I limit the sample to these interior hours. To measure the change in total gas output for the welfare analysis, I complement this approach with the a difference-in-differences design, including hours at the operating limits and interacting storage capacity with the predetermined cost shifter.

Storage displaces gas output in the afternoon and evening, and it displaces higher-cost generators more, as the model predicts. On interior hours, an additional 100 MWh of local storage lowers the capacity factor of a generator at average cost by 0.55 to 0.79 percentage points. The displacement is concentrated between 1 PM and 10 PM, with no significant change overnight. The interaction between storage and cost is negative in both designs, so storage shifts output toward lower-cost gas generators. The OLS and IV estimates are nearly identical, and the instruments are strong. These results hold when I control for the solar and wind capacity that often enters alongside storage. Aggregated across counties, the estimates imply that local storage displaced about 849 GWh of gas output per year. This is more than four times what the same storage discharged, likely because storage also relieves congestion and some gas output may shift to counties without storage. I therefore treat this aggregate as an upper bound.

In the welfare accounting, the competitive effect of storage is mainly a transfer. Local storage lowered gas operating profits by about \$47 million per year, of which \$34 million to \$55 million was market power rent. Storage owners kept about \$6 million of this transfer through energy arbitrage, and consumers and other generators gained the remaining \$41 million. My estimates suggest that most of it went to consumers, although the split between the two is imprecise. Storage also earns its arbitrage revenue from the markups it erodes. Valued at market prices, energy arbitrage earns storage owners \$5.8 million per year. Valued at the marginal cost of energy, the same operations are worth slightly less than zero, because the cost spread between charging and discharging hours barely covers round-trip losses. Storage earns its arbitrage profit from gas generators’ markups, so its competitive effect erodes its own profit.

The measured social benefits of storage do not cover its capital cost. Avoided emissions from displaced gas output are the largest measured benefit, at up to \$90 million per year. This figure

is an upper bound because it relies on the local displacement. Even at this bound, the measured benefits fall short of the \$179 million annualized capital cost of storage. Storage owners also earned \$35.6 million per year from ancillary services, which pay storage to hold capacity in reserve for grid reliability. These payments made up 86 percent of their total revenue of \$41.4 million, or \$54 per kW-year.³ Even this total covers only about a quarter of the annualized capital cost. At current penetration, the case for storage rests on its reliability services and on who gains from its competitive effect, rather than on energy arbitrage.

This paper contributes to three strands of literature. First, it contributes to the broad industrial organization literature on the effect of entry on incumbent firm behavior. A body of empirical work shows that incumbents respond to the threat or realization of entry across industries, from retail (Arcidiacono et al., 2020; Busso and Galiani, 2019; Taylor and Muehlegger, 2026) to airlines (Goolsbee and Syverson, 2008) and pharmaceuticals (Ellison and Ellison, 2011). Arcidiacono et al. (2020) study the entry of Walmart supercenters and find that nearby retailers lose revenue with no significant change in prices. Taylor and Muehlegger (2026) uses high-frequency microdata to evaluate how incumbents adjust to entry. I extend this framework to wholesale electricity markets, where I observe the incumbent’s output, pricing, and cost. Observing cost allows me to measure realized margins directly and to separate lost market power rent from a lower competitive price. I find that incumbents keep setting the price after entry, but at a smaller margin. A distinctive feature of storage as an entrant is that its competitive effect varies across hours of the day, because storage buys when incumbents face little competition and sells when they face the most.

Second, this paper contributes to the literature on competition and market power in wholesale electricity markets. A large body of work documents how generators exercise market power under transmission constraints and localized congestion (Borenstein, Bushnell and Wolak, 2002; Hortaçsu and Puller, 2008; Ito and Reguant, 2016; Woerman, 2023), and recent studies show that transmission expansion (Ryan, 2021; Doshi, 2025), financial arbitrageurs (Mercadal, 2022), and the price exposure of renewables (Fabra and Imelda, 2023) can reduce it. Closer to this paper, two recent studies examine how large-scale solar entry affects conventional generators. Jha and Leslie (2025) show that it softens evening competition through start-up costs, while Bushnell and Novan (2021) show that it shifts the marginal technology toward less efficient units. I study a different large-scale entry, storage, which competes in exactly the evening hours when solar entry softens competition. While transmission integrates markets across space, storage integrates them across hours of the day. I show that storage disciplines market power through a similar residual demand channel that this literature identifies as central to transmission-driven discipline.

Finally, this paper speaks to the growing literature on the economics of utility-scale storage. This work has largely quantified the aggregate, market-level impact of storage, documenting how it substitutes for or complements renewable capacity (Karaduman, 2023; Andres-Cerezo and Fabra,

³These figures are in line with industry estimates. Modo Energy reports that the average ERCOT battery earned \$141 per kW in 2022, \$192 per kW in 2023, and \$55 per kW in 2024, and that energy arbitrage provided about 15 percent of battery revenue in 2023 (Vermillion, 2024, 2025). My average spans 2018 to 2024, excludes Winter Storms Uri and Elliott, and is in 2020 dollars, which explains its lower level.

2025), lowers intraday price spreads (Lamp and Samano, 2022), shifts long-run investment equilibria (Butters, Dorsey and Gowrisankaran, 2025), and relieves local congestion (Kirkpatrick, 2026). Less is known about how storage entry changes the behavior of incumbents. Closest to this paper, Reaser and Taylor (2025) uses the entry of storage in Texas to estimate its effect on nodal prices, and builds a structural dispatch model that computes counterfactual prices without storage offers. I leverage the same source of variation, focusing on the marginal generators, and measure how storage changes marginal gas generator’s markups and output. I further separate the market power component of their lost profit from the effect of a lower competitive price. Other recent papers study different margins of storage entry. Dudek (2026) studies how storage reduces solar curtailment in California, and Visaggio (2025) studies how ownership structure shapes the market incentives of batteries paired with renewables in Texas. I complement these papers by focusing on the competition between storage and gas generators. I use unit-level storage dispatch to trace the channel from an individual storage unit’s entry to the behavior of nearby marginal gas generators. While prior work characterizes the theoretical price impact of storage (Andrés-Cerezo and Fabra, 2023), I provide the first causal evidence that observed storage entry lower realized markups. I also use a welfare framework that separates transfers from net social gains, and show that the private return to storage arbitrage depends on the market power it disciplines. This result helps explain why Karaduman (2023) and Butters, Dorsey and Gowrisankaran (2025) find that storage is rarely profitable without policy support. Storage earns most of its revenue from ancillary services rather than arbitrage, consistent with Reynolds (2025).

The remainder of this paper is organized as follows. Section 2 describes the Texas electricity market. Section 3 provides the data sources and presents the descriptive statistics on natural gas and storage units. Section 4 develops the model and derives theoretical predictions. Section 5 estimates the effect of storage on the markups and operating profits of gas generators and decomposes the lost profit into competitive and market power rents. Section 6 estimates the effect of storage on gas output. Section 7 presents the welfare analysis. Section 8 concludes.

2 Storage and The Texas Electricity Market

In this section, I provide an overview of the Texas electricity market, highlighting the unique features of ERCOT that make it an ideal setting to study the impact of storage entry on natural gas generator market power.

The Texas electricity market provides an ideal setting to study the competitive interactions between storage and natural gas generators. The Electric Reliability Council of Texas (ERCOT) manages this grid, delivering power to more than 26 million customers, about 90% of the state’s electric load, and operating a wholesale market that schedules the dispatch of over 1,500 generators in real time. Because it has limited interconnection with other regions, the market is relatively insulated from external shocks, creating a highly contained environment for empirical analysis. This system is particularly useful to study because of its generation mix; natural gas accounts for a

dominant 44.3% of the generation capacity, followed by wind (25.2%) and solar (13.2%) (ERCOT, 2025).⁴ At the same time, massive recent growth in battery storage capacity is actively altering these established competition dynamics. Furthermore, because ERCOT operates as an energy-only market with no capacity payments, it creates strong financial incentives for generators to respond strategically to real-time conditions.

The massive growth of utility-scale battery storage in ERCOT introduces new competitors to incumbent natural gas generators. Storage units operate as fully dispatchable assets that offer operational flexibility, functioning as generators when exporting and as load when charging. They participate in both the day-ahead and real-time markets providing both energy and ancillary services.⁵ Figure 1 documents this expansion, with ERCOT storage capacity growing to over 12 GWh since 2021.⁶

Because battery storage can respond quickly to price signals, these units can precisely target their dispatch during high-revenue periods where incumbent generators exercised market power. Panel B documents this pattern where storage discharges predominantly during the evening peak (5 PM to 9 PM). Panels C and D provide further descriptive evidence of competition during these hours: the share of storage units clearing at peak prices has increased, while the corresponding share of marginal natural gas generators has declined as more storage capacity enters. Before formally estimating the magnitude of this competitive effect, I first outline the features of the ERCOT market design that facilitate these interactions.

ERCOT shares a similar market feature with other deregulated electricity markets in the U.S. Like PJM, MISO, and CAISO, it runs a day-ahead and a real-time market and sets a separate wholesale electricity price at each node of the grid. This nodal price, also called the locational marginal price, reflects the cost of generation, transmission constraints, and line losses at that location.⁷ Generators submit offers for each hour of the next day in the day-ahead market. The real-time market then runs every five minutes, and the Security Constrained Economic Dispatch (SCED) clears it and sets the real-time nodal prices.

Transmission congestion creates local market power in every nodal market. When a constraint binds, it splits the grid into local markets and shields the generators within the local market from rivals on the other side of the constraint (LaRiviere and Lyu, 2022; Fell, Kaffine and Novan, 2021; Woerman, 2023; Doshi, 2025). Even small generators can then become pivotal and exercise market


⁴Figure F.22 shows the spatial distribution of each unit based on their fuel type in ERCOT. Table E.13 summarizes the generator units in ERCOT territory as of the end of 2024, while Figure F.24 details the installed capacity over my sample period: Natural gas units dominate the installed capacity (38.9% or 65.23 GW), followed by wind and solar PV units, respectively. Natural gas and coal capacity remain roughly the same over my sample, while solar PV and wind capacity grow significantly by about 25 GW and 20 GW.

⁵Because they serve as both generator and load, batteries can provide a wide range of regulation and reserve services. During summer 2023, for example, ERCOT Contingency Reserve Service (ECRS) accounted for more than 75% of total battery revenues, whereas real-time energy settlement contributed just under 15% (Miller and Nakajima, 2024). Figure F.23 illustrates how a single storage resource can simultaneously fulfill multiple AS responsibilities.

⁶In Q2 2025, ERCOT's installed storage capacity surpassed California's. A fundamental difference between the two markets is that ERCOT's expansion is driven by price signals and efficient interconnection queues, rather than the state-level procurement mandates or subsidies seen in California (Mathrani, 2025; Ahmed et al., 2026).

⁷Each generator is assigned to a settlement point and receives the nodal price at that location.

Figure 1: Storage expansion and its competition trends with natural gas generators in ERCOT



Notes: Panel A shows the cumulative storage capacity built and commercially operating over time. Panel B shows the aggregate storage net production pattern by hour. Net production is defined as aggregate storage discharge less aggregate storage charge. Panel C shows the share of storage units that receive the highest offered price in each hour. Panel D shows the share of natural gas generators that receive the highest offered price in each hour. Section C.1 provides the same yearly trend in hourly demand, net demand, aggregate solar, wind, and natural gas output.

power by raising their offers above marginal cost. Storage competes with gas generators inside these local markets. The mechanism I study therefore requires only nodal pricing, congestion, and storage that responds to prices, which every deregulated electricity market has.

ERCOT's market design differs from other markets in three ways that affect generators' incentives to raise their offers. First, it is an energy-only market. PJM, MISO, ISO-NE, and NYISO pay generators for available capacity through centralized auctions, and CAISO does so through bilateral resource adequacy contracts. ERCOT made no capacity payments during my sample, so generators recover their fixed costs from energy and ancillary service revenue, including scarcity prices in tight hours.⁸ Second, its offer cap was \$9,000/MWh until 2022 and \$5,000/MWh after, compared with \$1,000/MWh to \$2,000/MWh in markets subject to FERC Order 831.⁹ Third, ERCOT mitigates

⁸ Ancillary services in ERCOT include regulation (REGUP, REGDN) and operating reserves (Responsive Reserve [RRS], Non-Spin [NSPIN], and ERCOT Contingency Reserve Service [ECRS]).

⁹ Under FERC Order 831, non-ERCOT markets allow offer caps to increase to \$2,000/MWh to set the market

local market power more narrowly. When a transmission constraint is classified as non-competitive, SCED caps each offer at the greater of a cost-based mitigated offer cap and the price that would prevail without that constraint (ERCOT, 2024; Potomac Economics, 2025). Offers that affect only constraints classified as competitive are not mitigated. Other markets apply structural tests for pivotal suppliers to constraints as they bind, such as PJM’s three-pivotal-supplier test (Joskow, 2022). Together, these features leave gas generators in ERCOT more scope to exercise local market power than in other markets.

3 Data and Descriptive Evidence

This section describes the data sources and the construction of the panel dataset used to analyze the impact of storage entry on natural gas generator market power. I first outline the sample construction and the empirical definitions of local markets, marginal costs, and markups. I then provide descriptive statistics for natural gas generators and storage units.

3.1 Data Sources and Sample Construction

I combine multiple datasets to construct a panel at the generator-by-hour level covering January 1, 2018 to December 1, 2024. The main components are the generator and storage unit characteristics, supply offers and bids, the market outcomes, and marginal costs. The novel component of this dataset is that I observe each storage unit’s power capacity, storage duration, bid/offer curves, and unit-level charge and discharge quantity.

Generator and battery storage data. I identify unique generators from ERCOT’s 60-day SCED Report and the Monthly Outlook for Resource Adequacy Report.¹⁰ For all generators, I define effective capacity as the maximum production observed in the SCED Report, and I use the first date a generator produces electricity outside of testing as its initial commercial operating date.

Because ERCOT reports power capacity but not energy capacity for battery storage, I complement this data with the Energy Information Administration (EIA) Form 860 (EIA, 2025), the MODO Energy Battery Buildout Report (Modo Energy, 2026), and other public sources to determine each unit’s storage duration. Furthermore, because storage acts as both a generator and a load, it operates under different names when submitting supply offers versus load bids; I manually link these identifiers for each storage unit. Finally, I use location information from EIA Form 860 to assign the load zone and county for each generator, and the Decision Making Entity from the SCED Report to identify the parent company.¹¹

clearing price, but strictly on the condition that these higher offers are justified by verified short-run marginal costs.

¹⁰Two or more generators may be associated with a single storage or power plant unit. A generator’s name might also change before and after testing. For combined-cycle units, I use the maximum production from the generator combination with the highest possible output as the effective capacity.

¹¹The Decision Making Entity is often a project company, a marketer, or a battery optimizer rather than the owner. I use AI tools to match each entity to its parent company and verify each match against EIA Form 860, company filings, and press releases. I assign jointly owned generators to the company that submits their offers.

Market outcome data. I collect 15-minute interval data on offer and bid curves, dispatch instructions, operating status, and the identity of marginal generators directly from the SCED reports. I also extract 5-minute nodal prices at each settlement point, mapping them to individual generators and aggregating them to 15-minute intervals. I then aggregate these market outcomes to the hourly level, constructing variables for total hourly generation, average nodal price, and an indicator for whether a generator was marginal at any point during the hour.

To define the local markets created by transmission congestion, I group generators based on nodal price divergence and network connectivity at each interval, following [Mercadal \(2022\)](#) and [Woerman \(2023\)](#). I first group generators whose nodal prices are within \$1/MWh or 1% of the larger price. I then verify transmission connectivity within these groups using transmission lines and substation shapefiles collected by [Arkolakis and Walsh \(2024\)](#). Generators sharing both similar nodal prices and grid connections constitute a local market. Figure C.4 shows examples of market definition for two intervals, one with no congestion and one with congestion. The number of local markets varies from 1 to 137 across intervals, with a median of 8 local markets.¹²

I further use the local market definition to construct residual demand for each generator. I do so by stacking the bids of all other generators then compute the slope at the average nodal price in the local market using an Epanechnikov kernel for bids within \$5/MWh of the nodal price. The resulting slope represents the density of rival supply, measured in MW per (\$/MWh). This captures how many additional MW of rival capacity would enter the local market if the nodal price were to increase by one dollar. Appendix A provide more details on how I define the local market and the slope of residual demand.

Marginal cost. I construct the marginal cost of natural gas generator i at day t , c_{it} , as the sum of variable fuel cost plus emission permit costs following [Hausman \(2025\)](#) and [Doshi \(2025\)](#):

$$c_{it} = \underbrace{HR_{it} \cdot p_{it}^{Fuel}}_{\text{Fuel Cost}} + \underbrace{ER_{it}^{SO_2} p_t^{SO_2} + ER_{it}^{NO_x} p_t^{NO_x}}_{\text{Emission Cost}}$$

where HR_{it} is the generator's heat rate, p_{it}^{Fuel} is the generator's natural gas price, $ER_{it}^{SO_2}$ and $ER_{it}^{NO_x}$ are the generator's emission rate for SO_2 and NO_x , and $p_t^{SO_2}$ and $p_t^{NO_x}$ are the respective permit prices.

I use hourly production data from SCED Report and collect information on heat content and emissions to compute heat rates and emission rates at the plant-by-type-by-week level from the Environmental Protection Agency (EPA) Continuous Emission Monitoring System (CEMS) from 2018 to 2024.¹³ To compute fuel cost, I collect daily natural gas price from S&P Capital IQ. Natural

When a generator changes ownership during the sample, I assign it to the owner for most of the period. I define the parent as the operating company rather than the investment fund that owns it. For example, I code Calpine and Terra-Gen as separate parents, even though Energy Capital Partners owns both.

¹²Figure C.3 presents the distribution within my sample and across hours by year of sample. The number of local markets has increased over the years as more generators are added to the ERCOT grid.

¹³Heat rate and emission rate depends on the generator's operating condition as discussed by [Demirer and Karaduman \(2026\)](#). I compute the heat rate by dividing the weekly total heat content and emissions for that week by

gas units source their fuel from different hubs. I use the natural gas price location information for each unit.¹⁴ To compute emission cost, I collect permit prices for SO₂ and NO_x from S&P Capital IQ and apply the emission cost to units subject to the Acid Rain SO₂ Program and CSAPR NO_x Program.

Markup. I use realized markup as the primary outcome variable to measure market power, defined as:

$$\mu_{ith} = \mathbb{1}[q_{ith} > 0] (p_{ith} - c_{it}).$$

This markup is zero when the generator is not producing electricity, and equal to the difference between its nodal price p_{ith} and marginal cost c_{it} when it is producing. This measure captures both the extensive margin (whether the generator is producing) and the intensive margin (how much markup it earns when producing). This is consistent with the electricity market literature that measures markup over all unit-time rather than only when a unit is dispatched (Bushnell, Mansur and Saravia, 2008; Borenstein, Bushnell and Wolak, 2002; Reguant, 2014).

Other data. To control for local weather conditions, I construct a balanced hourly panel of local temperature using station data from the NOAA Integrated Surface Database (ISD) (NOAA, 2024), interpolating minor gaps and assigning local temperatures to each power plant based on its nearest weather station. I also construct hourly renewable energy potential at the county level from the global horizontal irradiance data from National Solar Radiation Database (NSRDB) (NREL, 2024) and wind speed from ERA5 global reanalysis data (C3S, 2018).

Sample construction. I limit my estimation sample to marginal natural gas generators, defined at the generator-by-hour level. To prevent endogenous selection from storage shifting the merit order, I restrict the sample to generator-hour pairs that set the nodal price prior to local storage entry, requiring at least 100 pre-treatment observations to ensure I capture established operational patterns. I exclude the period of Winter Storms Uri and Elliott from my estimation sample. I drop generators in Howard County, where a battery operated before my sample begins, because it has no untreated period. Next, I describe the gas generator’s characteristics.

3.2 Descriptive Statistics on Natural Gas Generators

Table 1 presents the descriptive statistics on natural gas generators. My sample consists of 553 natural gas generators (210 Simple Cycle, 294 Combined Cycle, and 49 Steam Turbine generators),

the weekly total net production for each plant-type combination. For units with no observation in CEMS, I assign the average heat rate for the plant with the same Type (Combined Cycle, Simple Cycle, or Gas Steam Turbine). I winsorize the extreme values to be within the reported heat rate range for the same generator type from EPA eGrid, to account for unreasonably high and low values from CEMS.

¹⁴For units missing this information in the S&P Capital IQ data, I assign Waha Hub for units in the West load zone, NGPL TX-OK for units in the North load zone, TETCO STX for units in the South load zone, and Katy for units in the Houston load zone.

Table 1: Summary statistics of natural gas generators

	All			Treated			Never-Treated		
	All	12AM –12PM	12PM –12AM	All	12AM –12PM	12PM –12AM	All	12AM –12PM	12PM –12AM
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Market Outcome and Characteristics (2018–2021)									
Nodal Price (\$/MWh)	28.7 (79.2)	21.7 (38.4)	35.7 (104.8)	29.0 (80.8)	21.9 (39.8)	36.1 (106.7)	27.7 (73.2)	21.1 (32.9)	34.5 (97.7)
Capacity Factor	0.11 (0.18)	0.08 (0.16)	0.13 (0.20)	0.10 (0.17)	0.08 (0.15)	0.11 (0.19)	0.14 (0.20)	0.11 (0.18)	0.17 (0.21)
Probability Operating	0.34 (0.47)	0.31 (0.46)	0.38 (0.48)	0.33 (0.47)	0.30 (0.46)	0.36 (0.48)	0.39 (0.49)	0.34 (0.47)	0.44 (0.50)
Probability Marginal	0.10 (0.29)	0.10 (0.30)	0.09 (0.29)	0.10 (0.30)	0.10 (0.30)	0.09 (0.29)	0.09 (0.29)	0.09 (0.29)	0.09 (0.29)
Realized Markup (\$/MWh)	-2.78 (79.88)	-9.76 (39.97)	4.20 (105.23)	-4.33 (81.49)	-11.38 (41.38)	2.74 (107.13)	2.74 (73.53)	-3.95 (33.84)	9.45 (97.96)
Resid. Demand Slope (MW per \$/MWh)	1234 (703)	1443 (663)	1064 (689)	1227 (705)	1435 (668)	1054 (689)	1259 (695)	1470 (646)	1094 (687)
Market Size (GW)	86.6 (29.5)	88.1 (28.5)	85.1 (30.3)	85.6 (30.5)	87.1 (29.6)	84.0 (31.4)	90.4 (24.8)	91.7 (23.6)	89.0 (25.8)
Panel B: Generator Characteristics									
Heat Rate (MMBtu/MWh)		10.7 (3.0)			11.0 (3.0)			9.8 (3.1)	
CO ₂ Rate (tons/MWh)		0.6 (0.2)			0.6 (0.2)			0.5 (0.2)	
NO _x Rate (lbs/MWh)		1.145 (1.974)			1.163 (1.752)			1.079 (2.625)	
SO ₂ Rate (lbs/MWh)		0.010 (0.044)			0.011 (0.050)			0.007 (0.006)	
Marginal Cost (\$/MWh)		31.79 (11.28)			32.68 (10.29)			28.63 (13.84)	
Avg. Capacity (MW)		368 (447)			334 (451)			488 (413)	
Panel C: Sample Composition									
# of Unique Generators		553			432			121	
# Simple Cycle		210			184			26	
# Combined Cycle		294			209			85	
# Steam Turbine		49			39			10	
# of Counties		49			30			19	
# in Houston Zone		220			208			12	
# in North Zone		136			56			80	
# in South Zone		152			140			12	
# in West Zone		45			28			17	

Notes. This table presents summary statistics of natural gas generators in ERCOT. Each entry reports the mean and counts of natural gas generators, with standard deviations in parentheses. Panels A and B report statistics computed over 2018–2021, the period before significant storage build-out in ERCOT. Panel C reports sample composition over the full estimation period (2018–2024). Columns (1)–(3) report statistics for all natural gas generators; Columns (4)–(6) for generators sharing the same county with storage; Columns (7)–(9) for generators located not in the same county with any storage. Panel A describes the variation in market outcome and characteristics within a day. Panel B describes the generator characteristics and marginal costs. Panel C describes sample composition.

with an average capacity of 366 MW and an average marginal cost of \$31/MWh. Geographically, these generators cluster heavily in the Houston, South, and North load zones.

Columns (1)–(3) detail within day variation in market outcomes: later hours (12 PM–12 AM) which include the evening peak hours consistently feature higher nodal prices, capacity factors, and markups compared to early hours. Figure C.5 presents a detailed daily pattern of these outcomes over the years. These elevated afternoon and evening markups link directly to tighter market conditions, specifically, smaller local market sizes and steeper residual demand that enhance generator market power. Columns (4)–(6) present identical statistics for generators that share the same county with storage units, while Columns (7)–(9) present statistics for generators located in different counties from any storage units. Because the operational patterns and market conditions of these groups closely mirror the overall sample, the treatment group in my difference-in-differences design is highly representative of ERCOT natural gas generators. The nodal price spread between early and late periods does not differ significantly between the treatment and control groups. Appendix D.1 test for selection in storage siting directly.

3.3 Descriptive Statistics of Battery Storage Units

This section describes the characteristics, ownership, siting, and operation of battery storage units in ERCOT.

Table 2, Column (1), shows that 13.1 GWh of storage (218 units) is built by the end of my sample, with 12.3 GWh (168 units) operating commercially. The average unit has 48.8 MW of capacity and a 74.5 MWh of energy capacity, this means that an average unit at full charge can produce up to 48.8 MW for up to 1.5 hours (i.e. 74.5 MWh divided by 48.8 MW). Units are spread across the state, with the largest shares in the West and South load zones.

Storage that shares a county with a natural gas generator provides my identifying variation. Column (2) shows 6.3 GWh of such capacity (106 units), with 6 GWh (75 units) operating commercially. Capacity and duration resemble the full sample. Because gas generators cluster near population centers, these units sit more often in the Houston and South zones and less often in the West. At the end of my sample, the top five county with largest storage capacity are Hidalgo, Brazoria, Wharton, Travis, and Harris county. Figure F.26 details the variation of commercially operating battery capacity across county in Texas within my sample period. Column (3) shows that the first entrant into each county varies widely in capacity.

Ownership and collocation differ across storage units. Panel B shows that about half of the units belong to storage-only firms, while the rest belong to firms that also own renewables. Storage often shares a node with solar PV, mainly in the North zone, and rarely with wind, which is entirely in the West. Storage and gas units rarely share a node.


Observed storage operation appears consistent with price arbitrage. Figure 2 shows that units tend to charge when nodal prices are low, demand is low, and renewable output is high. Units also tend to charge during less congested periods (larger local markets) and discharge during congested hours. Woerman (2023) documents that firms exercise more market power in smaller markets. This pattern suggests that storage entry may affect competition with gas generators where market power is greatest.

Table 2: Summary statistics of storage units

	All Storage (1)	In County w/ Gas (2)	First in County w/ Gas (3)	Collocated w/ Solar (4)	Collocated w/ Wind (5)
Panel A: Storage Characteristics					
Avg. Capacity (MW)	48.8 (51.0)	51.0 (53.3)	61.4 (61.7)	61.7 (38.9)	48.3 (46.9)
Avg. Energy Capacity (MWh)	73.5 (87.4)	79.6 (94.5)	89.1 (96.5)	84.4 (43.4)	60.8 (65.8)
Avg. Duration (hr)	1.36 (0.61)	1.36 (0.64)	1.34 (0.47)	1.46 (0.72)	1.33 (0.91)
Year of First Appearance	2022.8 (1.7)	2023.0 (1.6)	2022.4 (1.7)	2022.7 (1.5)	2021.5 (2.3)
Year of First Operation	2023.0 (1.6)	2023.0 (1.4)	2022.9 (1.6)	2023.4 (1.5)	2022.5 (2.7)
Cumulative Capacity Built (GWh)	13.1	6.3	2.7	2.4	0.5
Cumulative Capacity Operational (GWh)	12.3	6.0	2.7	2.0	0.5
Panel B: Ownership Structure					
By Parent Company					
Share w/ NG Portfolio	0.05	0.07	0.07	0.04	0.00
Share w/ Wind Portfolio	0.43	0.36	0.43	0.58	1.00
Share w/ Solar Portfolio	0.49	0.47	0.57	1.00	1.00
Share Independent	0.48	0.51	0.40	0.00	0.00
By Site Collocation (ERCOT Designation)					
Share Collocated w/ Natural Gas	0.02	0.05	0.03	0.00	0.00
Share Collocated w/ Solar	0.14	0.13	0.17	1.00	0.00
Share Collocated w/ Wind	0.05	0.00	0.00	0.00	1.00
Panel C: Sample Composition					
# of Battery Units (Operational)	168	75	30	24	9
# of Battery Units (Built)	218	106	32	27	9
# of Counties	73	30	30	16	6
# in Houston Zone	19	19	4	2	0
# in North Zone	30	16	10	14	0
# in South Zone	53	30	12	6	0
# in West Zone	66	10	4	2	9

Notes. Summary statistics of ERCOT storage units as of December 1, 2024. I report means and counts, with standard deviations in parentheses. Column (1) includes all units. Column (2) includes units in a county with a natural gas generator. Column (3) includes the first storage unit to enter each of those counties. Columns (4) and (5) include units collocated with solar PV and wind generators.

Figure 2: Battery storage capacity factor by market conditions



Notes: Each panel shows the conditional mean and 95% confidence interval of storage capacity factor by quantile of the explanatory variable. Capacity factor is energy charged or discharged divided by the maximum possible in each interval; positive values indicate charging and negative values indicate discharging. Market size is total installed capacity in the local market. Following [Lamp and Samano \(2022\)](#), I estimate $\text{StorageCF}_{it} = \sum_q \beta_q \mathbf{1}(X_{it} \in q\text{-th quantile of } \mathbf{X}) + \varepsilon_{it}$ for generator i in interval t . Nodal price and market size vary by generator and interval; demand and renewable output vary by interval. Standard errors are two-way clustered by generator and day.

4 Theoretical Framework

I build a simple two-period model to show how storage entry changes the competitive behavior of natural gas generators (gas generators, hereafter). I model the real-time market as a uniform-price auction in which a strategic gas generator competes with price-taking rivals, following the multi-unit auction framework of [Wilson \(1979\)](#) and [Hortaçsu and Puller \(2008\)](#) as applied to transmission-constrained electricity markets ([Ryan, 2021](#); [Woerman, 2023](#); [Doshi, 2025](#)). In this market, a gas generator’s markup is larger when it serves more of the demand left by its rivals and when that demand is less responsive to price. Storage changes markup through both channel. When batteries discharge at the peak, they serve part of the demand, so the gas generator generate less electricity. I refer to this as the *level* channel. Because batteries discharge more when the price is higher, the demand left for the gas generator also becomes more responsive to price. I refer to this as the *slope* channel. Both channels lower the markup when batteries discharge. When batteries charge, they add load and raise the level, so the two channels push the markup in opposite directions. The appendix extends the model to account for transmission congestion, strategic storage, and portfolio ownership characteristic.

4.1 Setup

The model has two periods, a low-demand charge period and a high-demand discharge period, $h \in \{L, H\}$, with inelastic load D_h and $D_L < D_H$. In each period, the generator’s local market clears as a uniform-price auction: suppliers submit offers, offers are dispatched in merit order until supply meets load, and every dispatched unit receives the clearing price p_h . Three groups submit offers. A competitive fringe of price-taking units offers at marginal cost, so its aggregate supply is $S_h^c(p) = \alpha_h + \gamma p$, where α_h is its level and $\gamma > 0$ is its slope ([Fabra and Imelda, 2023](#); [Ito and Reguant, 2016](#)). A storage fleet with net injection s_h charges in the charge period ($s_L \leq 0$) and discharges in the discharge period ($s_H \geq 0$). A strategic gas generator i with constant marginal cost c_i supplies q_{ih} . The market clears in each period,

$$(1) \quad S_h^c(p_h) + q_{ih} + s_h(p_h) = D_h.$$

The generator exercises market power by withholding. Every unit it sells receives the clearing price, so a higher offer raises the price on all of its inframarginal output, at the cost of the sales it loses to rivals. Given its rivals’ offers, choosing an offer is equivalent to choosing a point on its residual demand, the load left after rivals’ supply. The generator therefore chooses q_{it} to maximize its profit $\pi_{ih} = (p_h - c_i) q_{ih}$ subject to Equation (1). I focus on the case in which the generator’s offer is marginal and sets the price, which matches my empirical sample of marginal gas generators.¹⁵

¹⁵I abstract from demand uncertainty and forward contracts. With uncertainty, the markup rule below holds in expectation over residual demand ([Hortaçsu and Puller, 2008](#); [Woerman, 2023](#)). With a forward position, the quantity in the markup rule becomes output net of that position ([Woerman, 2023](#); [Doshi, 2025](#)).

Storage supply. Each battery maximizes arbitrage profit. To sell one MWh in the discharge period, a battery with round-trip efficiency $\eta \leq 1$ must buy $1/\eta$ MWh in the charge period, so its marginal cost is the opportunity cost $m_s = p_L/\eta$. Because the objective is linear in the arbitrage quantity, it discharges at full capacity when $p_H > m_s$ and stays idle otherwise. A single battery's offer is a step, consistent with awarded storage quantities clustering on the vertical segment of the offer curve observed in the storage bid data (Figure F.23). Batteries differ in efficiency, state of charge, and price expectations, so a higher price pulls successively more of them past their threshold, and the aggregate supply is smooth (Karaduman, 2023). I approximate it as

$$(2) \quad s_h(p_h) = \bar{s}_h + \sigma_h p_h, \quad \sigma_h \geq 0,$$

where $\bar{s}_h$ is the level of storage supply and σ_h is its slope. The level is positive when storage discharges and negative when it charges. The slope reflects batteries near their threshold, which respond to the price in both periods: more of them discharge at the peak, and fewer of them charge off-peak. Storage differs from the fringe because its offer price is an opportunity cost set by the price in the other period, and the same capacity sells at the peak and buys off-peak. Equation (2) takes the other period's price as given.

4.2 The Generator's Optimal Markup

The gas generator's optimal markup is proportional to the quantity it serves and inversely proportional to the slope of its residual demand. Substituting fringe supply and storage supply into Equation (1), the generator faces the residual demand

$$(3) \quad q_{ih} = A_h - E_h p_h, \quad A_h = D_h - \alpha_h - \bar{s}_h, \quad E_h = \gamma + \sigma_h,$$

where A_h is the level of residual demand and E_h is its slope. The first-order condition gives the markup rule

$$(4) \quad \mu_{ih} \equiv p_h - c_i = \frac{q_{ih}}{E_h},$$


where μ_{ih} is the generator's i markup at hour h . This rule is the residual-demand form of the Lerner index, and it holds in any oligopoly model in which a firm faces a downward-sloping residual demand. Under Cournot competition, rivals' quantities are fixed, so the slope of residual demand comes only from market demand, and the rule becomes $(p_h - c_i)/p_h = \theta_{it}/\varepsilon_h$, where θ_{it} is the generator's market share and ε_h is the price elasticity of market demand. When firms compete in supply schedules, as in electricity auctions, rivals' supply also responds to the price and adds to the slope (Klemperer and Meyer, 1989). Because load is inelastic, rivals' response is the only source of slope in my setting, so $E_h = \gamma + \sigma_h$ sums the offer slopes of the fringe and storage. The rule is standard in the empirical literature on electricity markets (Jha and Leslie, 2025; Woerman, 2023; Ryan, 2021; Doshi, 2025; Fabra and Imelda, 2023). Solving jointly with residual demand gives the

equilibrium price (p_h^*), output (q_{ih}^*), and markup (μ_{ih}^*),

$$(5) \quad p_h^* = \frac{1}{2} \left(\frac{A_h}{E_h} + c_i \right), \quad q_{ih}^* = \frac{1}{2} (A_h - E_h c_i), \quad \mu_{ih}^* = \frac{A_h}{2E_h} - \frac{c_i}{2}.$$

The markup rises with the level of residual demand and falls with its slope. I abstract from any effect of storage on the generator's own marginal cost.

Figure 3: Illustration of market power effect of storage entry



Notes: Residual demand of a natural gas generator before (gray) and after (red) storage entry. The generator sets $MR = c$, and the markup μ is the gap between the price and marginal cost c . Panels (a) and (b) illustrate a flatter residual demand around the pre-entry optimum to isolate the slope channel. Panels (c) and (d) add the level channel: an outward shift when storage charges and an inward shift when it discharges. The faded curve in panels (c) and (d) repeats the slope-only curve. Parameters are illustrative.

Figure 3 illustrates the two channels. In the figure, superscripts 0 and 1 denote outcomes before and after storage entry, and a prime denotes the outcome under the slope channel alone. Panels (a) and (b) isolate the slope channel. Residual demand becomes flatter but still passes through the pre-entry optimum, so it rotates around that point. A flatter curve raises the cost of withholding. The generator offers more ($q_h' > q_h^0$), and markup falls ($\mu_h' < \mu_h^0$) in both periods.

Panels (c) and (d) add the level channel, which works in opposite directions in the two periods. In the discharge period, storage serves part of the load and shifts residual demand inward. Output falls below its pre-entry level ($q_H^1 < q_H^0$), and the markup falls further ($\mu_H^1 < \mu_H^0$). In the charge period, storage adds load and shifts residual demand outward. Output rises further ($q_L^1 > q_L^0$), and the markup recovers. In this example, the level channel dominates, so the markup ends above its pre-entry level ($\mu_L^1 > \mu_L^0$). With less charging load or a more price-responsive storage fleet, the slope channel would dominate, and the markup would fall.

4.3 Storage Entry and the Generator's Outcomes

Local storage shifts both the level and the slope of residual demand. Let B be the installed battery energy capacity in the generator's local market, and let $\partial s_h / \partial B$ be the change in storage net injection at the clearing price. More storage raises net injection when it discharges, lowers it when it charges, and adds price-responsive capacity in both periods,

$$(6) \quad \frac{\partial s_H}{\partial B} \geq 0, \quad \frac{\partial s_L}{\partial B} \leq 0, \quad \frac{\partial E_h}{\partial B} = \frac{\partial \sigma_h}{\partial B} \geq 0.$$

The level channel moves residual demand by $-\partial s_h / \partial B$ at the clearing price, and the slope channel rotates residual demand around that point. The following comparative statics describe how storage affects the generator's output and markup.¹⁶

Output. Differentiating equilibrium output separates the two channels,

$$(7) \quad \frac{\partial q_{ih}}{\partial B} = \underbrace{-\frac{1}{2} \frac{\partial s_h}{\partial B}}_{\text{level}} + \underbrace{\frac{1}{2} \mu_{ih} \frac{\partial E_h}{\partial B}}_{\text{slope}}, \quad \frac{\partial^2 q_{ih}}{\partial B \partial c_i} = -\frac{1}{2} \frac{\partial E_h}{\partial B} \leq 0.$$

The slope term is positive because a flatter residual demand makes withholding more costly, so the generator offers more. In the charge period, both terms raise output. In the discharge period, the level term lowers output and the slope term raises it, so output falls only if the load storage takes over exceeds the extra output the generator offers as it withholds less, $\partial s_H / \partial B > \mu_{iH} \partial E_H / \partial B$. The cross-derivative implies that a flatter residual demand makes output more responsive to cost, so storage shifts output from higher-cost toward lower-cost gas generators.

Markup. Differentiating the equilibrium markup separates the same two channels,

$$(8) \quad \frac{\partial \mu_{ih}}{\partial B} = \underbrace{-\frac{1}{2E_h} \frac{\partial s_h}{\partial B}}_{\text{level}} - \underbrace{\frac{\mu_{ih}}{2E_h} \frac{\partial E_h}{\partial B}}_{\text{slope}}.$$

¹⁶Both expressions follow from differentiating Equation (5) and using $\partial A_h / \partial B = -\partial \bar{s}_h / \partial B = -\partial s_h / \partial B + p_h \partial E_h / \partial B$ and $A_h / E_h = 2p_h - c_i$.

The slope term is always negative because a flatter residual demand raises the cost of withholding. In the discharge period, both terms are negative, so the markup falls. In the charge period, the two terms oppose, and the markup rises only if the added charging load exceeds the flattening of residual demand, $-\partial s_L/\partial B > \mu_{iL} \partial E_L/\partial B$.

4.4 Summary of Theoretical Results

The model produces four predictions that motivate the empirical analysis. First, storage flattens the residual demand of gas generators in both periods. Second, in the discharge period, storage lowers the markup of gas generators, and it lowers their output when the level channel dominates. Third, in the charge period, storage raises the output of gas generators, but its effect on the markup is ambiguous. The markup rises only when the added charging load outweighs the flatter residual demand. Fourth, storage shifts output from higher-cost toward lower-cost gas generators.

5 Impact of Storage on Market Power

This section evaluates how local battery storage affects the markups and operating profits of incumbent natural gas generators. I test the model prediction on how storage affect market power of gas generator differently for each hour of the day. I estimate the marginal effect of storage separately for each hour of the day leveraging the staggered entry of battery capacity into the generator’s local market. I provide reduce-form evidence that storage entry reduce markup and operating profit during the evening peak.

5.1 Empirical Strategy

The treatment is the entry of storage into a generator’s county. I define the commercial operation date $COD_{m(i)}$ as the first date a same-county storage unit is awarded a positive quantity in the real-time market. The treatment is a continuous dose $B_{m(i),t}$, the storage capacity in county m of generator i at date t . The county boundary approximates the competitive neighborhood that transmission constraints create (Figure F.25). The interpretation these treatment effect is the marginal effect of battery capacity commercially operating in the generator’s local market more than 99% of the time.

For an hourly outcome Y_{iht} , the estimating equation is

$$(9) \quad Y_{iht} = \sum_{h=0}^{23} \beta_h \mathbb{1}\{\text{hour} = h\} B_{m(i),t} + \alpha_i + \eta_{\ell(i),t,h} + \mathbf{X}_{iht}\theta + \varepsilon_{iht},$$

where $\eta_{\ell(i),t,h}$ is a load-zone-by-date-by-hour fixed effect. The controls $\mathbf{X}_{iht}$ are system net load, a cubic in local temperature, and local renewable capacity and potential output: county solar and wind capacity, and their products with local irradiance and a normalized wind power curve. I

cluster standard errors at the county level. Section 6.1 extends this design to the capacity factor by interacting storage with predetermined cost.

Counties adopt storage at different dates and the dose grows over time. Under this staggered timing, the two-way fixed-effects estimator uses already-treated counties as controls, which can attenuate or reverse the sign of the true effect when effects are heterogeneous (Goodman-Bacon, 2021; De Chaisemartin and D’Haultfœuille, 2020; Callaway and Sant’Anna, 2021; Sun and Abraham, 2021; Borusyak, Jaravel and Spiess, 2024). I therefore use the two-stage estimator of Gardner et al. (2024). The first stage estimates the fixed effects and control coefficients on never-treated and not-yet-treated observations. The second stage regresses the first-stage residuals on the storage terms. The standard errors account for the first stage.

With a continuous and time-varying dose, β_h is a positively weighted average of the local slopes of the dose-response function (Yitzhaki, 1996; Gardner et al., 2024),

$$(10) \quad \hat{\beta}_h = \mathbb{E} \left[\sum_d w(d) \frac{\partial \mathbb{E}[Y_{iht}(d) - Y_{iht}(0) \mid B_{m(i),t} = d]}{\partial d} \right],$$

where the weights $w(d) \geq 0$ reflect the distribution of the dose among treated generators. The coefficient is therefore a marginal effect scaled to per 100 MWh rather than the effect of moving from no storage to 100 MWh. The identification results of Callaway, Goodman-Bacon and Sant’Anna (2024) assume a dose fixed at a single value per unit, which does not hold in this setting.

Identification rests on three assumptions. The first is parallel trends: absent local entry, the untreated outcome is additive in the fixed effects and controls, $\mathbb{E}[Y_{iht}(0) \mid \alpha_i, \eta_{\ell(i),t,h}, \mathbf{X}_{iht}] = \alpha_i + \eta_{\ell(i),t,h} + \mathbf{X}_{iht}\theta$. The second is no anticipation, $Y_{iht}(d) = Y_{iht}(0)$ for $t < \text{COD}_{m(i)}$. The third is no interference, so a generator’s outcome depends only on its own county’s dose.

I examine whether there is parallel trends in the pre-treatment period with an event study that replaces the dose with indicators for years relative to entry,

$$(11) \quad Y_{iht} = \sum_{\substack{k=-3 \\ k \neq -1}}^3 \tau_k \mathbb{1}\{K_{m(i),t} = k\} + \alpha_i + \eta_{\ell(i),t,h} + \mathbf{X}_{iht}\theta + \varepsilon_{iht},$$

where $K_{m(i),t}$ counts years relative to entry, binned up to three years pre and post, and the year before entry is the reference. I estimate it by hour group and at the daily level for operating profit. Anticipation is unlikely, because storage appears in the SCED data long before commercial operation, with a median gap of 204 days (Figure F.27), and the data are published with a sixty-day lag. The remaining concern is that storage enters where local conditions are changing. The generator and load-zone-by-date-by-hour fixed effects absorb permanent and regional differences, and the event study then provides a falsification test for county-specific trends. Storage is often sited alongside new solar capacity as a new cheap supplier which might displace gas output directly. To address this, I control for the local renewable capacity and potential.

I also test directly whether storage enters where prices or gas conduct make entry attractive

(Appendix D.1). Storage siting is not random across Texas. Counties with more solar capacity and wider pre-period price spreads receive more storage, and receive it earlier (Table D.3). Solar capacity is the strongest predictor, consistent with the co-location pattern above. Among counties with gas generator, pre-period prices do not predict whether a county receives storage (joint $p = 0.30$), although they still predict how much and when. The county fixed effects absorb these level differences. The estimates would be biased only if storage selects on trends, and I find no evidence that it does. County price spreads show no trend before entry (joint $p = 0.43$, Figure D.6). Within counties, batteries also do not locate closer to gas units with higher pre-period markups or prices (joint $p = 0.33$, Table D.4). The markup coefficient is close to zero and precisely estimated, so siting does not select on the conduct I measure.

For the markup, the outcome is $\mu_{iht} = p_{iht} - c_{it}$, set to zero when the generator does not run. For the residual demand slope, the outcome E_{iht} is the slope of the local residual demand curve constructed from the offer stack.¹⁷ For operating profit, I aggregate to the generator-day and use daily operating profit per MW of capacity, $\pi_{it} = \sum_h \mu_{iht} q_{iht} / K_i$, weighted by capacity, with load-zone-by-date fixed effects,

$$(12) \quad \pi_{it} = \beta^\pi B_{m(i),t} + \alpha_i + \eta_{\ell(i),t} + \mathbf{X}_{it}\theta + \varepsilon_{it}.$$

I also estimate (12) separately for the competitive rent and the market power rent defined in Section 5.2. Because the estimator is linear, the two effects sum to β^π .

5.2 Results


I present the marginal effects of nearby storage entry on generator markups and on the slope of their residual demand, followed by the hourly and daily impacts on operating profits. I then decompose profit impact into competitive and market power rents. After decomposing these profit changes, I conclude by estimating the aggregate rent transfer away from natural gas generators. Throughout this section, all treatment effects are scaled per 100 MWh of installed local storage capacity, and all values are reported in 2020 dollars.

Markups. Storage lowers the markups of gas generators in the late afternoon and evening. Figure 4 presents the two-stage estimates on the markup. Panel (a) shows that the markup does not change overnight or in the morning. Markup begins to fall at 1 PM, and the effect is largest at 7 PM, when an additional 100 MWh of local storage lowers the markup by \$2.71 per MWh (s.e. 1.17). The effect is also significant at 2 PM, 3 PM, 8 PM, and 9 PM. Over the evening peak from 5 PM to 9 PM, the markup falls by \$1.50 per MWh on average, about 5 percent of the pre-entry evening markup of \$31 per MWh. Panel (b) presents the corresponding event study estimates for the evening peak. The pre-treatment coefficients are near zero and statistically insignificant, supporting the identifying assumption. The post-entry coefficients grow more negative over time,

¹⁷Section 3 details the construction of E_{iht} . A larger value denotes a flatter, more elastic residual demand.

although they are imprecisely estimated. Finally, Panel (c) demonstrates that this markup compression occurs across all dose bins, yielding an average marginal evening effect of -1.31 (s.e. 0.67) per 100 MWh. The pattern is consistent across technologies (Appendix Figure D.7), with evening effects of -0.92 for combined cycle, -1.57 for simple cycle, and -1.70 for steam turbine generators.

Figure 4: Effect of local storage on markups




Notes. Two-stage difference-in-differences estimates. The markup is the nodal price minus marginal cost, set to zero when the generator does not run. Panel (a) reports the effect of 100 MWh of local storage by hour. Panel (b) reports a binary event study for 5 PM to 9 PM, with the year before entry as the reference. Panel (c) reports the effect of each dose bin relative to no storage for 5 PM to 9 PM. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Residual demand slope The model predicts that storage makes the residual demand of gas generators flatter, which is the slope channel behind the lower markup. To test this prediction, I use the slope of the local residual demand constructed from the bid data as the outcome in the two-stage design. Before storage entry, the slope is about 1,300 MW per \$/MWh overnight and about 750 to 900 MW per \$/MWh in the afternoon and evening. Figure 5 shows that each 100 MWh of local storage increases the slope by 1.3 to 4.3 MW per \$/MWh, or 0.2 to 0.5 percent of the baseline. The effect is statistically significant overnight and in the morning, and positive but

imprecisely estimated in the evening. The event study estimates exhibit no pre-trend in any hour group (Appendix Figure D.10). Across technology, the change in slope is stronger for combined cycle generators in the evening, when the model predicts market power binds. Their residual demand becomes flatter by 5.5 to 8.6 MW per \$/MWh per 100 MWh of storage between 6 PM and 11 PM (Appendix Figure D.9).

Figure 5: Effect of local storage on the residual demand slope




Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the slope of the local residual demand curve, in MW per \$/MWh, by hour, for all gas generators. A larger value is a flatter, more elastic residual demand. The slope is winsorized at the 99th percentile within technology. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Operating profit The markup results translate into lost operating profit in the same hours. Figure 6 presents the effect on hourly operating profit per MW of capacity, weighted by capacity. Gas generators lose operating profit from 3 PM to 8 PM, with the largest loss of \$1.55 per MW at 7 PM. Overnight, operating profit remains the same. Summed over the day, the hourly effects imply a loss of about \$4.9 per MW-day.


The daily estimates confirm this loss. I aggregate the panel to the generator-day and estimate the effect on daily operating profit per MW, weighted by capacity. An additional 100 MWh of local storage lowers daily operating profit by \$5.45 per MW-day (s.e. 1.85), or about 6 percent of the pre-entry average of \$84 per MW-day. This is close to the sum of the hourly effects. Figure 7 presents the event study and the dose response. The coefficients before entry are small and insignificant. After entry, operating profit falls by \$19 per MW-day in the first year and \$17 per MW-day in the second year, and the second-year effect is significant. The dose response shows that the loss grows with the installed dose. Counties with 200 to 400 MWh of storage lose \$28 per MW-day relative to counties without storage, and counties with more than 400 MWh lose the same amount.

Figure 6: Effect of local storage on hourly operating profit



Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on hourly operating profit per MW of nameplate capacity, weighted by capacity. The 95% confidence intervals use t_{G-1} critical values and county clustering.

Figure 7: Effect of local storage on daily operating profit



Notes. Two-stage difference-in-differences estimates on the generator-day panel. The outcome is daily operating profit per MW of nameplate capacity, weighted by capacity, with generator and load-zone-by-date fixed effects. Panel (a) reports a binary event study with the year before entry as the reference. Panel (b) reports the effect of each dose bin relative to no storage. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Why distinguish competitive rent and market power rent A fall in gas operating profit does not necessarily mean that storage reduces market power. A gas generator earns operating profit for two reasons. It earns a competitive rent when a more expensive unit sets the price, which it would earn even in a competitive market. Storage can lower competitive rent if its discharge lowers the load gas must serve, so a cheaper unit becomes marginal. It earns a market power rent when the price exceeds the competitive level. Storage can lower market power rents if it also

flattens residual demand, which narrows the gap between the price and the competitive level. Only the second is a reduction in market power. Separating the two shows whether storage disciplines the conduct of gas generators or simply adds supply.

I decompose each gas generator's operating profit into a competitive rent and a market power rent, following [Borenstein, Bushnell and Wolak \(2002\)](#). In each period, generator i belongs to unique local market $n(i)$, and $\mathcal{G}_{nth}$ denotes the set of gas generators in market n of day t and hour h . For generator i in hour h of day t , q_{ith} is its output in MWh, p_{ith} is the nodal price at its settlement point in \$/MWh, and $\bar{k}_{ith}$ is its available capacity in MW. Its marginal cost c_{it} in \$/MWh varies only by day, because it combines a weekly heat rate from CEMS with a daily fuel price.

I construct the competitive price from a cost-based supply stack. For each market n_{th} and 15-minute interval, I rank the generators in $\mathcal{G}_{nth}$ from lowest to highest marginal cost. I then fill the realized gas output $Q_{nth} = \sum_{j \in \mathcal{G}_{nth}} q_{jth}$ from the bottom of this stack. The competitive price is the marginal cost of the last generator needed to serve that output,

$$(13) \quad p_{nth}^c = \min \left\{ c_{jt} : j \in \mathcal{G}_{nth}, \sum_{l \in \mathcal{G}_{nth} : c_{lt} \leq c_{jt}} \bar{k}_{lth} \geq Q_{nth} \right\}.$$

This is the price that a cost-based dispatch of the same gas output would produce. It holds fixed the output of non-gas generators, storage, and imports into the market. I average the competitive price over the intervals in each hour.

Operating profit then splits into two parts by adding and subtracting $p_{n(i)ht}^c$,

$$(14) \quad \pi_{ith} \equiv (p_{ith} - c_{it}) q_{ith} = \underbrace{(p_{n(i)th}^c - c_{it}) q_{ith}}_{\text{competitive rent}} + \underbrace{(p_{ith} - p_{n(i)th}^c) q_{ith}}_{\text{market power rent}}.$$

The competitive rent is what generator i would earn on its realized output if the market cleared at cost. It is positive for units below the competitive margin. It can be negative for a unit that runs while its cost exceeds $p_{n(i)ht}^c$, for example because of minimum-run or commitment constraints. The market power rent is the price premium over the competitive benchmark, earned on each unit of output.

This central benchmark likely overstates the market power rent for two reasons. First, measured marginal cost omits startup, maintenance, and ramping costs, so the stack understates the competitive price. Second, p_{iht} is nodal while $p_{n(i)ht}^c$ is market-level, so the premium also absorbs local congestion and losses. I therefore treat the central estimate as an upper bound.


I construct a lower bound from the generator that actually sets the price. ERCOT identifies the marginal generator in each interval. When a gas generator sets the nodal price, I use its marginal cost as the competitive price. Because withholding pushes more expensive generators to the margin, this cost is at or above the true competitive price, so the benchmark understates the market power rent. When no gas generator sets the price, I keep the central benchmark, so the sample composition stays the same. A gas generator sets the nodal price in 58 percent of intervals.

I estimate the effect of storage on each rent component under both benchmarks, using the same specification, sample, and weights. Because the estimator is linear in the outcome, the two effects sum exactly to the effect on operating profit.

Rent decomposition The loss in operating profit falls mainly on the market power rent. Figure 8 presents the hourly effects on the two rents under the central benchmark. The market power rent falls from 2 PM to 8 PM, by up to \$1.60 per MW at 7 PM, and accounts for nearly all of the hourly profit loss. The competitive rent changes little in any hour. Table 3 presents the daily effects. Under the central benchmark, the market power rent falls by \$6.27 per MW-day (s.e. 1.74), while the competitive rent rises by an insignificant \$0.82. Under the price-setter benchmark, the market power rent falls by \$3.96 per MW-day (s.e. 1.80), and the competitive rent falls by an insignificant \$1.50. The market power rent therefore accounts for between 73 percent and all of the loss in operating profit, and it is significant under both benchmarks. The estimates by technology tell the same story. Under the price-setter benchmark, the market power rent accounts for 72 percent of the loss for combined cycle generators and 87 percent for simple cycle generators (Appendix Table D.5).

Storage lowers markups without displacing gas from being the price-setter. The probability that a gas generator sets the nodal price does not change significantly with local storage in any hour (Appendix Figure D.14). If storage reduced gas profits only by adding supply, it would push cheaper units or storage itself to the margin. I find no evidence of this in the data and gas generators continue to set the price, but at a smaller margin above cost. This pattern supports the interpretation that storage disciplines the market power of gas generators.

Figure 8: Effect of local storage on competitive and market power rents by hour



Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the competitive rent and the market power rent in Equation (14) under the central benchmark, per MW of nameplate capacity per hour, weighted by capacity. The two effects sum to the effect on hourly operating profit. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Table 3: Effect of local storage on daily operating profit and rents

	Operating profit (1)	Competitive rent (2)	Market power rent (3)	Competitive rent (price-setter) (4)	Market power rent (price-setter) (5)
\$/MW-day per 100 MWh	-5.45*** (1.85)	0.82 (0.79)	-6.27*** (1.74)	-1.50 (1.24)	-3.96** (1.80)
Share of profit effect	1.00	-0.15	1.15	0.27	0.73
Annual total (\$ million)	-47.4 (16.1)	7.1 (6.8)	-54.5 (15.2)	-13.0 (10.8)	-34.4 (15.6)
Per 100 MWh of storage (\$ thousand/year)	-4,057 (1,378)	609 (585)	-4,666 (1,297)	-1,114 (921)	-2,943 (1,337)
Observations (generator-days)	671,972				
Counties	49				

Notes. Two-stage difference-in-differences estimates on the generator-day panel, weighted by capacity, with generator and load-zone-by-date fixed effects and the local renewable and temperature controls. Columns (2) and (3) use the central competitive benchmark. Columns (4) and (5) use the price-setter benchmark, which sets the competitive price to the marginal cost of the gas generator setting the nodal price and keeps the central benchmark when no gas generator sets it. The annual total multiplies the daily effect by installed storage and capacity on treated generator-days and divides by the 6.9 sample years. Dollar amounts are in 2020 dollars, deflated by year with the CPI-U before estimation. The last rows divide the annual total by the average installed storage in gas counties, 1,170 MWh. Standard errors clustered at the county level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Aggregate transfer Scaling the daily effects by installed storage and generator capacity gives the aggregate transfer from gas generators. Local storage lowered the operating profit of gas generators by \$47.4 million per year (s.e. 16.1), or about \$4.1 million per year per 100 MWh of storage. Of this, the market power rent accounts for between \$34.4 million per year (s.e. 15.6) under the price-setter benchmark and \$54.5 million per year (s.e. 15.2) under the central benchmark. Aggregating over the technology-specific estimates yields a similar total loss of \$53.2 million, so the aggregate number is robust to pooling across technologies. These amounts are transfers. The rents gas generators lose accrue to storage owners through arbitrage revenue and to consumers and other generators through lower prices, and I account for them in the welfare analysis of Section 7.

Summary of market power results Local storage reduces the markups and operating profits of gas generators in the late afternoon and evening, when storage discharges and gas generators would otherwise exercise market power. Storage also flattens the residual demand of gas generators, consistent with the slope channel, although the pooled evening effect is imprecise. The daily profit loss of \$5.45 per MW-day per 100 MWh of storage grows with the installed dose. Between 73 percent and all of this loss falls on the market power rent, and gas generators continue to set the price after storage enters. Storage therefore disciplines the market power of gas generators rather than only adding supply.

6 Impact of Storage on Gas Generator Output

I take the model to the data using a panel of marginal gas generators in ERCOT at the hourly level. I estimate the effect of storage separately by hour of the day, with q_{iht} , μ_{iht} , and c_{it} the output, markup, and marginal cost of generator i on date t at hour h . Storage follows a daily cycle, so I group hours into charging periods, overnight (midnight to 6 AM) and midday (10 AM to 4 PM), and a discharging period, the evening peak (5 PM to 9 PM). I normalize output as a capacity factor, $y_{iht} = q_{iht}/K_i$, where K_i is the capacity of generator i . The capacity factor puts generators of very different sizes on a common scale, and multiplying by K_i recovers output in MWh.

I estimate the generator's optimal output rule from the theoretical model by instrumental variables. The estimates show how output responds to storage, to the generator's own cost, and to their interaction, as the model predicts.

6.1 Empirical Strategy

I bring the theoretical prediction from the previous section to the data. Equilibrium output in Equation (5) is linear in marginal cost, $q_{ih}^* = \frac{1}{2}(A_h - E_h c_i)$. I let local storage capacity shift the level and slope of residual demand linearly, $A_h(B) = A_h^0 + a_h B$ and $E_h(B) = E_h^0 + e_h B$, where $a_h = \partial A_h / \partial B$ and $e_h = \partial E_h / \partial B \geq 0$. Substituting gives

$$(15) \quad q_{ih}^* = \frac{1}{2}(A_h^0 - E_h^0 c_i) + \frac{1}{2} a_h B - \frac{1}{2} e_h B c_i.$$

The first term is the pre-entry output rule. The second term shifts output with storage, and the third term is the cross-derivative in Equation (7), so higher-cost generators lose more output as storage flattens residual demand. Normalizing Equation (15) by capacity yields the optimal output rule expressed in capacity factor.

I estimate the empirical analog of Equation (15) on interior hours,

$$(16) \quad y_{iht} = \psi_h B_{m(i),t} + \phi_h B_{m(i),t} c_{it} + \rho c_{it} + \mathbf{X}_{iht} \theta + \alpha_i + \nu_{\ell(i),y(t),h} + u_{iht}.$$

The outcome y_{iht} is the capacity factor in percentage points. The treatment $B_{m(i),t}$ is the cumulative storage energy capacity in the generator's county $m(i)$, in units of 100 MWh. The cost c_{it} is marginal cost in \$/MWh. I center c_{it} at the mean of its technology on interior hours, so ψ_h measures the effect of storage on a generator with typical cost for its technology. The generator fixed effect α_i absorbs permanent differences across generators, including the pre-entry level A_h^0 . The load-zone-by-month-by-hour fixed effect $\nu_{\ell(i),y(t),h}$ absorbs regional shifts in load and fringe supply. The controls $\mathbf{X}_{iht}$ are system net load, a cubic in local temperature, and local renewable capacity and potential output: county solar and wind capacity, and their products with local irradiance and a normalized wind power curve.

The coefficients map one to one into the model. The coefficient ψ_h equals $\frac{1}{2}(a_h - e_h \bar{c}_k)$, which is the output effect in Equation (7) evaluated at the mean cost $\bar{c}_k$ of technology k . The model

predicts $\psi_h < 0$ in discharge hours when the level channel dominates and $\psi_h > 0$ in charge hours. The coefficient ϕ_h equals $-\frac{1}{2}e_h$, the cross-derivative of output in storage and cost, and the model predicts $\phi_h \leq 0$. The coefficient ρ equals $-\frac{1}{2}E_h^0$, half the pre-entry slope of residual demand, and the model predicts $\rho < 0$.

I cluster standard errors at the county level. I estimate the equation pooled, by technology type $k \in \{\text{Combined Cycle, Simple Cycle, or Gas Steam Turbine}\}$, and by technology and hour. Some estimates rely on as few as 18 county clusters, so I construct confidence intervals with critical values from a t distribution with $G - 1$ degrees of freedom, where G is the number of counties behind each estimate (Bester, Conley and Hansen, 2011; Colin Cameron and Miller, 2015).

Interior hours are hours in which the generator runs strictly between its low and high sustained limits. The first-order condition in Equation (4) holds with equality only in the interior. At the low sustained limit, a committed generator’s output reflects the startup and minimum-load costs of staying online (Jha and Leslie, 2025). At the high sustained limit, capacity binds.

Marginal cost is potentially endogenous. Realized heat rates depend on realized output, so the constructed cost is mechanically a function of output (Demirer and Karaduman, 2026). The constructed cost omits maintenance, startup, and ramping costs, so it is a noisy proxy. The identity of the marginal unit may also respond to the shocks in u_{iht} .¹⁸ I instrument c_{it} and $B_{m(i),t}c_{it}$ with z_{it} and $B_{m(i),t}z_{it}$, where $z_{it} = \text{HR}_i^{\text{pre}} \times P_t^{\text{gas}}$ is the product of the generator’s pre-entry heat rate and the Henry Hub gas price, centered within technology like cost. The instrument shifts cost strongly, because fuel cost dominates marginal cost and equals heat rate times gas price by construction. The pre-entry heat rate does not move with the current operating point, which breaks the mechanical link between cost and output.

Identification rests on four assumptions. Let W_{iht} collect the fixed effects and controls. First, the instruments must be relevant. Second, they must satisfy the exclusion restriction, $\mathbb{E}[z_{it}u_{iht} \mid W_{iht}] = 0$ and $\mathbb{E}[z_{it}B_{m(i),t}u_{iht} \mid W_{iht}] = 0$. Third, storage must be conditionally exogenous, $\mathbb{E}[B_{m(i),t}u_{iht} \mid W_{iht}] = 0$. Fourth, the generator must optimize in the interior.

Each assumption faces a specific threat. First, the Henry Hub price shifts the costs of all gas generators at once and may shift residual demand directly. Replacing the month fixed effects with load-zone-by-date-by-hour fixed effects absorbs this common shock. Second, storage enters as an exogenous regressor, so ψ_h is exposed to comparisons with already-treated counties and to siting on local trends. The two-stage design below addresses both. Third, storage can change which generators operate in the interior. I report how sensitive the estimates are to excluding hours near the operating limits, and I estimate the effect on operating states directly (Appendix Table D.10). Fourth, ρ and ϕ_h equal the residual demand parameters only under constant marginal cost. With increasing marginal cost, $MC = c + \kappa q$, the cost coefficient identifies $-E/(2 + E\kappa)$, which approaches $-1/\kappa$ when residual demand is flat. I therefore use ϕ_h to test the sign of the slope channel, and I measure the slope directly from offer curves in Section 5.2.

¹⁸Ganapati, Shapiro and Walker (2020) instrument estimated marginal costs with energy-price shocks to address measurement error, and Fabra and Reguant (2014) treat hourly marginal cost as endogenous because the identity of the marginal unit is set by unobserved supply and demand shocks.

I also estimate the effect of storage on output with the two-stage difference-in-differences design of Section 5.1. I carry the cost gradient from the structural model into this design by adding an interaction with predetermined cost,

$$(17) \quad y_{iht} = \sum_{h=0}^{23} \mathbf{1}\{\text{hour} = h\} (\beta_h B_{m(i),t} + \gamma_h \tilde{z}_{it} B_{m(i),t}) + \lambda \tilde{z}_{it} + \alpha_i + \eta_{\ell(i),t,h} + \mathbf{X}_{iht} \theta + \varepsilon_{iht},$$

where $\tilde{z}_{it}$ is the predetermined cost shifter z_{it} , centered within technology, and the fixed effects and controls are as in Equation (9). The coefficient β_h is the displacement at mean predetermined cost, and γ_h is the reduced-form analog of ϕ_h . I use predetermined cost rather than realized cost because realized cost responds to storage. The first stage estimates λ together with the fixed effects and control coefficients on never-treated and not-yet-treated observations. I cluster standard errors at the county level.

The outcome is the capacity factor, including hours with zero output, estimated by technology and weighted by capacity. Weighting makes β_h a capacity-weighted average of generator-level effects, so multiplying it by capacity and summing recovers the change in MWh. This is the displacement I carry into the welfare analysis. I also estimate the effect on indicators for running, running at minimum load, running in the interior, and running at maximum output, which show how storage changes the way generators operate. I test for pre-trends in output with the event study in Equation (11), with the predetermined cost shifter in the first stage.

6.2 Results

This section presents the estimates on the output of gas generators. First, I compare the structural estimates of Equation (16) with the two-stage difference-in-differences (2SDD) estimates of Equation (17). I then present the impact on hourly output by technology and on operating margins, and I compare the structural cost gradient with the measured residual demand slope from Section 5.2. I then estimate the aggregate displacement that informs the welfare analysis.

Table 4 presents the pooled estimates. Column (1) and Column (2) present the OLS and IV estimates of Equation (16) on interior hours. An additional 100 MWh of local storage lowers the capacity factor of a generator with an average marginal cost by 0.79 percentage points. The OLS and IV estimates are nearly identical, which suggests that simultaneity in the constructed marginal cost is modest. The instruments are highly relevant, yielding first-stage cluster-robust F -statistics of 327 for cost and 2,447 for the interaction, with partial R^2 of 0.40 and 0.81. The cost coefficient is negative consistent with the optimal output rule. The interaction term is also negative, so storage displaces higher-cost generators more.

Column (3) applies the two-stage estimator to the same interior sample. The storage coefficient is -0.55 (s.e. 0.17), about 70 percent of the IV estimate, and the confidence intervals overlap. The two-stage estimator uses only generators in not-yet-treated counties as controls, so it avoids the comparisons with already-treated generators that affect the storage coefficient in Equation (16). The interaction with the predetermined cost shifter is -0.013 (s.e. 0.003), close to the structural

Table 4: Local storage effect on capacity factor: structural and two-stage estimates

	Capacity factor (pp)			
	OLS (1)	IV (2)	2SDD (3)	2SDD (4)
Storage B (100MWh)	-0.785*** (0.291)	-0.790*** (0.292)	-0.553*** (0.165)	-0.904*** (0.266)
Cost $\tilde{c}$ (\$/MWh)	-0.070*** (0.019)	-0.070*** (0.017)		
Cost $\tilde{c} \times B$	-0.016*** (0.004)	-0.016*** (0.003)		
Predetermined cost $\tilde{z} \times B$			-0.013*** (0.003)	-0.004 (0.008)
Generator FE	✓	✓	✓	✓
Load zone-Month-Hour FE	✓	✓		
Load zone-Date-Hour FE			✓	✓
Local renewable capacity and potential	✓	✓	✓	✓
Local temperature	✓	✓	✓	✓
Sample	Interior	Interior	Interior	All hours
Weights	None	None	None	Capacity
Observations	3,770,307	3,770,307	3,705,777	13,472,250
Counties	49	49	49	49
First-stage F , $\tilde{c}$		327.3		
First-stage F , $\tilde{c} \times B$		2447.3		
Partial R^2 , $\tilde{c}$		0.397		
Partial R^2 , $\tilde{c} \times B$		0.810		

Notes. The outcome is the hourly capacity factor in percentage points. B is cumulative storage energy capacity in the generator's county in 100 MWh. Columns (1) and (2) estimate Equation (16), with marginal cost centered within technology and instrumented in Column (2) by the pre-entry heat rate times the Henry Hub price. Columns (3) and (4) use the two-stage difference-in-differences estimator, with storage interacted with predetermined cost, the pre-entry heat rate times the Henry Hub price, centered within technology. Interior hours have output strictly between the low and high sustained limits. Standard errors clustered at the county level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

interaction of -0.016 . This implies that higher-cost generators are displaced more under both designs. Column (4) is my preferred specification for the welfare analysis. It includes all hours, including hours with zero output and hours at an operating limit, and weights by capacity. Storage lowers the capacity factor by 0.90 percentage points (s.e. 0.27).

The two-stage estimate is more robust to the sample of operating hours. Appendix Table D.7 shows that excluding hours within 2 percent of capacity of either operating limit shrinks the IV estimate to -0.26 and makes it insignificant. The same restriction leaves the two-stage estimate at -0.44 and significant (Appendix Table D.9). Much of the IV estimate is driven by generators operating close to their limits. The IV estimate is otherwise stable across choice of fixed effect and sample weights. It is -0.81 when I include hours at the operating limits, -0.66 with capacity weights, and -0.81 with load-zone-by-date-by-hour fixed effects that absorb the common natural gas price shock.

Table 5 presents the welfare specification by technology. Storage lowers the capacity factor of combined cycle generators by 0.37 percentage points (s.e. 0.17), or about 3.4 MWh at their mean

capacity of 921 MW. The estimates for simple cycle (-0.93) and steam turbine (-0.09) generators are not statistically significant. The pooled estimate of -0.90 exceeds the combined cycle and steam turbine estimates because the pooled first stage shares fixed effects across technologies and so also compares generators of different types. For aggregation, I therefore use the technology-specific estimates.

Table 5: Local storage effect on capacity factor by technology


	Capacity factor (pp)			
	All gas (1)	Combined Cycle (2)	Simple Cycle (3)	Steam Turbine (4)
Storage B (100MWh)	-0.904*** (0.266)	-0.367** (0.172)	-0.931 (0.573)	-0.090 (0.496)
Predetermined cost $\tilde{z} \times B$	-0.004 (0.008)	-0.013 (0.012)	-0.007 (0.009)	0.011 (0.012)
Generator FE	✓	✓	✓	✓
Load zone-Date-Hour FE	✓	✓	✓	✓
Local renewable capacity and potential	✓	✓	✓	✓
Local temperature	✓	✓	✓	✓
Sample	All hours	All hours	All hours	All hours
Weights	Capacity	Capacity	Capacity	Capacity
Observations	13,472,250	3,540,612	7,155,074	2,424,701
Counties	49	32	24	18

Notes. Two-stage difference-in-differences estimates on all hours, weighted by nameplate capacity. Column (1) pools all gas generators and repeats Column (4) of Table 4. Standard errors clustered at the county level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Hourly output displacement Storage displaces gas output in the afternoon and evening. Figure 9 presents the hourly two-stage estimates by technology. Combined cycle output falls significantly from 1 PM to 10 PM, with the largest effect of 0.78 percentage points at 7 PM, or about 7.2 MWh at mean capacity. Simple cycle output falls significantly from 1 PM to 9 PM, with the largest effect of 2.06 percentage points at 4 PM, or about 2.8 MWh at mean capacity. None of the overnight estimates is significant, and none of the steam turbine estimates is significant in any hour. The confidence intervals use t_{G-1} critical values, where G is the number of counties behind each estimate. The displacement works mainly through fewer hours at maximum output rather than through shutdowns. For combined cycle generators, the probability of running at maximum output in the evening falls by 1.4 percentage points (Appendix Table D.10). Steam turbines instead run more often in the evening but further from their maximum output, which offsets their output loss and explains why their estimates are not significant.

The theoretical model predicts that storage displaces higher-cost generators more. Panel (b) of Figure 9 evaluates the displacement at the 25th and 75th percentiles of predetermined cost within each technology. Consistent with the prediction, the higher-cost simple cycle generator loses 0.6 to 0.8 percentage points more capacity factor than the lower-cost one from 1 PM to 8 PM, and the cost interaction is significant in these hours. The higher-cost combined cycle generator loses 0.1 to

Figure 9: Hourly displacement by technology



Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the capacity factor, in percentage points, by hour and technology, on all hours and weighted by capacity. Panel (a) evaluates the effect at the mean predetermined cost of each technology. Panel (b) evaluates it at the 25th and 75th percentiles of predetermined cost. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level. The shaded band marks the evening peak.

0.3 percentage points more in every hour, although the interaction is not significant hour by hour. Steam turbines show the opposite gradient, which is significant only in a few overnight and morning hours. Event studies by technology show no pre-trends in output (Appendix Figure D.17).

I perform two robustness checks that support a causal reading of the afternoon displacement. First, the local renewable controls matter little for the afternoon estimates (Appendix Figure D.18). They lower the combined cycle estimates by about 0.1 percentage points in every hour and remove most of the overnight simple cycle displacement. The afternoon simple cycle displacement changes little, from 2.3 to 2.0 percentage points at 4 PM. This implies that the afternoon displacement is unlikely to be driven by solar capacity that enters alongside storage. Second, the displacement is larger in the Houston and North load zones, where it reaches 1.5 percentage points in the evening,

than in the wind-rich West and South zones. It is also larger on days with high solar or wind output, when arbitrage is more valuable (Appendix Figure D.19).


Comparison with the measured residual demand slope The residual demand slope implied by the structural estimates and the two-stage estimates on the measured slope in Figure 5 agree in sign. The negative interaction implies a flatter residual demand for combined cycle and simple cycle generators. The measured slope agrees for combined cycle generators overnight and in the evening, where the effect of 7.5 MW per \$/MWh is significant, and for simple cycle generators at midday and in the evening (Appendix Figure D.9). In magnitude, the slope implied by the cost coefficient is about 1.7 MW per \$/MWh for a combined cycle generator at mean capacity, roughly three orders of magnitude below the measured slope of 1,300 MW per \$/MWh. The static model with constant marginal cost cannot explain this discrepancy. With increasing marginal cost, $MC = c + kq$, the optimal output becomes $q^* = (A - Ec)/(2 + Ek)$, and when residual demand is flat the cost response approaches $-1/k$. The cost coefficient then reflects the curvature of the generator’s own cost, factoring in heat rate curvature and ramping and minimum load constraints, rather than the residual demand alone. Therefore, I use the measured slope for the slope channel and interpret the structural interaction only for its sign.

Counterfactual output I construct counterfactual output without storage for each generator and hour as observed output minus the estimated effect of local storage. The counterfactual uses the unit-specific effect from the welfare specification, which lets displacement vary with each generator’s predetermined cost. Output without storage cannot be negative or exceed capacity, so I bound the counterfactual at these limits. The lower bound binds when a generator is off and the estimate implies that storage raised its output, which would make output without storage negative. The upper bound almost never binds. The lower bound binds in 30 percent of generator-hours with local storage for steam turbines, 7 percent for simple cycle units, and about 1 percent for combined cycle units. It raises displacement from 811 GWh to 849 GWh per year, mostly among steam turbines. Because the bound makes the totals nonlinear in the coefficients, I compute standard errors by simulation. I draw the hourly coefficients 500 times from their estimated joint normal distribution, recompute the bounded totals in each draw, and report the standard deviation across draws (Krinsky and Robb, 1986). Appendix Table D.12 reports the unbounded totals.

Aggregate displacement Local storage displaced about 849 GWh of gas output per year over the sample (s.e. 305). Of the total, 371 GWh came from combined cycle generators and 440 GWh from simple cycle generators, which are 0.25 and 2.84 percent of their observed output. Each is significant at the 10 percent level, and the total is significant at the 5 percent level. The steam turbine estimate of 38 GWh is not distinguishable from zero, and the bound adds 27 GWh of it, so I do not interpret it. Figure 10 shows that the displaced output is concentrated in the afternoon and evening. Average installed storage in counties with gas generator was 1,170 MWh, so the estimates imply about 73 GWh of displaced gas output per 100 MWh of storage per year.

The point estimate is large relative to what storage supplies. Storage in counties with gas generator discharged 196 GWh per year, about 0.46 full cycles per day, so the estimated displacement is more than four times its discharge. Even the lower bound of the 95 percent confidence interval, about 250 GWh per year, exceeds that discharge. Two mechanisms can explain this. First, storage might relieve local congestion and allow imported energy to replace local gas output. Gas output might also shift to counties without storage, which the county-level design cannot separate from a true reduction. I therefore treat the aggregate as an upper bound on the storage-induced gas output displacement. The welfare analysis in Section 7 values storage by its observed operations rather than by this displacement for the same reason.

Figure 10: Gas output displaced by local storage, by hour



Notes. Annual displaced output, in GWh, from the technology-specific two-stage estimates by hour, on all hours and weighted by capacity. Counterfactual output is observed output net of the estimated storage contribution, bounded between zero and capacity.

Summary of displacement results Overall, both empirical strategies provide consistent evidence that battery storage displaces natural gas output during peak hours. The structural estimates confirm that generator output is responsive to marginal costs, consistent with standard optimal bidding behavior, and that storage disproportionately displaces higher-cost units. The two-stage estimates confirm the displacement and the cost gradient, and they remain robust even when excluding observations near generators’ operational limits. Furthermore, the negative cost interaction agrees in sign with the measured flattening of residual demand in Section 5.2, consistent with the slope channel. For the subsequent welfare analysis, I rely on the technology-specific, capacity-weighted two-stage estimates of the capacity factor. Because this specification holds the generator-hour population fixed and retains observations at zero output or physical limits, it correctly aggregates to the total change in output (MWh) that determines avoided emissions and the

comparison measure of production cost.

7 Welfare and Distributional Effects

This section asks who gains and loses from the competitive effect of storage, and whether storage raises total surplus. I first set out an accounting identity that separates transfers from changes in total surplus. I then measure each component for the storage installed in counties with gas generator, the storage behind the estimates in Sections 6.2 and 5.2. All amounts are annual averages over the sample, in 2020 dollars, the dollar year of the AP4 damages and the social cost of carbon. I deflate prices, costs, and profits by year with the annual average CPI-U ([U.S. Bureau of Labor Statistics, 2025](#)) before estimation, so the estimated effects are also in 2020 dollars.

7.1 Welfare Accounting

Load is inelastic, so storage changes total surplus only by changing the cost of serving load and the damages from emissions. I write the change in total surplus as

$$(18) \quad \Delta W = V^S + \Delta \text{Damage}^{\text{local}} + \Delta \text{Damage}^{\text{CO}_2} - K,$$

where V^S is the value of the energy storage moves across hours, ΔDamage are avoided damages, and K is the annualized capital cost of storage.

The same inelastic load implies that any change in prices is a transfer. When storage lowers the price, consumers pay less, and generators and storage owners receive less, by the same amount. The value of energy shifting is therefore split among four groups,

$$(19) \quad V^S = \Delta CS + \Delta \pi^{\text{gas}} + \pi^S + \Delta \pi^{\text{other}},$$

where ΔCS is the gain to consumers, $\Delta \pi^{\text{gas}}$ is the change in gas operating profit, π^S is the arbitrage revenue of storage owners, and $\Delta \pi^{\text{other}}$ is the change in the profit of the remaining generators. Equation (19) is the key to the precision of the transfer results. I measure V^S and π^S from observed storage operations, so they carry no sampling error. The sum $\Delta CS + \Delta \pi^{\text{other}} = V^S - \Delta \pi^{\text{gas}} - \pi^S$ is therefore as precise as the estimated loss of gas producers, even though its split between consumers and other generators is not.

7.2 Measurement

Value of energy shifting I value storage operations at the social marginal cost of energy rather than at market prices. Let s_{mht} be the net discharge of storage in county m at hour h of date t , positive when storage discharges. I compute

$$(20) \quad V^S = \sum_{m,h,t} \tilde{p}_{mht} s_{mht},$$

where $\tilde{p}_{mht}$ is the marginal cost of the energy that storage displaces when it discharges and draws on when it charges. V^S therefore measures the avoided production cost, net of round-trip losses. I use observed storage operations rather than the estimated displacement. Local storage displaced about 849 GWh of gas output per year, but storage discharged only 196 GWh, and part of the difference may be gas output that moved to counties without storage (Section 6.2).

The marginal source of energy differs across hours. When a gas generator sets the nodal price, I use the competitive price $p_{n(m)ht}^c$ of Section 5.2, the marginal cost of the last gas generator in a cost-based dispatch of realized gas output. When gas does not set the price, the marginal source is usually renewable, with a marginal cost near zero. In these hours I use the nodal price, floored at zero and capped at the gas benchmark:

$$(21) \quad \tilde{p}_{mht} = g_{mht} p_{n(m)ht}^c + (1 - g_{mht}) \min\{p_{n(m)ht}^c, \max\{\lambda_{mht}, 0\}\},$$

where g_{mht} is the share of the county's gas generators that ERCOT identifies as setting the price, λ_{mht} is the average nodal price at the county's gas nodes, and $n(m)$ is the local market that contains county m . The floor removes negative prices, which reflect renewable production credits, a transfer rather than a cost. The cap prevents the markups of other price setters from counting as costs. As a comparison, I also report V^S with the gas benchmark in every hour, and the avoided production cost implied by the local displacement, with replacement output valued at the competitive price.

Transfers Gas producers lose the operating profit $\Delta\pi^{\text{gas}}$ estimated in Section 5.2. Storage owners earn arbitrage revenue $\pi^S = \sum_{m,h,t} \lambda_{mht}^S s_{mht}$, where λ_{mht}^S is the nodal price at the storage node. Consumers gain from lower prices. I estimate the effect of storage on each county's daily nodal price with the two-stage design and apply it to local load,

$$(22) \quad \Delta CS = -\hat{\beta}^p \sum_{m,t} B_{mt} L_{mt},$$

where $\hat{\beta}^p$ is the effect of 100 MWh of local storage on the county's daily price, B_{mt} is installed storage capacity in county m on date t , and L_{mt} is the county's load on that date. The daily price averages the nodal prices at all non-storage nodes in the county and weights the hours of the day by the county's own hourly load. I observe hourly county load through 2023 and project 2024 load from each county's 2023 load profile, scaled by ERCOT load. Counties without gas generator join the control group until they receive storage, which raises the number of counties to 133. I weight counties equally, because the load-weighted estimate shows a price trend before entry in the largest counties (Appendix D.4). The remaining generators take the residual of Equation (19), $\Delta\pi^{\text{other}} = V^S - \Delta CS - \Delta\pi^{\text{gas}} - \pi^S$.

Avoided damages I compute avoided emissions by multiplying the displaced output of each generator-hour by that generator's weekly emission rates of NO_x , SO_2 , and CO_2 from CEMS, and a $\text{PM}_{2.5}$ factor from [U.S. Environmental Protection Agency \(2020\)](#). I value local pollutants with

the AP4 source-receptor matrices, which give the damage in each receptor county per ton emitted at each source (Tschofen, Azevedo and Muller, 2019). I match each generator to its AP4 source by plant code. Fifteen plants, holding 5 percent of gas capacity, have no AP4 source and take the source of the nearest plant of the same technology, in the same county when possible. I value CO₂ at a social cost of carbon of \$190 per metric ton (EPA, 2023). Because the damages follow the estimated change in gas output in every hour, any gas output that storage draws on when charging is netted out. They use the local displacement, so they share its caveat: if gas output moved to counties without storage, they overstate the system-wide reduction. I therefore treat them as upper bounds. The displacement is bounded so that output without storage lies between zero and capacity, and its standard errors come from the simulation in Section 6.2.

Storage cost The investment cost comprises a power cost of 1,440 \$/kW and an energy cost of 461.08 \$/kWh (EIA, 2024). I apply these unit costs to the average installed capacity of storage in counties with gas generator over the sample, 759 MW and 1,170 MWh, an average duration of 1.54 hours, so costs and benefits refer to the same capacity. Applying a capital recovery factor for a 15-year asset life, following the NREL Annual Technology Baseline assumption for utility-scale lithium-ion storage, at a 7% real discount rate gives an annualized cost of about \$179 million per year.¹⁹

7.3 Welfare Results

Table 6 presents the welfare accounting. I discuss the transfers first, because they are the competitive effect this paper measures, and then the change in total surplus.

The competitive effect of storage moves surplus away from gas producers. Gas producers lose \$47.4 million per year (s.e. 16.1), of which the market power rent accounts for \$34.4 million to \$54.5 million. Storage owners keep \$5.8 million per year from energy arbitrage, or \$8 per kW-year. By Equation (19), consumers and the other generators together gain \$41.1 million per year, with the same standard error as the gas producers' loss. This is the precise statement of who gains: almost all of the surplus storage takes from gas producers goes to the buyers and other sellers of energy, not to storage owners.

The split between consumers and other generators is less precise. My estimate for consumers is \$71.6 million per year (s.e. 55.5), which leaves the other generators with a loss of \$30.5 million. A loss for other generators is plausible, because lower prices also reduce the revenue of generators that do not lose market power, such as wind and solar units selling at the same nodes. The consumer estimate implies about \$94,000 per MW of storage per year, close to the \$124,000 per MW per year that Kirkpatrick (2026) estimates for the cost of serving load in California. Because the consumer estimate is not statistically significant, I interpret only the combined transfer.

¹⁹I annualize using the capital recovery factor $CRF = r(1+r)^N / [(1+r)^N - 1]$, with asset life $N = 15$ and real discount rate $r = 0.07$, so that $CRF \approx 0.110$. The annualized cost excludes fixed operating costs.

Table 6: Welfare and distributional effects of local storage

	Annual (2020 \$ million)	Per 100 MWh of storage (2020 \$ million/year)
<i>Transfers among market participants</i>		
Gas producers	-47.4 (16.1)	-4.05 (1.38)
Competitive rent, central benchmark	7.1 (6.8)	0.61 (0.59)
Market power rent, central benchmark	-54.5 (15.2)	-4.66 (1.30)
Market power rent, price-setter benchmark	-34.4 (15.6)	-2.94 (1.34)
Storage owners	5.8	0.50
Consumers and other generators	41.2 (16.1)	3.52 (1.38)
Consumers	71.6 (55.5)	6.12 (4.75)
Other generators (residual)	-30.5	-2.61
<i>Change in total surplus</i>		
Value of storage energy shifting	-0.5	-0.04
Avoided local damages	13.2 (5.0)	1.13 (0.43)
Avoided CO ₂ damages	76.8 (28.7)	6.57 (2.45)
Total benefit, local damages only	12.7 (5.0)	1.09 (0.43)
Total benefit, including CO ₂	89.6 (33.2)	7.66 (2.84)
Annualized storage cost	-107.8	-9.21
Net benefit, including CO ₂	-18.2 (33.2)	-1.56 (2.84)
Net benefit, lower capital cost	13.0 (33.2)	1.11 (2.84)
<i>Comparison measures of production cost</i>		
Value of storage energy shifting, gas benchmark	-0.5	-0.05
Avoided production cost, local displacement	5.7 (2.2)	0.49 (0.18)
<i>Storage owner revenue and cost recovery</i>		
Storage revenue, energy arbitrage	5.8	0.50
Arbitrage revenue minus value of energy shifting	6.3	0.54
Storage revenue, ancillary services	35.6	3.04
Storage revenue, total	41.4	3.54
Share of annualized cost recovered	38%	
Share recovered, net of investment tax credit	42%	
Share recovered, lower capital cost	54%	
Net benefit, ancillary services at their price	17.4	1.48

Notes. Annual averages over 2018 to 2024 for storage in counties with gas generator, in 2020 dollars. Energy shifting is valued at the social price of energy (Equations (20) and (21)). Avoided damages apply CEMS emission rates and AP4 damages to the bounded local displacement, with CO₂ at \$190 per metric ton. Storage cost annualizes the EIA investment cost at 7 percent over 15 years. Other generators are the residual of the surplus identity in Equation (19). Ancillary service revenue does not enter total surplus. Per 100 MWh values divide by 1,170 MWh of installed storage. Standard errors in parentheses come from 500 draws of the estimated displacement coefficients for the damages and production cost rows, and from the clustered regressions for gas producers and consumers.

The private return to storage Storage earns its arbitrage revenue from the markups it erodes. The difference between what storage earns and what its operations are worth is

$$(23) \quad \pi^S - V^S = \sum_{m,h,t} (\lambda_{mht}^S - \tilde{p}_{mht}) s_{mht},$$

the gap between the market price and the social marginal cost of energy, weighted by storage operations. Storage discharges in the evening, when gas generators set the price above their cost, and charges when that gap is small. The same operations that earn \$5.8 million at market prices are worth $-\$0.5$ million at marginal cost, so the private return of storage arbitrage is entirely the markup of the marginal gas generators. As storage erodes those markups, it also erodes the price spread it earns. The private return to arbitrage therefore declines as storage disciplines gas generators.

Storage creates little surplus by shifting energy. Valued at the social price of energy, storage operations slightly raise production costs, by \$0.5 million per year. Storage charges overnight and in the morning and discharges in the evening peak (Figure 1). The cost spread between these hours is small, and storage loses about 19 percent of its energy to round-trip efficiency losses, so the two roughly cancel. About half of the charging occurs in hours when gas does not set the nodal price, typically when renewable output is on the margin. Valuing that energy at the renewable margin changes the result little, because the nodal price in these hours is close to the cost of the marginal gas generator. With the gas benchmark in every hour, the value is also $-\$0.5$ million. The comparison based on the local displacement, which also counts any reallocation of output from higher-cost to lower-cost gas generators, gives a gain of \$5.7 million per year (s.e. 2.2), still small.

Avoided damages are the largest measured benefit, but they rest on the local displacement. Avoided local damages total \$13.2 million per year (s.e. 5.0), and avoided CO₂ damages \$76.8 million per year (s.e. 28.7), for a total of \$89.6 million, or \$12.7 million without CO₂. These are upper bounds. The displacement is more than four times what storage discharged, and the damages assume that none of the displaced gas output moved elsewhere.

The measured benefits do not cover the capital cost of storage. The annualized capital cost is \$179 million per year, or \$15.3 million per 100 MWh of storage. Even with avoided damages at their upper bound, the net benefit is $-\$90$ million per year. This comparison omits the main service that ERCOT storage provides. Storage in counties with gas generator cycled only 0.46 times per day and earned \$35.6 million per year from ancillary services, 86 percent of its \$41.4 million total revenue. Ancillary services hold capacity in reserve for grid reliability, and their social value does not appear in energy-market accounts. Valuing them at their market prices, which likely overstates their value when reserve prices exceed the cost of the next available provider, raises the net benefit to $-\$54$ million per year. These estimates imply that energy arbitrage and its competitive effect alone do not justify the investment, not that the investment lowers welfare.

Ancillary services provide most of storage revenue. Together with arbitrage, they cover about 23 percent of the annualized capital cost of storage in counties with gas generator, or \$54 per kW-

year.²⁰ The gap is partly a feature of the sample, which excludes Winter Storms Uri and Elliott, when reserve and energy prices reached their caps. Storage in these counties earned \$82 million from ancillary services in these 12 days, about 22 percent of all its revenue over the sample period.

The avoided local damages accrue mostly outside the counties that host storage. Storage counties receive 29 percent of the avoided damages, the rest of Texas 19 percent, and other states 52 percent, because fine particles formed from Texas emissions travel downwind. Per resident, the benefit is concentrated in Texas: \$0.22 per person per year in counties with storage and \$0.19 in the rest of Texas, compared with \$0.02 outside the state (Appendix Figure D.21).

Summary of welfare results The competitive effect of storage is a transfer. Storage moved about \$47 million per year away from incumbent gas generators, most of it market power rent. Storage owners kept about \$6 million, and consumers and the other generators gained the remaining \$41 million, with most of it going to consumers by my estimates, although that split is imprecise. Storage created little value by shifting energy at marginal cost, and its measured benefits, even with avoided damages at their upper bound, fall short of its capital cost. Its private return from arbitrage comes from the same markups it disciplines.

8 Conclusion

This paper shows that battery storage disciplines the market power of natural gas generators. Using the staggered rollout of county-level storage capacity in ERCOT, I find that local storage makes the residual demand of nearby gas generators more elastic and displaces their output in the afternoon and evening. It lowers their evening markups and their daily operating profits by about 6 percent per 100 MWh of storage. Between 73 percent and all of the profit loss is market power rent, and gas generators continue to set the price at a smaller margin. These results support a simple mechanism: storage flattens the residual demand of a strategic generator and raises the cost of withholding output. The competitive effect of storage moved about \$47 million per year away from gas producers, most of it market power rent. Storage owners kept about \$6 million, and consumers and the other generators gained the rest. By contrast, storage created little value by shifting energy at marginal cost. Its arbitrage revenue comes from the markups it erodes, and it earns most of its revenue from ancillary services rather than energy. Even with ancillary services, storage revenue covers about a quarter of its annualized capital cost over the sample.

These results carry three policy implications. First, evaluations of storage should count its competitive effect. Standard cost-benefit analyses value storage by the production cost and emissions it avoids. In my setting, the production cost savings are small relative to the transfer away from gas producers. The avoided emissions are larger, but they rest on a local displacement that I treat as

²⁰The capital cost is the full resource cost of storage. The investment tax credit introduced by the Inflation Reduction Act, available to standalone storage from 2023, shifts part of this cost from storage owners to taxpayers and does not change total surplus. Net of the credit, storage revenue covers about 25 percent of the annualized cost to owners, compared with 23 percent before the credit, because most storage in the sample entered before standalone storage became eligible.

an upper bound. Omitting competition therefore understates what consumers gain from storage. Second, storage complements market power mitigation. ERCOT combines permissive offer caps with energy-only pricing, and storage lowers markups in the congested local markets where offer mitigation is hardest to target. Interconnection and siting policies that speed storage into these markets would deliver competitive benefits beyond the energy it supplies. Third, consistent with results from prior works ([Karaduman, 2023](#); [Butters, Dorsey and Gowrisankaran, 2025](#)), energy arbitrage alone cannot justify storage investment, because its revenue shrinks as storage erodes markups. The design and pricing of ancillary service and reliability products will therefore shape how much storage enters and, through entry, how much competitive discipline it provides.

My analysis has limitations that point to future work. First, the county-level design measures local effects. If gas output shifts to counties without storage, the local displacement and the avoided emissions overstate the system-wide effects. The welfare accounts also omit the reliability value of ancillary services. Second, I abstract away from storage ownership, and my estimates average across owners. My model implies that an owner of both storage and other generator discharges less, because discharging lowers the price its other generator receives. Storage owned by incumbent gas generators or by large renewable portfolios may therefore discipline markups less than storage owned by independent firms. Third, I do not model curtailment directly. When gas does not set the nodal price, I value the energy storage shifts at the nodal price floored at zero, which approximates the near-zero cost of renewable energy that would otherwise be curtailed. This approximation is coarse in congested hours, when the nodal price can exceed the cost of curtailed energy, so my estimates may understate the value of storage in those hours. Curtailment is concentrated in the West and South, where storage is common but gas generators are few, so this value likely falls largely outside my sample. Fourth, the estimates reflect storage at its current penetration. As storage grows, it may set prices more often and exercise market power of its own. Finally, the estimates come from an energy-only market with high offer caps and narrow local mitigation. Markets with capacity payments and tighter mitigation leave less market power for storage to discipline, so the mechanism should carry over, but the effect is likely smaller. Where storage also earns capacity payments, its revenue mix and the share of its capital cost it recovers would also differ.

References

- Ahmed, Mayesha, Caleb Crain, Shreya Maggo, and Nabila Safitri. 2026. “Replacing California’s Peaker Plants with Long-Duration Energy Storage.”
- Andrés-Cerezo, David, and Natalia Fabra. 2023. “Storing Power: Market Structure Matters.” *The RAND Journal of Economics*, 54(1): 3–53.
- Andres-Cerezo, David, and Natalia Fabra. 2025. “Storage and Renewable Energies: Friends or Foes?” *Working Paper*.
- Arcidiacono, Peter, Paul B. Ellickson, Carl F. Mela, and John D. Singleton. 2020. “The Competitive Effects of Entry: Evidence from Supercenter Expansion.” *American Economic Journal: Applied Economics*, 12(3): 175–206.
- Arkolakis, Costas, and Conor Walsh. 2024. “Clean Growth.” *NBER WORKING PAPER SERIES*.
- Bester, C. Alan, Timothy G. Conley, and Christian B. Hansen. 2011. “Inference with Dependent Data Using Cluster Covariance Estimators.” *Journal of Econometrics*, 165(2): 137–151.
- Borenstein, Severin, James B Bushnell, and Frank A Wolak. 2002. “Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market.” *American Economic Review*, 92(5): 1376–1405.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. “Revisiting Event-Study Designs: Robust and Efficient Estimation.” *The Review of Economic Studies*, 91(6): 3253–3285.
- Bushnell, James, and Kevin Novan. 2021. “Setting with the Sun: The Impacts of Renewable Energy on Conventional Generation.” *Journal of the Association of Environmental and Resource Economists*, 8(4): 759–796.
- Bushnell, James B, Erin T Mansur, and Celeste Saravia. 2008. “Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets.” *American Economic Review*, 98(1): 237–266.
- Busso, Matias, and Sebastian Galiani. 2019. “The Causal Effect of Competition on Prices and Quality: Evidence from a Field Experiment.” *American Economic Journal: Applied Economics*, 11(1): 33–56.
- Butters, R. Andrew, Jackson Dorsey, and Gautam Gowrisankaran. 2025. “Soaking up the Sun: Battery Investment, Renewable Energy, and Market Equilibrium.” *Econometrica*, 93(3): 891–927.
- C3S. 2018. “ERA5 Hourly Data on Single Levels from 1940 to Present.”
- Callaway, Brantly, and Pedro H.C. Sant’Anna. 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics*, 225(2): 200–230.
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H C Sant’Anna. 2024. “Difference-in-Differences with a Continuous Treatment.”
- Colin Cameron, A., and Douglas L. Miller. 2015. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources*, 50(2): 317–372.
- De Chaisemartin, Clément, and Xavier D’Haultfœuille. 2020. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” *American Economic Review*, 110(9): 2964–2996.
- Demirer, Mert, and Ömer Karaduman. 2026. “Do Mergers and Acquisitions Improve Efficiency? Evidence from Power Plants.” *Journal of Political Economy*.
- Doshi, Gaurav. 2025. “Wiring America: Electricity Grid Expansion and Fossil Fuel Market Power.”
- Dudek, Matthew. 2026. “Smoothing the Green Transition: Which Battery Locations Provide

- the Greatest Returns to Renewable Generation in California.” *Working Paper*.
- EIA.** 2024. “EIA - U.S. Battery Storage Market Trends.” <https://www.eia.gov/analysis/studies/electricity/batterystorage/>.
- EIA.** 2025. “Annual Electric Power Industry Report, Form EIA-860.”
- Ellison, Glenn, and Sara Fisher Ellison.** 2011. “Strategic Entry Deterrence and the Behavior of Pharmaceutical Incumbents Prior to Patent Expiration.” *American Economic Journal: Microeconomics*, 3(1): 1–36.
- EPA.** 2023. “EPA Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances.” EPA.
- ERCOT.** 2024. “ERCOT Nodal Protocols, Section 6: Adjustment Period and Real-Time Operations.”
- ERCOT.** 2025. “Fact Sheet.”
- Fabra, Natalia, and Imelda.** 2023. “Market Power and Price Exposure: Learning from Changes in Renewable Energy Regulation.” *American Economic Journal: Economic Policy*, 15(4): 323–358.
- Fabra, Natalia, and Mar Reguant.** 2014. “Pass-Through of Emissions Costs in Electricity Markets.” *American Economic Review*, 104(9): 2872–2899.
- Fell, Harrison, Daniel T. Kaffine, and Kevin Novan.** 2021. “Emissions, Transmission, and the Environmental Value of Renewable Energy.” *American Economic Journal: Economic Policy*, 13(2): 241–272.
- Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker.** 2020. “Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes.” *American Economic Journal: Applied Economics*, 12(2): 303–342.
- Gardner, John, Neil Thakral, Linh T Tô, and Luther Yap.** 2024. “Two-Stage Differences in Differences.”
- Goodman-Bacon, Andrew.** 2021. “Difference-in-Differences with Variation in Treatment Timing.” *Journal of Econometrics*, 225(2): 254–277.
- Goolsbee, Austan, and Chad Syverson.** 2008. “How Do Incumbents Respond to the Threat of Entry? Evidence from the Major Airlines*.” *Quarterly Journal of Economics*, 123(4): 1611–1633.
- Hausman, Catherine.** 2025. “Power Flows: Transmission Lines, Allocative Efficiency, and Corporate Profits.” *American Economic Review*.
- Hortaçsu, Ali, and Steven 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.
- Ito, Koichiro, and Mar Reguant.** 2016. “Sequential Markets, Market Power, and Arbitrage.” *American Economic Review*, 106(7): 1921–1957.
- Jha, Akshaya, and Gordon Leslie.** 2025. “Start-Up Costs and Market Power: Lessons from the Renewable Energy Transition.” *American Economic Review*, 115(2): 690–724.
- Joskow, Paul L.** 2022. “From Hierarchies to Markets and Partially Back Again in Electricity: Responding to Decarbonization and Security of Supply Goals.” *Journal of Institutional Economics*, 18(2): 313–329.
- Karaduman, Ömer.** 2023. “Economics of Grid-Scale Energy Storage in Wholesale Electricity Markets.” *Working Paper*.
- Kirkpatrick, A. Justin.** 2026. “Estimating the Congestion Benefits of Batteries When Network Connections Are Unobserved.” *The Energy Journal*, 01956574261469461.
- Klemperer, Paul D., and Margaret A. Meyer.** 1989. “Supply Function Equilibria in Oligopoly under Uncertainty.” *Econometrica*, 57(6): 1243–1277.
- Krinsky, Itzhak, and A. Leslie Robb.** 1986. “On Approximating the Statistical Properties of

- Elasticities.” *The Review of Economics and Statistics*, 68(4): 715–719.
- Lamp, Stefan, and Mario Samano.** 2022. “Large-Scale Battery Storage, Short-Term Market Outcomes, and Arbitrage.” *Energy Economics*, 107: 105786.
- LaRiviere, Jacob, and Xueying Lyu.** 2022. “Transmission Constraints, Intermittent Renewables and Welfare.” *Journal of Environmental Economics and Management*, 112: 102618.
- Mathrani, Amit.** 2025. “Texas: A High Stakes Frontier for US Battery Energy Storage Systems.” <https://www.rabobank.com/knowledge/d011484585-texas-a-high-stakes-frontier-for-us-battery-energy-storage-systems>.
- Mercadal, Ignacia.** 2022. “Dynamic Competition and Arbitrage in Electricity Markets: The Role of Financial Players.” *American Economic Journal: Microeconomics*, 14(3): 665–699.
- Miller, David, and Mia Nakajima.** 2024. “Differences in Battery Participation in the CAISO and ERCOT Markets - Gridmatic.” <https://www.gridmatic.com/differences-in-battery-participation-in-the-caiso-and-ercot-markets/>.
- Modo Energy.** 2026. “ERCOT Annual Buildout Report: Battery Capacity Reaches 14 GW Entering 2026 - Research.” <https://modoenergy.com/research/en/ercot-battery-buildout-2025-annual-report>.
- NOAA.** 2024. “Global Surface Hourly (Integrated Surface Database).”
- NREL.** 2024. “National Solar Radiation Database (NSRDB).”
- Potomac Economics.** 2025. “2024 STATE OF THE MARKET REPORT FOR THE ERCOT ELECTRICITY MARKETS.”
- Reaser, Robert, and Reid B. Taylor.** 2025. “The Short-Run Impacts of Grid-Scale Storage Entry in the Texas Electricity Market.” *Working Paper*.
- Reguant, Mar.** 2014. “Complementary Bidding Mechanisms and Startup Costs in Electricity Markets.” *The Review of Economic Studies*, 81(4): 1708–1742.
- Reynolds, Stanley S.** 2025. “Keeping the Lights On: Battery Storage, Operating Reserves, and Electricity Scarcity Pricing.” *Working Paper*.
- Ryan, Nicholas.** 2021. “The Competitive Effects of Transmission Infrastructure in the Indian Electricity Market.” *American Economic Journal: Microeconomics*, 13(2): 202–242.
- Sun, Liyang, and Sarah Abraham.** 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” *Journal of Econometrics*, 225(2): 175–199.
- Taylor, Reid, and Erich Muehlegger.** 2026. “The Effects of Competition in the Retail Gasoline Industry.” *Journal of the Association of Environmental and Resource Economists*.
- Tschofen, Peter, Inês L. Azevedo, and Nicholas Z. Muller.** 2019. “Fine Particulate Matter Damages and Value Added in the US Economy.” *Proceedings of the National Academy of Sciences*, 116(40): 19857–19862.
- U.S. Bureau of Labor Statistics.** 2025. “Historical Consumer Price Index for All Urban Consumers ({CPI-U}): U.S. City Average, All Items.” <https://www.bls.gov/cpi/tables/supplemental-files/historical-cpi-u-202412.pdf>.
- U.S. Environmental Protection Agency.** 2020. “Estimating Particulate Matter Emissions for eGRID.” Clean Air Markets Division Draft White Paper.
- Vermillion, Brandt.** 2024. “Battery Energy Storage in ERCOT’s Real-Time Market: What’s Changing? - Research.” <https://modoenergy.com/research/en/ercot-battery-energy-storage-systems-real-time-prices-offers-owners-eolian-engie-jupiter-power>.
- Vermillion, Brandt.** 2025. “ERCOT: Annual Battery Energy Storage Revenue Report 2024 - Research.” <https://modoenergy.com/research/en/ercot-battery-energy-storage-revenues-energy-markets-ancillary-services-capacity-location-owners-operators-annual-report-2024>.
- Visaggio, Pietro.** 2025. “Pairing Batteries with Renewables: How Ownership Shapes Operational Incentives and Market Outcomes.” *Working Paper*.

- Wilson, Robert.** 1979. “Auctions of Shares.” *The Quarterly Journal of Economics*, 93(4): 675.
- Woerman, Matt.** 2023. “Market Size and Market Power: Evidence from the Texas Electricity Market.” *Working Paper*.
- Yitzhaki, Shlomo.** 1996. “On Using Linear Regressions in Welfare Economics.” *Journal of Business & Economic Statistics*.

Supplementary Appendix

A Data Appendix

This appendix describes how I construct two variables that the main analysis relies on: the local market of each generator and the slope of the residual demand it faces. For each, I state the rule I apply to the data and then walk through a stylized example. The numbers in the examples are illustrative and do not come from a specific interval.

A.1 Local Market Definition

I define local markets separately for every 15-minute interval. I index intervals by $r \in \{1, 2, 3, 4\}$ within hour h of day t , so interval-level variables carry the subscript thr and average up to the hourly subscript th of the main text. Congestion changes from one interval to the next, so the boundaries of the markets it creates also change. A local market is a set of plants that face similar nodal prices and that the transmission network connects. I build it in three steps, following [Mercadal \(2022\)](#) and [Woerman \(2023\)](#).

The unit for which the market is assigned is at the plant level. The set of plants includes every generation resource in the SCED report, including wind, solar, and storage, not only gas generators. I assign each plant the nodal price at its settlement point, averaged from the 5-minute SCED prices to the 15-minute interval. All resources of a plant share its local market.

Step 1: price groups. I first link plants whose nodal prices are close. Two plants j and k are price neighbors in interval thr if

$$(A.1) \quad |p_{jthr} - p_{kthr}| \leq \max\{1, 0.01 \cdot \max(|p_{jthr}|, |p_{kthr}|)\}.$$

The absolute threshold of \$1/MWh binds when both prices are below \$100/MWh. The relative threshold of 1 percent binds at higher prices, so the rule does not split plants on small price gaps during scarcity. I then form a graph in which each edge links two price neighbors, and I define each connected component of this graph as a price group. Two plants can therefore belong to the same group even if their prices differ by more than the threshold, as long as a chain of price neighbors links them. A plant with no price neighbor forms its own group.

Step 2: network connectivity. I next check that the plants in each price group are physically connected. I build the transmission network from the line data of [Arkolakis and Walsh \(2024\)](#), keeping the lines that cross or touch Texas. Each line is an edge, and its two endpoints are nodes. In the data, the nodes consists of substations in Texas. I attach each plant to its three nearest lines within 30 km by linking the plant to both endpoints of each line. Two isolated plants have no line within this radius, and I connect them to the nearest line manually. I then split each price group

by the connected components of this network. Two plants in the same price group share a local market only if a path of transmission lines connects them.

Step 3: market identifier. Each set that survives both steps is a local market. I denote by $n(i)$ the local market of plant i in interval thr , as in the main text. The identifier is specific to the interval and carries no meaning across intervals. I therefore never use it as a fixed effect. I use it only to define the set of rivals a generator faces within an interval and to compute market size, the total capacity of the plants in the market. Across the sample, the number of local markets ranges from 1 to 137 per interval, with a median of 8 (Figure C.3).

Sylized Example. Table A.1 shows how the rule assigns eight plants in one interval. Plants A to D sit at prices between \$25.00 and \$26.90/MWh. Each adjacent pair is within \$1/MWh, so the four plants form one price group. A and D differ by \$1.90/MWh, which exceeds the threshold, but the chain through B and C links them. Plant H sits at \$31.00/MWh and has no price neighbor, so it forms its own group. Plants E to G sit above \$150/MWh, where the 1 percent rule binds. The threshold between E and F is \$1.51/MWh, and their gap of \$1.20/MWh falls within it. F and G are also neighbors, so E, F, and G form one price group. The network then splits this group, because G sits on a part of the network that no line connects to E and F. The interval therefore has four local markets. Figure C.4 shows the same rule applied to two actual intervals 15 minutes apart.

Table A.1: Stylized example of the local market definition

Plant	Nodal price (\$/MWh)	Gap to previous plant (\$/MWh)	Threshold (\$/MWh)	Price group	Network component
A	25.00			1	1
B	25.60	0.60	1.00	1	1
C	26.40	0.80	1.00	1	1
D	26.90	0.50	1.00	1	1
H	31.00	4.10	1.00	2	1
E	150.00			3	1
F	151.20	1.20	1.51	3	1
G	152.40	1.20	1.52	3	2
Local markets: $\{A, B, C, D\}$, $\{H\}$, $\{E, F\}$, $\{G\}$					

Notes. Plants are sorted by nodal price within each price range. The gap is the price difference to the plant in the row above. The threshold is $\max\{1, 0.01 \cdot \max(|p_j|, |p_k|)\}$ from Equation (A.1). A price group is a connected component of the graph that links price neighbors. A local market is the intersection of a price group and a network component.

A.2 Residual Demand Slope

I measure the slope of each gas generator's residual demand as the density of its local rivals' offers around its nodal price. In the model, the generator faces residual demand $q_{ih} = A_h - E_h p_h$ with slope $E_h = \gamma + \sigma_h$, the sum of the fringe slope and the storage slope (Equation (3)). Load is

inelastic in real time, so the slope of residual demand equals the slope of rival supply. My measure estimates this slope directly from the offers that rivals submit to SCED, following [Hortaçsu and Puller \(2008\)](#).

Each resource submits an offer curve in every interval as a set of price-quantity points on a monotone step function. I convert each curve into incremental blocks. Block b of resource j in interval thr has quantity q_{jbthr} , the increase in cumulative MW from the previous point, and price p_{jbthr} . I drop blocks with zero or negative increments. The supply stack includes the offers of every generation resource, of all fuel types. Storage offers enter the stack only in intervals when the battery discharges, defined as a positive net base point, because a charging battery does not supply energy in that interval.

For gas generator i in local market $n(i)$ at nodal price p_{ithr} , I compute

$$(A.2) \quad E_{ithr} = \sum_{j \in n(i), j \neq i} \sum_b q_{jbthr} K_{\Delta}(p_{jbthr} - p_{ithr}), \quad K_{\Delta}(u) = \frac{3}{4\Delta} \left(1 - \frac{u^2}{\Delta^2}\right) \mathbf{1}\{|u| \leq \Delta\},$$

where K_{Δ} is the Epanechnikov kernel with half-width $\Delta = \$5/\text{MWh}$. The kernel integrates to one over $[-\Delta, \Delta]$, so E_{ithr} is measured in MW per $\$/\text{MWh}$. Equivalently, E_{ithr} is the slope at p_{ithr} of the smoothed rival supply curve $\tilde{S}_{-i}(p) = \sum_{j \in n(i), j \neq i} \sum_b q_{jbthr} F_{\Delta}(p - p_{jbthr})$, where F_{Δ} is the cumulative distribution function of K_{Δ} . The smoothing replaces each jump in the offer stack with a gradual ramp over $\$10/\text{MWh}$, so the stack has a well-defined slope at every price. With a uniform kernel, E_{ithr} would equal the rival MW offered within $\$5/\text{MWh}$ of p_{ithr} divided by the $\$10/\text{MWh}$ window. The Epanechnikov kernel gives more weight to blocks close to the price and less to blocks near the edge of the window. A larger E_{ithr} means that a one-dollar increase in price draws in more rival supply, so withholding costs the generator more sales.

I evaluate the slope at each generator's own nodal price. I first compute the kernel density of the full local stack, including the generator's own offer, and then subtract the density of its own blocks. I set the rare negative differences to zero. I compute E_{ithr} for gas generators (simple cycle, combined cycle, and gas steam) in intervals when they produce. To build the hourly outcome used in the main text, I average across the intervals within the hour in which the generator produces,

$$E_{ith} = \frac{1}{|\mathcal{R}_{ith}|} \sum_{r \in \mathcal{R}_{ith}} E_{ithr},$$

where $\mathcal{R}_{ith}$ is the set of intervals in hour h of day t with positive output. I then winsorize E_{ith} at the 99th percentile within technology.

Stylized Example. Table [A.2](#) and Figure [A.1](#) walk through the calculation for one gas generator. The generator produces at a nodal price of $\$40/\text{MWh}$ and offers its own 50 MW block at $\$40/\text{MWh}$. A 1,000 MW block of inframarginal rivals at $\$22/\text{MWh}$ lies far below the window and receives no weight. Four rival generator blocks fall within the $\$35$ to $\$45/\text{MWh}$ window. Their kernel weights are 0.144 for the blocks $\$1/\text{MWh}$ away and 0.054 for the blocks $\$4/\text{MWh}$ away. A fifth block at

\$47/MWh lies outside the window and receives no weight. The rival slope is

$$E_{ithr} = 100(0.054) + 200(0.144) + 150(0.144) + 300(0.054) = 72.0 \text{ MW per } \$/\text{MWh}.$$


A uniform kernel over the same window gives $750/10 = 75.0$. A 500 MW block at \$40/MWh owned by a plant in a different local market does not enter E_{ithr} , because congestion separates it from the generator. A 250 MW battery block at \$42/MWh receives a weight of 0.126 and adds 31.5 MW per \$/MWh when the battery discharges. The slope then rises to 103.5, or 44 percent. When the battery charges, its block does not enter and the slope stays at 72.0. Panel (a) of Figure A.1 plots the rival offer stack as a step function together with its smoothed version $\tilde{S}_{-i}(p)$. The slope of the smoothed curve at \$40/MWh is the 72.0 in the table. Panel (b) plots the generator's residual demand, load minus the smoothed rival stack, with and without the discharging battery. The battery shifts residual demand inward and flattens it at the generator's price. This is the slope channel in the data: discharging storage places price-responsive supply near the generator's price and flattens its residual demand.

Table A.2: Stylized example of the residual demand slope

Block	Owner	Offer price (\$/MWh)	Quantity (MW)	Kernel weight $K_{\Delta}(p_{jb} - p_i)$	Contribution (MW per \$/MWh)
<i>Rivals in the same local market</i>					
R0	Inframarginal rivals	22	1,000	0	0
R1	Gas generator	36	100	0.054	5.4
R2	Gas generator	39	200	0.144	28.8
R3	Gas generator	41	150	0.144	21.6
R4	Gas generator	44	300	0.054	16.2
R5	Gas generator	47	400	0	0
E_{ithr} without storage					72.0
S	Storage, discharging	42	250	0.126	31.5
E_{ithr} with discharging storage					103.5
<i>Excluded from E_{ithr}</i>					
Own	Generator i	40	50	0.150	—
X	Other local market	40	500	0.150	—
S	Storage, charging	42	250	0.126	—

Notes. Generator i produces at a nodal price of $p_{ithr} = \$40/\text{MWh}$. The kernel is Epanechnikov with half-width $\Delta = \$5/\text{MWh}$, so $K_{\Delta}(u) = 0.15(1 - u^2/25)$ for $|u| \leq 5$ and zero otherwise. The contribution is quantity times kernel weight. In the data, I compute the density of the full local stack, here $72.0 + 50(0.150) = 79.5$, and subtract the generator's own density of 7.5.

Figure A.1: Stylized example of the residual demand slope



Notes: The example uses the offer blocks in Table A.2. Panel (a) plots the rival offer stack of generator i as a step function (gray) and its kernel-smoothed version $\tilde{S}_{-i}(p)$ (black). The dashed line is the tangent at the generator's nodal price $p_{ithr} = \$40/\text{MWh}$, whose slope is $E_{ithr} = 72.0$ MW per $\$/\text{MWh}$. The shaded band is the $\$5/\text{MWh}$ window around the nodal price. Panel (b) plots residual demand, a load of 1,700 MW minus the smoothed rival stack, without storage (gray) and with a 250 MW battery block discharging at $\$42/\text{MWh}$ (red). Dashed lines are tangents at the nodal price. The battery shifts residual demand inward and raises its slope from 72.0 to 103.5 MW per $\$/\text{MWh}$, which makes residual demand flatter.

B Detailed Theoretical Derivations

This appendix extends the model of Section 4 in three directions. It shows how transmission congestion determines how much of a battery's discipline reaches a gas generator, generalizes storage from a price taker to a strategic storage, and derives how ownership changes the competitive effect of storage. Throughout, I keep the notation of the main text: $h \in \{L, H\}$ indexes the period, i is the strategic gas generator with marginal cost c_i , output q_{ih} , and markup μ_{ih} , and the generator faces the residual demand $q_{ih} = A_h - E_h p_h$.

B.1 Congestion and the Local Market

Congestion determines how much of a remote battery's discipline reaches the gas generator, which is why I define the market locally. Consider two regions, a local region ℓ that holds generator i and a remote region r , joined by a line of capacity $\bar{F}$ that carries flow F_h from r to ℓ . Superscripts denote the region. Region ℓ holds generator i , a fringe $S_h^{c,\ell}(p) = \alpha_h^\ell + \gamma^\ell p$, and load D_h^ℓ . Region r holds a fringe $S_h^{c,r}(p) = \alpha_h^r + \gamma^r p$, load D_h^r , and a storage fleet with supply $s_h(p) = \bar{s}_h + \sigma_h p$ as in Equation (2). Nodal market clearing, in the form of Equation (1), and the line constraint are

$$(B.3) \quad (\ell) \quad S_h^{c,\ell}(p_h^\ell) + q_{ih} + F_h = D_h^\ell,$$

$$(B.4) \quad (r) \quad S_h^{c,r}(p_h^r) + s_h(p_h^r) = D_h^r + F_h,$$

$$(B.5) \quad F_h \leq \bar{F}, \quad p_h^\ell - p_h^r = \lambda_h \geq 0, \quad \lambda_h(\bar{F} - F_h) = 0,$$

so the price gap equals the congestion shadow price λ_h .

When the line is slack, remote storage reaches the generator in full. With $\lambda_h = 0$, both regions clear at one price p_h , and adding (B.3) and (B.4) eliminates the flow. The generator faces the residual demand $q_{ih} = A_h - E_h p_h$ with

$$(B.6) \quad A_h = (D_h^\ell + D_h^r) - (\alpha_h^\ell + \alpha_h^r + \bar{s}_h), \quad E_h = \gamma^\ell + \gamma^r + \sigma_h.$$

The flow adjusts to carry storage output into region ℓ , so storage enters the slope in full.

When the line binds, remote storage does not reach the generator at all. The import into ℓ is fixed at $F_h = \bar{F}$, so (B.3) gives $A_h = D_h^\ell - \alpha_h^\ell - \bar{F}$ and $E_h = \gamma^\ell$. When storage discharges, it clears within r , lowering p_h^r and widening λ_h , but it leaves the generator's output and markup unchanged. Storage located in ℓ , by contrast, enters (B.3) directly and gives $E_h = \gamma^\ell + \sigma_h$ regardless of the line.

These three cases order storage by its share of the slope of residual demand. The share σ_h/E_h is $\sigma_h/(\gamma^\ell + \sigma_h)$ for local storage, the smaller $\sigma_h/(\gamma^\ell + \gamma^r + \sigma_h)$ for remote storage on a slack line, and zero for remote storage behind a binding constraint. Local storage therefore delivers the full effect, remote storage an attenuated one, and constrained remote storage none. Storage entry can itself change congestion, because discharging storage in r lowers p_h^r and raises desired exports. I abstract from this by assuming no constraint is marginally binding, following [Woerman \(2023\)](#).

The same logic implies that the effect of storage capacity B should scale with its size relative to the local market, which motivates defining treatment at the county level.

B.2 Strategic Storage

Strategic storage weakens the discipline that price-taking storage imposes, and the price-taking model is its limiting case. I follow [Andres-Cerezo and Fabra \(2025\)](#); [Andrés-Cerezo and Fabra \(2023\)](#), in which storage internalizes the price effect of its own dispatch, and focus on the discharge period. Storage now chooses its net injection s_H as a quantity rather than through the supply schedule in Equation (2). Netting out the fringe, Equation (1) gives the price $p_H = (\tilde{A}_H - q_{iH} - s_H)/\gamma$, where $\tilde{A}_H = D_H - \alpha_H$ is the level of residual demand before storage. Generator i has marginal cost c_i , and storage has marginal cost $m_s = p_L/\eta$.

A conduct parameter $\kappa \geq 0$ indexes how much storage withholds. Storage perceives marginal revenue $p_H - \kappa s_H/\gamma$, so the two first-order conditions are

$$(B.7) \quad \mu_{iH} = p_H - c_i = \frac{q_{iH}}{\gamma} \quad (\text{gas generator}), \quad p_H - m_s = \kappa \frac{s_H}{\gamma} \quad (\text{storage}).$$

Storage is a price taker when $\kappa = 0$ and a Cournot competitor when $\kappa = 1$. Solving,

$$(B.8) \quad s_H^*(\kappa) = \frac{\tilde{A}_H + \gamma c_i - 2\gamma m_s}{1 + 2\kappa}, \quad q_{iH}^*(\kappa) = \frac{1}{2}(\tilde{A}_H - s_H^* - \gamma c_i), \quad \mu_{iH}^*(\kappa) = \frac{q_{iH}^*}{\gamma}.$$

More strategic storage discharges less and leaves a higher gas markup,

$$(B.9) \quad \frac{\partial s_H^*}{\partial \kappa} < 0, \quad \frac{\partial \mu_{iH}^*}{\partial \kappa} = \frac{\tilde{A}_H + \gamma c_i - 2\gamma m_s}{\gamma(1 + 2\kappa)^2} > 0 \quad \text{whenever } s_H^* > 0.$$

Competitive storage delivers the largest markup reduction. As κ grows, storage withholds more, and the gas markup returns toward its level without storage. Storage that is present but strategic can therefore provide little discipline. The main text studies the case $\kappa = 0$, in which the slope σ_h comes from aggregating heterogeneous price-taking batteries.

B.3 Ownership

Portfolio ownership reduces the discipline storage provides through two wedges. Suppose the storage owner also sells generation G_H in the discharge period and G_L in the charge period, and its storage is large enough to move prices along the residual demand of Equation (3), $\partial p_h / \partial s_h = -1/E_h$. Arbitrage links the two periods through $s_L = -s_H/\eta$. Maximizing portfolio profit $\sum_h p_h(s_h + G_h)$ over s_H gives the first-order condition

$$(B.10) \quad \underbrace{\eta p_H - p_L}_{\text{arbitrage spread}} + \underbrace{\left(\frac{s_L}{E_L} - \eta \frac{s_H}{E_H} \right)}_{\text{storage-side market power}} + \underbrace{\left(\frac{G_L}{E_L} - \eta \frac{G_H}{E_H} \right)}_{\text{portfolio}} = 0.$$

A price-taking battery with no other generation keeps only the arbitrage spread and discharges fully, as in Section 4.1. The other two terms reduce discharge.

The storage-side term makes a large storage owner withhold. Because $s_L < 0 < s_H$, the term is negative. The owner recognizes that discharging lowers the discharge-period price and charging raises the charge-period price, so it discharges less and responds less to price. In the notation of the main text, this lowers both the level $\bar{s}_H$ and the slope σ_h of storage supply.

The portfolio term makes an owner with discharge-period generation withhold. Its discharge-period component, $-\eta G_H/E_H$, is negative, because discharging lowers the price at which the owner’s other generation sells. This applies to owners of gas generation and equally to owners of wind or solar, whose output also clears at the discharge-period price. A battery co-located with solar faces a related limit, because its charging is tied to the solar profile, which lowers its responsiveness to price.


Both wedges point in the same direction. Storage owned by firms with large generation portfolios or large storage positions should discharge less and impose less discipline on gas markups than storage owned by independent, storage-only firms.

C Additional Descriptive Statistics

C.1 Trends in Demand and Supply Patterns

Figure C.2 documents the hourly feature of the Texas electricity market over my sample period. First, in Panel B and C, the daily renewable production pattern have changed significantly over the years. Before 2021, the average renewable generation is relatively flat throughout the day, with its highest production at night due wind domination in the renewable generation. As more solar PV capacity is installed, we see a significant increase in renewable generation during the daytime hours, peaking around 1-2 pm. Figure C.2a shows that the hourly demand pattern has also evolved over the years, generally increasing across all hour of day. The hourly pattern shows a typical daily cycle with peaks in the morning and evening hours. These two things combined has implications for the net demand pattern in Figure C.2d, which has evolved from a relatively flat pattern to a more pronounced “duck curve” pattern, with a significant dip during the daytime hours when solar PV production is high and a steep ramp in the morning and evening hours when demand is high and solar PV production is low.


Figure C.2: Hourly demand and supply patterns in ERCOT



Notes: These figures show the average renewable generation, load, net load, and natural gas-based generation for each hour of the day. The net load is defined as the total load minus the total renewable generation.


C.2 Local Market Descriptive Statistics

Figure C.3: Local market statistics



Notes: Panel (a) and (b) summarize the distribution of the number of local markets and the trends in hourly variation of the number of local markets by year.


Figure C.4: Example of local market definition



Notes: Panel (a) and (b) show an example of market definition for intervals 15 minutes apart. Each dot is a generator, with colors representing distinct local markets. In panel (a), there is no congestion and all generators are in the same local market, while in panel (b), congestion splits the generators into 22 local markets. Across the sample, the number of local markets varies from 1 to 137, with a median of 8. Figure C.3 presents the distribution of local markets statistics within my sample.

C.3 Natural Gas Generators Descriptive Statistics

Figure C.5: Natural gas unit markup and capacity factor by hour of day



Notes: Each panel shows the average of each variable for each hour of the day, plotted separately for the pre-2021 period and each year from 2021 through 2024. The data exclude period of extreme event.

D Additional Results

D.1 Selection in Battery Siting

I test whether batteries were sited where prices or gas generator conduct made entry attractive. The generator fixed effects in the main design absorb permanent differences across counties, so selection on levels does not bias the estimates. Selection on trends would. I therefore run three tests. The first relates storage entry to pre-period county characteristics. The second tests whether county price spreads were already changing before entry. The third tests whether, within a county, batteries locate near gas units with high pre-period markups.

County entry I relate storage entry in county m to its pre-period characteristics,

$$(D.11) \quad Y_m = \mathbf{P}_m' \pi + \mathbf{W}_m' \omega + \zeta_{z(m)} + e_m,$$

where Y_m is one of three outcomes: an indicator that the county receives storage during the sample, the log of installed storage in MWh, and the year of first storage entry. The vector $\mathbf{P}_m$ holds four pre-period price measures: the mean nodal price, the mean daily price spread, the mean evening price, and the standard deviation of the daily price. I measure them on 2018–2019 untreated days, in units of \$10/MWh. The vector $\mathbf{W}_m$ holds load growth from 2016 to 2019 in logs and the 2019 solar and wind capacity in GW. The term $\zeta_{z(m)}$ is a load zone fixed effect. I estimate Equation (D.11) by OLS with heteroskedasticity-robust standard errors. For the entry indicator, each coefficient in π is the change in the probability of entry per \$10/MWh. I estimate the first two outcomes on all counties and on counties with gas generation, and the entry year on treated counties only. The test of interest is a Wald test that $\pi = 0$.

Storage siting is not random, but the selection works mainly through solar capacity. Table D.3 shows that solar capacity is the strongest predictor of entry in every sample. Among counties with gas generation, the price measures do not jointly predict whether a county receives storage ($F = 1.28$, $p = 0.30$). They do predict how much storage a county receives ($p = 0.03$) and when it arrives ($p = 0.04$). These are differences in levels, which the fixed effects absorb.

Pre-entry price spreads I test for selection on trends with an event study of the county daily price spread,

$$(D.12) \quad S_{mt} = \sum_{\substack{k=-3 \\ k \neq -1}}^3 \tau_k \mathbb{1}\{K_{m,t} = k\} + \alpha_m + \eta_{\ell(m),t} + \mathbf{X}_{mt} \theta + e_{mt},$$

where S_{mt} is the daily maximum minus the daily minimum of the county mean nodal price on day t , in \$/MWh. The term $K_{m,t}$ counts years relative to first local storage entry, as in Equation (11). The county fixed effect α_m and the load-zone-by-date fixed effect $\eta_{\ell(m),t}$ absorb permanent county differences and regional daily shocks, and $\mathbf{X}_{mt}$ holds the same controls as the main estimates. I

Table D.3: Pre-period county characteristics and storage entry

Sample	Has storage		Log storage (MWh)		Entry year	
	All (1)	Gas (2)	All (3)	Gas (4)	All treated (5)	Gas treated (6)
Mean price (\$10/MWh)	0.661 (0.500)	0.699 (1.388)	1.485 (2.272)	-3.777 (6.980)	-0.838 (3.109)	-2.569 (4.567)
Mean daily spread (\$10/MWh)	0.044** (0.022)	0.250 (0.161)	0.158 (0.100)	0.898 (0.885)	-0.163 (0.123)	-0.805 (0.727)
Mean evening price (\$10/MWh)	-0.355 (0.268)	-0.840 (0.938)	-0.688 (1.203)	0.279 (4.686)	0.395 (1.668)	1.582 (3.521)
SD of daily price (\$10/MWh)	-0.123 (0.089)	-0.181 (0.249)	-0.436 (0.413)	-0.196 (1.670)	0.648 (0.464)	2.003 (1.196)
Load growth, 2016–2019 (log)	0.980 (0.660)	1.947 (1.320)	3.644 (3.508)	9.066 (6.647)	11.638 (7.740)	2.647 (11.552)
Solar capacity, 2019 (GW)	0.472*** (0.164)	4.134** (1.982)	3.253*** (0.744)	9.581 (12.409)	-1.104* (0.618)	-29.754*** (5.059)
Wind capacity, 2019 (GW)	0.008 (0.062)	0.001 (0.226)	0.097 (0.308)	-0.213 (1.184)	0.342 (0.351)	0.443 (0.766)
Load zone FE	✓	✓	✓	✓	✓	✓
Joint test, price variables F (p-value)	2.22 (0.073)	1.28 (0.295)	4.62 (0.002)	2.97 (0.033)	3.90 (0.009)	3.20 (0.036)
Observations	108	46	108	46	49	30
R ²	0.141	0.221	0.191	0.208	0.218	0.412

Notes. One observation per county. Characteristics are measured over 2018–2019 untreated days. Prices are in 2020 dollars and scaled to \$10/MWh. Gas columns restrict to counties with gas generation, the estimation sample. Entry year columns use treated counties only. Solar capacity in gas counties is small, so the solar coefficient in Column (6) reflects a narrow range of variation. All columns include load zone fixed effects. The joint test is a Wald test that the four price coefficients are zero. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

estimate Equation (D.12) with the two-stage estimator, with the first stage on untreated county-days. I construct confidence intervals from 200 draws of a county block bootstrap of both stages. If storage entered where arbitrage value was rising, the pre-entry coefficients τ_{-3} and τ_{-2} would be negative.


County price spreads show no trend before entry. Figure D.6 shows that the pre-entry coefficients are 1.71 (s.e. 2.29) and 1.14 (s.e. 4.07), and a joint test does not reject that both are zero ($p = 0.43$). After entry, the spread falls by \$9.44 per MWh in the year of entry (s.e. 5.83), consistent with arbitrage narrowing the gap between peak and off-peak prices, although the estimate is not statistically significant.

Within-county proximity I test whether batteries locate near gas units with high pre-period markups. For gas unit j in county $m(j)$, I estimate

$$(D.13) \quad D_j = \mathbf{C}_j' \boldsymbol{\kappa} + \mathbf{G}_j' \boldsymbol{\varphi} + \alpha_{m(j)} + e_j,$$

where D_j is either an indicator for a battery within 10 km of the unit or the log distance to the nearest battery, among batteries operating by December 2024. The vector $\mathbf{C}_j$ holds the unit's pre-

Figure D.6: County daily price spread before and after storage entry



Notes. Two-stage difference-in-differences estimates of the county daily price spread, the daily maximum minus minimum of the county mean nodal price, by year relative to first local storage entry. The first stage uses untreated county-days, with county and load-zone-by-date fixed effects and the same controls as the main estimates. The 95% confidence intervals use standard errors from a county block bootstrap of both stages. The year before entry is the reference.

period conduct: its evening markup, the nodal price minus marginal cost from 5 PM to 9 PM when the unit runs, and its mean evening price, both in \$10/MWh and measured on untreated hours. The vector $\mathbf{G}_j$ holds the unit's capacity factor, pre-period heat rate, and capacity in 100 MW. The county fixed effect $\alpha_{m(j)}$ restricts the comparison to units in the same county. I cluster standard errors by county and test $\kappa = 0$ with a Wald test.

Batteries do not locate closer to gas units with higher pre-period markups. Table D.4 shows that the markup and price coefficients are small and jointly insignificant in both columns ($p = 0.33$ and $p = 0.36$). The markup coefficient in Column (1) is 0.002, with a 95% confidence interval from -0.003 to 0.008 . A \$10/MWh higher evening markup therefore raises the probability of a battery within 10 km by at most 0.8 percentage points, against a mean of 28 percent. Siting does not select on the conduct I measure.

Table D.4: Battery proximity and gas unit pre-period characteristics

	Battery within 10 km (1)	Log distance to battery (2)
Evening markup (\$10/MWh)	0.002 (0.003)	-0.010 (0.009)
Mean evening price (\$10/MWh)	0.007 (0.009)	0.005 (0.015)
Capacity factor	0.036 (0.139)	-0.326 (0.247)
Pre-period heat rate	0.005 (0.004)	-0.018 (0.019)
Capacity (100 MW)	-0.003 (0.002)	0.012* (0.007)
County FE	✓	✓
Mean of outcome	0.284	2.928
Joint test, markup and price F (p-value)	1.13 (0.333)	1.06 (0.356)
Observations	473	473
Adjusted R ²	0.736	0.841


Notes. One observation per gas unit. Characteristics are measured over untreated hours. The evening markup is the nodal price minus marginal cost from 5 PM to 9 PM when the unit runs. Batteries are those operating by December 2024. All columns include county fixed effects, so the comparison is within county. The joint test is a Wald test that the markup and price coefficients are zero. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.2 Additional Market Power Results

This appendix reports the market power and residual demand slope results by technology and the event studies not shown in the main text.


Markups by technology and hour group Figure D.7 presents the hourly markup effects by technology. All three technologies show the same pattern as the pooled estimate, with no change overnight and a fall in the late afternoon and evening. The evening effects average -0.92 for combined cycle, -1.57 for simple cycle, and -1.70 for steam turbine generators. Figure D.8 presents the markup event studies for the overnight, midday, and evening hour groups. The coefficients before entry are small and insignificant in every hour group. After entry, markups fall in all three hour groups. The overnight and midday declines in the first year after entry are individually significant, while the evening coefficients are larger but imprecise. The overnight decline does not appear in the dose estimates by hour, and it reverses by the third year after entry.

Figure D.7: Effect of local storage on markups by technology



Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the markup, in \$/MWh, by hour and technology. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Figure D.8: Event study: markups by hour group



Notes. Binary event studies of the markup for midnight to 6 AM, 10 AM to 4 PM, and 5 PM to 9 PM, with years relative to the first local storage entry binned at three years and the year before entry as the reference. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Residual demand slope Figure D.9 presents the effect of storage on the measured residual demand slope by technology. The residual demand of combined cycle generators becomes significantly flatter between 6 PM and 11 PM and steeper at 3 PM. The effects for simple cycle and steam turbine generators are imprecise. Figure D.10 presents the event study for the slope by hour group. The coefficients two and three years before entry are small and insignificant in every hour group, so treated and untreated counties followed parallel paths before entry. After entry, the midday slope rises by 12 and 28 MW per \$/MWh in the second and third year. The post-entry effects in the other hour groups are imprecise.

Operating profit and rents by technology Figure D.11 presents the hourly operating profit effects by technology. Combined cycle and simple cycle generators lose operating profit in the evening, by \$0.83 and \$1.07 per MW per hour on average from 5 PM to 9 PM. Simple cycle generators also gain a small amount overnight. The steam turbine estimates are imprecise. Table D.5 presents the daily effects by technology. Under the central benchmark, the market power rent falls by \$6.24 per MW-day for combined cycle and \$6.69 per MW-day for simple cycle generators, and the competitive rent does not change significantly. For steam turbines, the operating profit effect is not significant.


Market power rent event study Figure D.12 presents the event study for the daily market power rent. The coefficients before entry are small and insignificant. After entry, the market power rent falls by \$7 to \$12 per MW-day in the year of entry and the two years after it. The coefficient three or more years after entry is positive, and all post-entry coefficients are imprecise.

Table D.5: Effect of local storage on daily operating profit and rents by technology

	Operating profit (1)	Competitive rent (2)	Market power rent (3)	Competitive rent (price-setter) (4)	Market power rent (price-setter) (5)
<i>Combined Cycle</i>					
\$/MW-day per 100 MWh	-6.38** (2.46)	-0.14 (0.51)	-6.24*** (2.24)	-1.76 (2.49)	-4.62** (2.19)
Share of profit effect	1.00	0.02	0.98	0.28	0.72
Annual total (\$ million)	-36.5 (14.1)	-0.8 (2.9)	-35.7 (12.8)	-10.1 (14.2)	-26.4 (12.5)
Per 100 MWh of storage (\$ thousand/year)	-3,122 (1,207)	-68 (250)	-3,054 (1,099)	-861 (1,219)	-2,261 (1,071)
<i>Simple Cycle</i>					
\$/MW-day per 100 MWh	-6.53** (2.72)	0.15 (1.52)	-6.69*** (2.16)	-0.82 (1.93)	-5.71** (2.71)
Share of profit effect	1.00	-0.02	1.02	0.13	0.87
Annual total (\$ million)	-13.9 (5.8)	0.3 (3.2)	-14.2 (4.6)	-1.7 (4.1)	-12.1 (5.7)
Per 100 MWh of storage (\$ thousand/year)	-1,185 (494)	28 (275)	-1,213 (392)	-149 (350)	-1,036 (492)
<i>Steam Turbine</i>					
\$/MW-day per 100 MWh	-3.31 (5.75)	-2.99*** (0.92)	-0.32 (5.54)	-1.63 (1.00)	-1.68 (4.95)
Share of profit effect	1.00	0.90	0.10	0.49	0.51
Annual total (\$ million)	-2.8 (4.9)	-2.6 (0.8)	-0.3 (4.7)	-1.4 (0.9)	-1.4 (4.2)
Per 100 MWh of storage (\$ thousand/year)	-242 (421)	-219 (67)	-24 (405)	-120 (73)	-123 (362)
Observations (generator-days)	203,395				
Counties	32				

Notes. Two-stage difference-in-differences estimates on the generator-day panel by technology, weighted by capacity. Columns are defined as in Table 3. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure D.9: Residual demand slope by technology




Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the local residual demand slope faced by each technology, in MW per \$/MWh, by hour. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Alternative competitive benchmark Table D.6 reports the decomposition under a third benchmark, which sets the competitive price to the highest marginal cost among gas generators actually running. This benchmark attributes only 18 percent of the profit loss to market power, and the market power rent is not significant. It is, however, mechanically biased by the treatment. Storage discharge displaces exactly the most expensive running generators, so the benchmark falls when storage enters, and the decomposition records that fall as a loss of competitive rent. I therefore use the price-setter benchmark as the lower bound in the main text. Figure D.13 presents the hourly rents under the price-setter benchmark, which follow the same pattern as under the central benchmark.

Gas price-setting Figure D.14 presents the effect of local storage on the probability that a gas generator sets the nodal price. Before storage entry, a gas generator sets the price in 53 to 69 percent of intervals, depending on the hour. The effect of storage is small and not significant in

Figure D.10: Event study: residual demand slope



Notes. Binary event study from Equation (11) by hour group, with years relative to the first local storage entry binned at three years. The year before entry is the reference. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.


Table D.6: Rent decomposition under the highest-running-cost benchmark

	Operating profit (1)	Competitive rent (highest running cost) (2)	Market power rent (highest running cost) (3)
\$/MW-day per 100 MWh	-5.45*** (1.85)	-4.47** (2.11)	-0.98 (2.53)
Share of profit effect	1.00	0.82	0.18
Annual total (\$ million)	-47.4 (16.1)	-38.9 (18.3)	-8.6 (22.0)
Per 100 MWh of storage (\$ thousand/year)	-4,057 (1,378)	-3,324 (1,567)	-733 (1,885)
Observations (generator-days)	671,972		
Counties	49		

Notes. Two-stage difference-in-differences estimates on the generator-day panel, weighted by capacity. The competitive price is the highest marginal cost among gas generators running in the local market and interval. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


any hour, ranging from -0.04 to 0.26 percentage points per 100 MWh.

Figure D.11: Effect of local storage on hourly operating profit by technology




Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on hourly operating profit per MW of nameplate capacity, by technology, weighted by capacity. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Figure D.12: Event study: daily market power rent




Notes. Binary event study of the daily market power rent per MW of capacity under the central benchmark, weighted by capacity, with the year before entry as the reference. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Figure D.13: Hourly rents under the price-setter benchmark



Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the competitive and market power rents under the price-setter benchmark, per MW of capacity per hour, weighted by capacity. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Figure D.14: Effect of local storage on gas price-setting



Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the share of intervals in the hour in which a gas generator sets the nodal price, in percentage points. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

D.3 Additional Gas Generator Output Results

This appendix reports the full structural estimates, the robustness of the two-stage estimates, and additional results on heterogeneity.

Structural estimates Table D.7 reports the structural estimates of Equation (16) across samples, weights, and fixed effects. The IV estimate is -0.79 on interior hours (Column 2) and -0.81 when I include hours at the operating limits (Column 3). Excluding hours within 2 percent of capacity of either limit shrinks it to -0.26 and makes it insignificant (Column 4). Capacity weights give -0.66 (Column 5), and load-zone-by-date-by-hour fixed effects, which absorb the common natural gas price shock, give -0.81 (Column 6). The cost coefficient and the interaction term are negative and significant in every column except the 2 percent band.

Table D.7: Structural estimates across samples, weights, and fixed effects

	OLS	Capacity factor (pp)				
	(1)	(2)	(3)	IV (4)	(5)	(6)
Storage B (100MWh), $\hat{\delta}^{(1)}$	-0.785*** (0.291)	-0.790*** (0.292)	-0.811*** (0.261)	-0.259 (0.175)	-0.656** (0.263)	-0.814** (0.343)
Cost $\tilde{c}$ (centered, \$/MWh), $\hat{\delta}^{(2)}$	-0.070*** (0.019)	-0.070*** (0.017)	-0.070*** (0.019)	-0.081*** (0.013)	-0.100*** (0.017)	-0.070*** (0.020)
Cost $\tilde{c} \times B$, $\hat{\delta}^{(3)}$	-0.016*** (0.004)	-0.016*** (0.003)	-0.019*** (0.003)	0.000 (0.008)	-0.011*** (0.004)	-0.017*** (0.004)
Generator FE	✓	✓	✓	✓	✓	✓
Load zone-Month-of-sample-Hour FE	✓	✓	✓	✓	✓	
Load zone-Date-Hour FE						✓
Local renewable capacity and potential	✓	✓	✓	✓	✓	✓
Local temperature	✓	✓	✓	✓	✓	✓
Weights	None	None	None	None	Capacity	None
Sample	Interior	Interior	Running	Interior 2% band	Interior	Interior
Observations	3,770,307	3,770,307	5,439,859	1,744,113	3,770,307	3,761,401
First-stage F , $\tilde{c}$		327.3	626.1	207.8	407.0	298.0
First-stage F , $\tilde{c} \times B$		2,447.3	2,351.4	1,762.3	534.2	9,073.3
Partial R^2 , $\tilde{c}$		0.397	0.376	0.374	0.410	0.565
Partial R^2 , $\tilde{c} \times B$		0.810	0.802	0.846	0.853	0.835

Notes. Estimates of Equation (16). The outcome is the hourly capacity factor in percentage points. Marginal cost is centered within technology and instrumented by the pre-entry heat rate times the Henry Hub price. Column (3) uses all hours with positive output, and Column (4) drops interior hours within 2 percent of capacity of either operating limit. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table D.8 reports the structural estimates by technology. None of the technology-specific storage coefficients is statistically significant. The interaction term is negative and significant for combined cycle and simple cycle generators, so storage displaces the higher-cost units within these technologies more. It is positive for steam turbines. Figure D.15 presents the hourly structural estimates. The displacement at mean cost is between 0.2 and 0.4 percentage points for combined cycle generators and between 0.4 and 1.2 percentage points for simple cycle generators, and it is imprecise in most

hours. The cost gradient has the predicted sign in every hour for both technologies. Figure D.16 shows that the instruments explain 32 to 46 percent of the remaining variation in cost and 79 to 87 percent of the remaining variation in the interaction for these two technologies in every hour.

Table D.8: Structural estimates by technology


	Capacity factor (pp)			
	All gas (1)	Combined Cycle (2)	Simple Cycle (3)	Steam Turbine (4)
Storage B (100MWh), $\hat{\delta}^{(1)}$	-0.790*** (0.292)	-0.298 (0.209)	-1.022 (0.641)	0.169 (0.547)
Cost $\tilde{c}$ (centered, \$/MWh), $\hat{\delta}^{(2)}$	-0.070*** (0.017)	-0.092*** (0.016)	-0.038 (0.026)	-0.120*** (0.037)
Cost $\tilde{c} \times B$, $\hat{\delta}^{(3)}$	-0.016*** (0.003)	-0.009** (0.004)	-0.024*** (0.005)	0.029*** (0.007)
Generator FE	✓	✓	✓	✓
Load zone-Month-of-sample-Hour FE	✓	✓	✓	✓
Local renewable capacity and potential	✓	✓	✓	✓
Local temperature	✓	✓	✓	✓
Sample	Interior	Interior	Interior	Interior
Observations	3,770,307	1,821,957	1,608,347	339,526
First-stage F , $\tilde{c}$	327.3	250.7	114.3	89.7
First-stage F , $\tilde{c} \times B$	2,447.3	1,879.4	12,777.0	348.2
Partial R^2 , $\tilde{c}$	0.397	0.413	0.367	0.187
Partial R^2 , $\tilde{c} \times B$	0.810	0.845	0.818	0.647

Notes. Estimates of Equation (16) on interior hours, separately by technology. Cost and the instrument are centered within technology. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Two-stage estimates across samples Table D.9 reports the pooled two-stage estimates across samples and weights. The storage coefficient is -0.55 on interior hours, -0.34 on hours with positive output, and -0.44 when I exclude hours within 2 percent of either operating limit. It is -0.79 on all hours and -0.90 on all hours with capacity weights. Unlike the IV estimate, the two-stage estimate remains significant when I exclude hours near the operating limits.

Operating margins Table D.10 presents the effect of storage on the operating state of gas generators. For combined cycle generators, the evening displacement works through fewer hours at maximum output, which fall by 1.4 percentage points. For simple cycle generators, storage lowers the probability of running at minimum load during midday by 1.6 percentage points. Steam turbines behave differently. In the evening peak, storage increases the probability that a steam turbine runs by 1.7 percentage points and runs in the interior by 2.7 percentage points, while lowering the probability of running at maximum output by 1.6 percentage points. Steam turbines therefore run more but further from their maximum output. This shift changes which steam turbines appear in the interior sample, which is why I use all hours for the welfare specification.


Figure D.15: Hourly structural displacement by technology



Notes. Estimates of Equation (16) by technology and hour on interior hours, in percentage points of capacity factor per 100 MWh of local storage. Panel (b) evaluates the displacement at the 25th and 75th percentiles of marginal cost within each technology. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Reconciling the structural and two-stage estimates Table D.11 changes one feature at a time by technology and hour group. For combined cycle generators, the four estimates are close in every hour group, so the choice of specification matters little. For simple cycle generators, the two-stage estimate on the interior sample is larger and more precise than the IV estimate, at 0.8 to 1.6 percentage points, and changing the sample to all hours moves the estimate toward the welfare specification. For steam turbines, the estimates on the interior sample are unreliable because storage changes which steam turbines operate in the interior. Table D.10 shows that storage moves steam turbines away from maximum output and into the interior, which changes the composition of the interior sample. The estimates on all hours do not have this problem.

Figure D.16: First-stage strength by hour



Notes. Partial R^2 of the excluded instruments for the cost term and for the cost-by-storage interaction in each hourly first stage, by technology.

Bounded counterfactual Table D.12 compares displaced generation with and without the bounds on counterfactual output. Bounding raises the total from 811 GWh to 849 GWh per year, mostly among steam turbines, whose counterfactual output hits zero in 30 percent of generator-hours with local storage.

Table D.9: Two-stage estimates across samples and weights

	Capacity factor (pp)				
	2SDD (1)	2SDD (2)	2SDD (3)	2SDD (4)	2SDD (5)
Storage B (100MWh)	-0.553*** (0.165)	-0.342** (0.129)	-0.437** (0.212)	-0.794*** (0.224)	-0.904*** (0.266)
Predetermined cost $\tilde{z} \times B$	-0.013*** (0.003)	-0.009** (0.004)	-0.001 (0.010)	-0.002 (0.004)	-0.004 (0.008)
Generator FE	✓	✓	✓	✓	✓
Load zone-Date-Hour FE	✓	✓	✓	✓	✓
Local renewable capacity and potential	✓	✓	✓	✓	✓
Local temperature	✓	✓	✓	✓	✓
Sample	Interior	Running	Interior 2% band	All hours	All hours
Weights	None	None	None	None	Capacity
Observations	3,705,777	5,362,817	1,696,824	13,472,250	13,472,250
Counties	49	49	49	49	49

Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the capacity factor, in percentage points, with storage interacted with predetermined cost. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table D.10: Effect of local storage on operating margin

	Running	At minimum load	Interior	At maximum output
<i>Combined Cycle</i>				
Overnight (0–6)	-0.390 (0.296)	-0.424 (0.481)	-0.520 (0.653)	0.554 (0.492)
Midday (10–16)	-0.162 (0.274)	-0.032 (0.323)	0.301 (0.520)	-0.431 (0.457)
Evening (17–21)	-0.408 (0.278)	0.276 (0.311)	0.697 (0.632)	-1.381** (0.640)
<i>Simple Cycle</i>				
Overnight (0–6)	2.585 (3.174)	-0.901 (0.570)	3.510 (3.668)	-0.024 (0.252)
Midday (10–16)	1.899 (2.087)	-1.593** (0.757)	3.250 (2.622)	0.242 (0.424)
Evening (17–21)	1.549* (0.857)	-1.516 (1.037)	2.433 (2.283)	0.632 (1.165)
<i>Steam Turbine</i>				
Overnight (0–6)	1.109* (0.569)	-0.233 (1.065)	2.125*** (0.720)	-0.784* (0.374)
Midday (10–16)	1.651* (0.817)	-0.047 (0.655)	1.706** (0.806)	-0.008 (0.343)
Evening (17–21)	1.712*** (0.490)	0.637 (0.625)	2.672*** (0.744)	-1.598*** (0.338)

Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on indicators for running, running at minimum load, running in the interior, and running at maximum output, in percentage points, on hours with both operating limits observed. The last three columns sum to the first. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table D.11: Reconciling the structural and two-stage estimates

	(1) IV Interior Unweighted	(2) 2SDD Interior Unweighted	(3) 2SDD All hours Unweighted	(4) 2SDD All hours Capacity-weighted
<i>Combined Cycle</i>				
Overnight (0–6)	-0.344 (0.263)	-0.289 (0.217)	-0.216 (0.186)	-0.203 (0.243)
Midday (10–16)	-0.216 (0.180)	-0.170 (0.172)	-0.203 (0.177)	-0.225 (0.154)
Evening (17–21)	-0.356 (0.230)	-0.341* (0.201)	-0.451** (0.187)	-0.438** (0.184)
<i>Simple Cycle</i>				
Overnight (0–6)	-1.169 (0.753)	-1.485*** (0.206)	-0.789 (0.677)	-1.202 (0.931)
Midday (10–16)	-0.988 (0.667)	-1.596*** (0.213)	-0.674** (0.317)	-0.774 (0.542)
Evening (17–21)	-0.641 (0.607)	-0.812*** (0.250)	-0.813 (0.582)	-0.802** (0.357)
<i>Steam Turbine</i>				
Overnight (0–6)	-0.847 (0.626)	-0.126 (0.157)	-0.435*** (0.067)	-0.320 (0.372)
Midday (10–16)	0.736 (0.652)	4.382*** (1.150)	0.291 (0.225)	0.434 (0.706)
Evening (17–21)	-0.282 (0.756)	3.160* (1.665)	-0.626*** (0.202)	0.170 (0.606)

Notes. Effect of 100 MWh of local storage on the capacity factor, in percentage points, by technology and hour group. Each column changes one feature relative to the previous one: the estimator, then the sample, then the weights. Column (4) is the specification used in the welfare analysis. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Table D.12: Displaced gas generation with bounded and unbounded counterfactual output

	Displaced generation (GWh per year)		Generator-hours bounded (%)
	Bounded	Unbounded	
Combined Cycle	371.0 (191.4)	370.4 (203.3)	1.3
Simple Cycle	440.2 (217.4)	429.8 (258.1)	7.1
Steam Turbine	37.8 (66.4)	10.9 (93.5)	29.7
All gas	849.0 (304.8)	811.1 (341.6)	10.1

Notes. Annual gas generation displaced by local storage, from the technology-specific two-stage estimates with the cost interaction. The bounded counterfactual keeps output without storage between zero and capacity. Standard errors for bounded totals come from 500 draws of the hourly coefficients; standard errors for unbounded totals use the delta method. The last column is the share of generator-hours with local storage in which a bound binds.

Event study estimates by technology Figure D.17 presents event studies of the capacity factor from Equation (11), by technology and hour group, with the predetermined cost shifter in the first stage. The pre-entry coefficients are within about one percentage point of zero in every panel, which supports parallel trends in output. A few coefficients two years before entry for simple cycle and steam turbine generators are marginally different from zero, but they are small relative to the post-entry effects. After entry, simple cycle output falls in every hour group, by about 5.5 percentage points at midday and 4 percentage points in the evening two years after entry. Combined cycle output falls by about 2 percentage points at midday and in the evening one and two years after entry, although most of these coefficients are imprecise. Steam turbine output rises after entry at midday and in the evening, consistent with the operating margin results in Table D.10. These binary event studies measure the effect of entry rather than the dose, so their magnitudes are not comparable to the per 100 MWh estimates.


Figure D.17: Event studies of the capacity factor, by technology and hour group



Notes. Two-stage difference-in-differences estimates of the effect of local storage entry on the hourly capacity factor, in percentage points, by years relative to first local entry. The year before entry is the reference. Rows are technologies; columns are hour groups. Estimates use all hours, weighted by capacity, with the predetermined cost shifter in the first stage. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Local renewable controls Figure D.18 compares the welfare specification with and without the local renewable controls. The controls lower the combined cycle estimates by about 0.1 percentage points in every hour. For simple cycle generators, they remove most of the overnight displacement but leave the afternoon displacement nearly unchanged. Solar capacity that enters alongside storage therefore does not explain the afternoon displacement.

Figure D.18: Displacement with and without local renewable controls




Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the capacity factor, by hour and technology, on all hours and weighted by capacity. The controls are county solar and wind capacity and their potential output, capacity times irradiance or a normalized wind power curve. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Heterogeneity Figure D.19 presents the displacement for all gas generators by region and by renewable conditions. In the Houston and North load zones, storage lowers the capacity factor by 1.5 percentage points in the evening and 1.2 percentage points at midday, and the evening effect is significant in every hour. In the West and South load zones, the evening effect is 1.1 percentage points and not significant. The displacement is larger on days with above-median solar output, by about 0.4 percentage points at midday and 0.3 percentage points in the evening. It is also larger on

days with above-median wind output, especially overnight, when it is 0.8 percentage points against 0.3 percentage points on low-wind days.

Figure D.19: Displacement by region and renewable conditions




Notes. Two-stage difference-in-differences estimates of the effect of 100 MWh of local storage on the capacity factor of all gas generators, by hour, on all hours and weighted by capacity. Panel (a) splits the sample into the Houston and North load zones and the West and South load zones. Panels (b) and (c) split days at the sample median of statewide daily maximum solar irradiance and wind speed. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

D.4 Additional Welfare Results

Effect on nodal prices Figure D.20 presents the effect of local storage on the county nodal price, which enters the consumer surplus calculation in Equation (22). I average the nodal prices at all non-storage nodes in each county and hour. The sample includes counties with gas generation and, until they receive storage, counties without gas generation. Panel (a) shows that the price falls in every hour, by \$0.08 to \$1.15 per MWh per 100 MWh of storage. The largest decline is at 7 PM, when storage discharges. The declines overnight and late in the evening are individually significant, while those from 8 AM to 8 PM are not. Panel (b) presents the event study for the evening peak. The coefficients before entry are small and insignificant. After entry, the evening price falls by \$1.6 to \$3.9 per MWh, although the individual coefficients are imprecise.

Figure D.20: Effect of local storage on county nodal prices




Notes. Two-stage difference-in-differences estimates on the county-by-hour panel of nodal prices, in 2020 dollars per MWh, averaging all non-storage settlement points in each county, with county and load-zone-by-date-by-hour fixed effects and the local renewable and temperature controls. The sample includes counties with gas generation and, until they receive storage, counties without gas generation. Counties are weighted equally. Panel (a) reports the effect of 100 MWh of local storage by hour. Panel (b) reports a binary event study for 5 PM to 9 PM, with the year before entry as the reference. The 95% confidence intervals use t_{G-1} critical values and standard errors clustered at the county level.

Consumer surplus across specifications The consumer surplus estimate in Table 6 weights counties equally. Weighting by county load instead gives a consumer loss of \$15.3 million per year (s.e. 34.4). The load-weighted estimate shows a significant price decline two years before entry, because prices in the largest counties were already rising relative to their load zones before storage entered. Controlling for county load does not remove this trend, while the unweighted estimate shows no trend before entry. Other choices change the estimate little. Adjusting the county price for changes in which nodes report gives \$68.8 million (s.e. 55.9), and weighting the hours of the day equally gives \$61.4 million (s.e. 49.2). Following the design of Kirkpatrick (2026), with fixed effects for county, hour, quarter, year, and weekend and county-by-hour slopes on the load zone price, gives -\$26.1 million (s.e. 103.6), which is consistent with the main estimate given its precision. None of

these estimates is statistically significant, which is why I interpret only the combined transfer to consumers and other generators.

Distribution of avoided damages across Texas counties Figure D.21 maps the avoided local damages across Texas counties. Panel (a) shows the total avoided damages in each receptor county, which are largest in the most populated counties of Harris, Dallas, Tarrant, and Bexar. Panel (b) divides the damages by county population. Per resident, the avoided damages are more evenly spread and remain highest in and around the counties that host gas generation and storage.

Figure D.21: Avoided local damages by Texas county



Notes. Annual avoided damages from NO_x , SO_2 , and $\text{PM}_{2.5}$ in each receptor county, in 2020 dollars, from the avoided emissions of gas generators displaced by local storage and the AP4 source-receptor matrices. Panel (b) divides by county population. The estimates use the local displacement and are an upper bound if displaced gas generation moves to counties without storage.

E Additional Tables


Table E.13: Generation Mix Summary

	Capacity			Mean Generation		Units	
	GW	Duration (hr)	%	GWh	%	Count	%
<i>Panel A: Fuel Type</i>							
Natural Gas	65.23		38.9	20.482	44.5	159	24.2
Coal	20.19		12.1	8.252	17.9	17	2.6
Wind	41.57		24.8	10.453	22.7	192	29.2
Solar PV	26.63		15.9	2.062	4.5	133	20.2
Batteries	8.04	1.33	4.8	0.043	0.1	142	21.6
Nuclear	5.23		3.1	4.666	10.1	2	0.3
Other	0.64		0.4	0.079	0.2	12	1.8
Total	167.54		100.0	46.036	100.0	657	100.0
<i>Panel B: Top 5 Natural Gas Firms</i>							
Luminant	12.91		19.8	4.767	21.7	18	11.3
Calpine Corporation	10.04		15.4	6.877	31.3	17	10.7
Constellation Energy	6.09		9.3	2.007	9.1	10	6.3
NRG Energy	5.64		8.6	0.943	4.3	12	7.6
City of San Antonio	5.53		8.5	1.449	6.6	10	6.3
Others	25.03		38.4	5.943	27.0	92	57.9
Total	65.23		100.0	21.985	100.0	159	100.0
<i>Panel C: Top 5 Storage Firms</i>							
Engie S.A.	1.83	1.02	22.8	0.017	27.8	24	16.9
Enel S.p.A.	1.19	1.69	14.8	0.014	22.6	15	10.6
Plus Power	0.79	1.92	9.8	0.009	14.8	4	2.8
Key Capture Energy	0.63	1.83	7.8	0.002	3.5	11	7.8
Luminant	0.44	2.05	5.5	0.003	5.4	4	2.8
Others	3.16	1.22	39.3	0.016	25.8	84	59.1
Total	8.04		100.0	0.061	100.0	142	100.0

Notes. Panel A shows capacity, generation, and number of units by fuel type. Panel B shows capacity, generation, and number of units by firms for natural gas unit, while Panel C shows specifically for storage firms. Discrepancies in generation between each panel arise from differences in operating intervals when aggregating generation over multiple units. I define capacity using maximum observed generation from my data, this creates discrepancies in generation capacity compared to (ERCOT, 2025). The largest firm in ERCOT owns 15 percent of total capacity and contributes, on average, about 20 percent of electricity generation in each interval. The second to fifth largest firms own between 4 to 6 percent of total capacity, contributing on average between 5 to 14.6 percent of electricity generation. For storage firms, the largest owns 21.9 percent of storage capacity, while the second to fifth largest own between 3.6 to 10 percent of storage capacity. Based on just the top five storage firms, there is significant heterogeneity in storage characteristics in terms of unit capacity and duration.

F Additional Figures


Figure F.22: Location of Generator in ERCOT



Notes: This figure shows the spatial distribution of each generator unit based on their fuel type in ERCOT. Natural gas and coal units are concentrated in the populated area of Texas, where the demand center is, such as the cities of Austin, Dallas, Houston, and San Antonio. Solar PV and storage are relatively spread evenly throughout the state, while wind is concentrated in the West and South load zone, the area with the highest wind potential.

Figure F.23: Example Storage Offer

(a) Example of unconstrained storage offer



(b) Example of energy constrained storage offer




Figure F.24: Installed capacity by Fuel Type from 2018-2024






Figure F.25: Treatment assignment and market interaction frequency




Notes: This figure shows which natural gas generators are treated by the LILY_BESS1 storage unit. The blue dot shows the location of the storage unit. The solid black triangle represents generators located in the same county as the storage unit. The red triangle represents generators sharing the same local market with the storage unit more than 97% of the intervals prior to storage COD. The void triangle represents generators that do not share the same local market with the storage unit more than 97% of the intervals prior to storage COD.

Figure F.26: Storage Capacity Variation across County



Notes: This figure shows the evolution of cumulative capacity of commercially operating battery storage for each county in my sample.

Figure F.27: Distribution of storage first observation to COD



Notes: This figure shows the distribution of the time gap between the first date a storage unit shows up in the SCED data (synchronized) and the COD date.