Data Centers: Location Choice and Grid Externalities*

Coly Elhai, *Harvard University*

Yixin Zhou, *Harvard University*

November 3, 2026

[Click here for latest version]

Abstract

Data centers already account for 5% of U.S. electricity consumption, a share that is projected to rise dramatically as new facilities come online. The price and emissions effects of this load growth depend on where it occurs: data centers drive up prices more in locations with steep electricity supply curves and increase emissions more in places with carbon-intensive marginal generation. We combine a dispatch model with a model of firm location choice to quantify data centers' effects on wholesale electricity prices and emissions, how these effects are mediated by firm location choice, and the role of public policies in shaping both location choice and social welfare. We find that data centers disproportionately locate in places with relatively flat supply curves, a pattern that has attenuated their effects on consumer electricity prices. Although single-state data center moratoria can reduce local electricity prices, they can magnify the national price and emissions impacts of data center growth by pushing facilities into locations with steeper supply curves or more carbon-intensive marginal generation.

Keywords: Data centers; electricity; location choice; economic geography

JEL Codes: L94, Q41, R30

*We are very thankful to Jesse Shapiro, Jim Stock, Anna Russo, Ed Glaeser, Gabriel Kreindler, Rebecca Diamond, and Owen Zidar for their valuable feedback. The data centers database used in this paper was built in companion work with Garrison Schlauch, Simon Greenhill, Kevin Butler, Gianluca Guidi, and Falco Bargagli Stoffi. Elhai gratefully acknowledges support from the Lab for Economic Applications and Policy and the National Science Foundation Graduate Research Fellowship Program.

1 Introduction

Data center electricity demand is rising rapidly. It is also highly geographically concentrated: almost half of all U.S. data center capacity is located in just three states. Data centers account for 5% of total U.S. electricity consumption, but in some states, that share is already over 15% and could rise to nearly 50% in the next few years (EPRI 2024). How does the geography of data center development shape its effects on electricity prices and emissions?

We estimate how data center electricity demand affects prices and emissions by combining a new dataset on the data center buildout with a model that approximates the grid as a set of regional electricity supply curves. We find that where data centers are located is an important determinant of their impacts. A data center drawing one gigawatt of power would raise electricity costs by over $3 per megawatt-hour in New England, but just $0.30 in the upper Midwest. That same data center would add about 3.7 megatons of CO_2 per year if built in New England, versus almost twice that in the upper Midwest. Data center location preferences have attenuated the price effects of the data center buildout so far because developers tend to choose regions with flatter electricity supply curves, which also generally have lower electricity price levels and more existing electricity infrastructure. Because development geography matters, policies that affect where data centers are located change the total costs of data center development. We build a model of data center location choice with endogenous electricity prices to show that, on average, proposed state-level data center moratoria would increase the total electricity cost and emissions impacts of data centers. However, some moratoria could substantially reduce data centers' emissions effect by redirecting facilities toward locations with cleaner electric grids.

We use a dispatch model of the electricity grid, which builds region-level electricity supply curves out of power plant-level data, allowing us to identify which power plants respond as electricity demand varies. Using this structure, we can estimate the total wholesale price and emissions effects of the data center buildout over time.[1] We calculate that

[1] We focus on wholesale prices, which capture the cost of generating electricity. Retail rates also include the costs of building electricity infrastructure, which data centers may raise or lower depending on how much infrastructure must be upgraded to serve them and how utilities allocate costs across customer classes. Policies have been proposed or implemented in several states and at the federal level that would require data centers to pay the full cost of those upgrades. The generation-cost margin we model does not depend on cost allocation policies.

in 2025, wholesale electricity prices were about $1.60 per megawatt-hour higher than they would have been without any data centers. This translates to residential consumer costs from data centers of about $2.3 billion per year. Actual price effects are close to what could have been achieved if data centers were reallocated under a price-minimizing objective. We find that the climate costs of the data center buildout are much larger. Assuming fixed generator entry, data center electricity demand was associated with annual emissions of about 200 megatons in 2025, about an eighth of the total from the U.S. power sector, and equivalent to $38 billion in climate damages. Even the climate damage-minimizing data center placement would have guaranteed climate damages over $30 billion per year.

We introduce a structural model of data center location choice in order to estimate how counterfactual policies would reallocate data center development across space, and what this would mean for data centers' price and emissions effects. In our model, data center developers choose where to build as a function of local characteristics such as electricity prices, infrastructure, and land values. Their entry decisions affect regional electricity demand, which in turn shifts electricity prices and aggregate emissions. We nest our dispatch model within a location choice model, endogenizing electricity prices, then estimate it using maximum likelihood with an inner fixed point. Our results indicate data center electricity price elasticities between -0.2 and -2.0, values that overlap with estimates for other electricity-intensive industries (Kahn and Mansur 2013). We also find strong preferences for utilities with large non-residential electricity sales, utilities where firms have already built campuses, and locations with better connected transmission networks.

Using our estimated parameters, we model the impacts of policies that would change the data center location decision. We model both backward-looking counterfactuals, in which we reallocate campuses built 2021–2025, and forward-looking counterfactuals, which reallocate campuses planned for 2026–2030. We find that the existing grid lacks the generation capacity to support all of the capacity associated with data center campuses slated to open between 2026 and 2030. Our forward counterfactuals therefore assume that each project is realized at 50% of planned total capacity by 2030. This assumption brings our totals in line with data center capacity projections from other sources, and is consistent with the phased opening plans employed by large data center campuses. Our data still imply that the grid would need to add new electricity generation exceeding the size of New England's entire

fossil fleet for the full planned data center buildout to be realized. In some versions of the forward counterfactuals, we augment the grid with proposed generation that could come online by 2030.

We find that state-level data center bans, which have been proposed or implemented in 16 states as of September 2026 (National Conference of State Legislatures 2026; Office of the Texas Governor 2026), cut electricity prices in the implementing state but can substantially increase out-of-state electricity costs and carbon emissions. For example, if Virginia had had a ban on new data center development in place 2021–2025, Virginia residents would have saved $25 million in electricity costs per year. However, the ban would have induced data centers to relocate to locations with greater price increases and emissions impacts, thus increasing residential electricity bills outside the state by $80 million and increasing total climate damages by $349 million per year. In contrast, a ban in Michigan in effect 2021–2025 would have cut CO_2 emissions by about 0.8 megatons per year, amounting to $156 million in annual avoided climate damages, while raising total residential electricity costs by only $9 million. However, most proposed bans increase the costs of data center development: the average proposed ban increases total residential electricity costs by $1 million and increases climate damages by $58 million.

By comparing future and past counterfactual scenarios, we find that emissions scale linearly with data center capacity, while price effects scale more than linearly because electricity supply curves are convex. This convexity implies that backward-looking, reduced-form approaches to estimating data center effects on prices will tend to underpredict price effects going forward. Generation growth reduces the price impacts of the data center buildout while also changing the geography of data centers' effects.

This paper contributes to an emerging literature on data centers and their impacts on society. Guidi et al. (2025); Muller (2026) and Knittel, Senga and Wang (2025) estimate the emissions costs of data centers, while other papers estimate the price effects of data centers using dispatch models (Kay, Reaser and Taylor 2026; Darmouni, Lehner and Zhang 2026) and a range of identification strategies (Benetton, Compiani and Morse 2023; Meeks et al. 2026; Watten, Bistline and Blanford 2026; Rana, Kaps and Netessine 2026; Mamkhezri, Sun and Yang 2026; Alvarez et al. 2026). Our paper is the first to explore how price and emissions effects depend on location choice and to model location choice in order to quantify the

impacts of policy interventions. We also demonstrate how location selection can complicate reduced-form approaches to estimating data centers' price effects.

We also contribute to the literature on electricity grid modeling. Other work (Borenstein, Bushnell and Wolak 2002; Cicala 2022; Gonzales, Ito and Reguant 2023; Hausman 2025; Ham, Kay and Hausman 2026; Kay, Reaser and Taylor 2026) has established how to use economic dispatch models to estimate the effects of changes in regional integration, market liberalization, and electricity demand on wholesale electricity prices. We embed the dispatch model machinery within a model of location choice, allowing electricity prices to move endogenously and allowing firms to incorporate price movement into their location decisions.

Finally, we contribute to work on spatial environmental economics. This field, summarized in Balboni and Shapiro (2025), includes research on how environmental costs vary across space (Holland et al. 2016; Borenstein and Bushnell 2022) and how location choice interacts with energy and climate (Kahn and Mansur 2013; Arkolakis and Walsh 2023; Saussay and Sato 2024; Castro-Vincenzi 2024; Hsiao 2025; Pan Fang and Greenstein 2025). We contribute an important example of how policy can reduce climate costs by reallocating firms across space.

Our paper proceeds as follows. Section 2 details the data sources we use in our analysis. In Section 3, we provide institutional details on data center development and the relationship between data centers and the grid. In Section 4, we describe our dispatch model and what we learn from it. We present estimates of the price and emissions impacts of data centers to date, along with results on how data center location choice correlates with grid features. Next, in Section 5, we describe and estimate our model of data center location choice. We use the estimates from our model to describe consumer welfare under a range of policy counterfactuals in Section 6. Section 7 concludes.

2 Data

2.1 Data centers

We use a database built in companion work, Schlauch et al. (2026), that combines data center location and characteristics from several commercially available datasets. We hand-validate location information, manually collect opening dates (and announcement dates for a subset of campuses), and merge sources to create a comprehensive dataset of 2,741 operating and planned data center campuses with a total power capacity of over 200 gigawatts (GW).

The Schlauch et al. (2026) dataset is primarily backward-looking. As a complement, and to enhance our coverage of planned data centers, we merge this database with Aterio's data center locations and power demand product (Aterio 2026). This proprietary dataset is updated daily and includes a broad set of under-construction, planned, and announced data centers. We use the dataset's status and likelihood fields to separate projects that are likely to be completed from those that are more speculative.

Because many of our analyses are at the utility-state level, we assign data centers to utility-states using campus lat-longs and utility service territory shapefiles from Homeland Infrastructure Foundation-Level Data (2023). Utility territories often overlap, so some campuses are assigned to multiple utility-states. We drop some campus-utility matches based on utility size and type, but are still left with multiple utility matches for about 20% of campuses.[2] For campuses matched to multiple utilities, we allocate campus megawatts (MW) proportional to each utility's 2016 non-residential sales.

2.2 Electricity load and generation

Data from the Energy Information Administration (EIA) allow us to track retail electricity sales over time. We use EIA 861 annual data on state and utility-state level retail electricity sales and revenues by customer class (residential, commercial, and industrial). For some analyses, we supplement with EIA 861M data, which is analogous to EIA 861 but is available at the monthly level for more recent periods, though it also excludes some smaller

[2]For campuses matched with multiple utilities, we drop matched candidate utilities that account for less than 10% of the cross-candidate sum of 2016 non-residential sales. We also drop federal utilities when a non-federal match is available.

utilities.

To measure generation costs and emissions, we combine several datasets from the U.S. Environmental Protection Agency (EPA) and EIA. EPA's Continuous Emissions Monitoring System (CEMS) gives us hourly data on emissions and generation for the universe of utility-scale fossil fuel generators in the country, which we then aggregate to the plant-level. EIA 860 gives a full inventory of power plants (including non-fossil generators), along with their location, capacity, technology, and operating years. EIA 923 provides monthly plant-level generation data including fuel type, fuel consumption, and fuel costs. EPA's eGRID dataset allows us to map individual plants to their respective regions. These datasets allow us to construct measures of plant-level costs and emissions for fossil fuel plants, which we then aggregate into regional supply curves. Plant marginal costs are computed as the product of CEMS-derived heat rates (Btu/kWh) and EIA 923 fuel costs ($/MMBtu), deflated to $2025.[3]

For additional plant characteristics, we use data from FERC Form 1 via the Public Utilities Data Liberation Project (PUDL).

2.3 Electricity prices

Retail rates

We again use EIA 861 data to examine retail rate trends over time. By dividing customer-class and year-level revenues by sales, we can compute the average retail price per megawatt-hour (MWh) of electricity over time. Although this approach does not distinguish between fixed and marginal electricity rates, it gives us a holistic sense of electricity costs for firms without requiring us to identify the precise retail tariffs that data centers are subject to in each utility. We again supplement with EIA 861M data where we need more recent information on prices.

Wholesale prices

While retail price data captures the prices that most consumers actually pay for electricity, it is wholesale prices that we would expect to respond most quickly to the changes brought by

[3] We cap CEMS heat rates at 30,000 Btu/kWh (roughly twice the least efficient coal plant operating at load) to prevent low-utilization plant-months from producing implausible marginal costs; this cap affects fewer than 2% of plant-month observations.

data centers. We use wholesale price data from Millstein, O'Shaughnessy and Wiser (2026), which aggregates nodal price data into gridcell-year level series of average prices, average prices across the 100 highest-price hours in the year, and negative price frequency. There is limited data coverage in those areas of the country not served by a Regional Transmission Organization (RTO) such as the Southeast (Figure C.6), where competitive wholesale markets do not exist.[4]

2.4 Local characteristics and miscellaneous

We gather a range of additional variables to supplement core data on data centers and electricity. We use high-resolution gridded land price estimates from Nolte (2020). From the Bureau of Economic Analysis (BEA), we collect county-level GDP data. We also incorporate market-clearing prices and volumes for the Regional Greenhouse Gas Initiative (RGGI 2026).

3 Background and descriptive facts

3.1 How data centers move electricity prices

Data center pressure on electricity prices has become extremely salient in the past few years (Goldmacher 2026; Pontecorvo 2026; Saul and Natter 2026). Average residential electricity prices have risen in real (Figure C.2a) and nominal terms (Figure C.2b) since 2021, coinciding with the recent boom in data center construction. Although this timing alone is not proof that data centers are at fault, surging data center electricity demand can quickly increase prices. By pushing the grid higher up electricity supply curves, additional demand raises wholesale electricity prices, which capture the cost of electricity generation. Utilities can pass these costs on to consumers without the need for a formal rate case. Utilities differ in how much electricity they procure on spot markets versus in longer-term contracts (see, e.g., Bray 2026), but all customers eventually face higher costs when more expensive plants are deployed more frequently. In an article titled "AI data centers are sending power bills

[4]We use "RTO" as a generic term for the entities with operational control over organized wholesale markets. Technically, FERC designates organizing entities that do not cross state lines as Independent System Operators (ISOs) instead.

soaring", Bloomberg journalists argue that wholesale price increases caused by data centers are behind the broader upward trend in residential electricity prices (Saul et al. 2025).

We focus on the wholesale market effects of data centers in this project because, on top of being salient and quick to realize, they are also easily measured and modeled. We can approximate the price effect of additional load by tracing out regional supply curves, and use variation in the shape of these curves across space to predict how price effects might differ by geography. These same supply curves also allow us to examine the emissions effects of additional load.

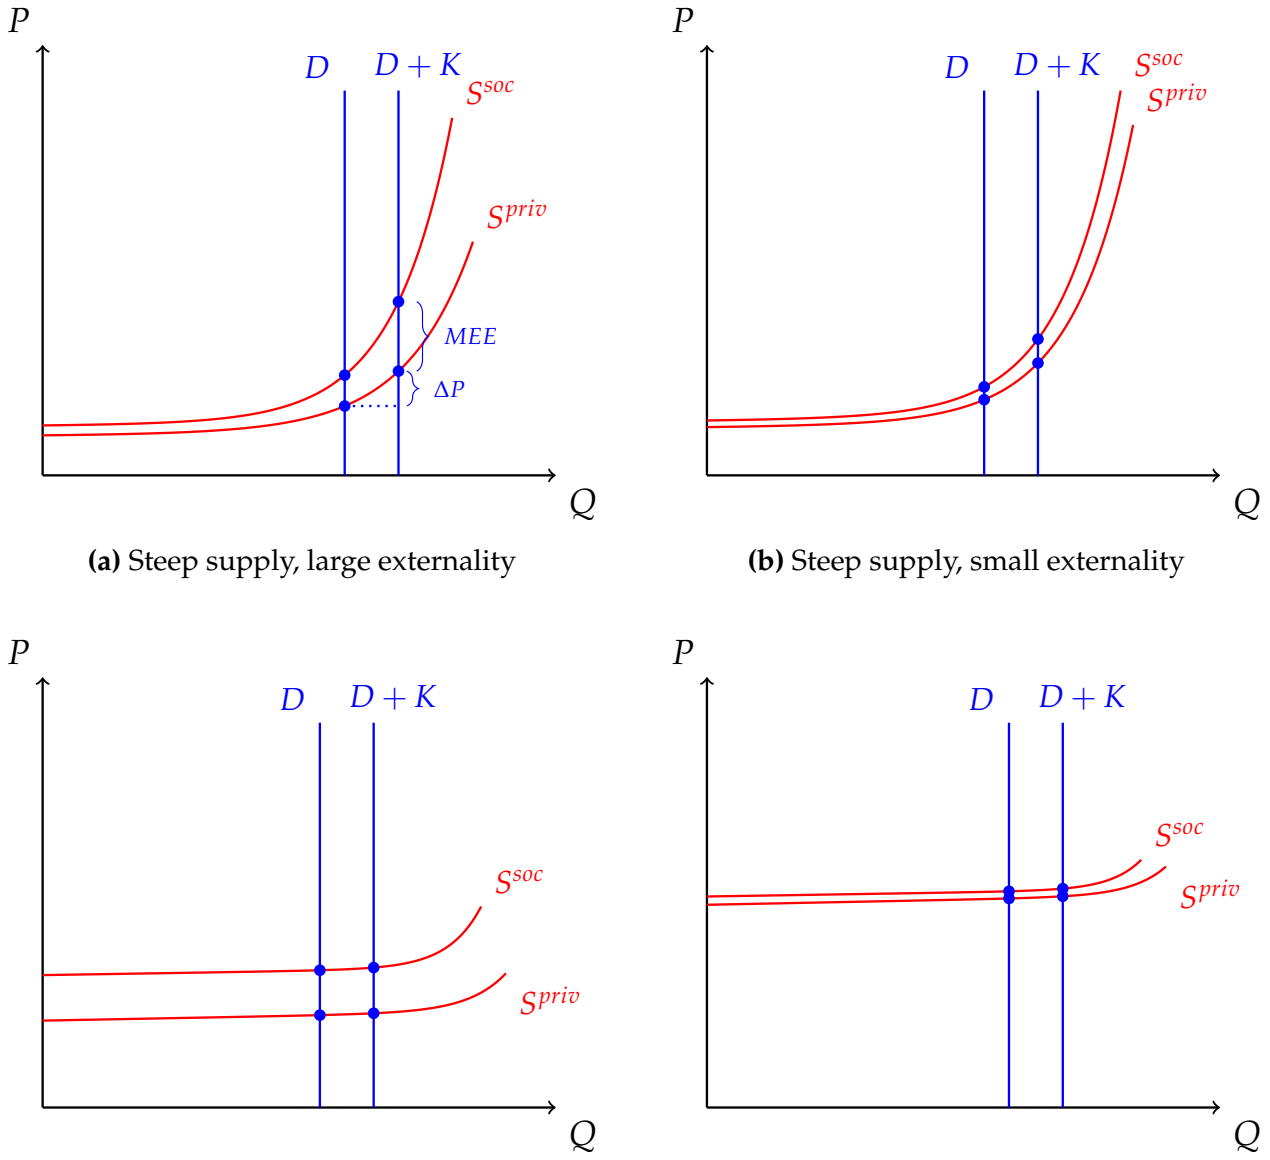
We use a stylized model (Figure 1) to demonstrate how supply curve characteristics determine the emissions and price impacts of data center entry. We illustrate four locations and one firm, building a data center with capacity K. Each location has different local supply characteristics. Supply curve slopes at baseline demand are steep in the top panels and shallow in the bottom panels, while emissions externalities are dirtier in the left panels and cleaner in the right panels.

In all panels of Figure 1, adding data center load shifts the demand curve to the right by the size of the data center load, K. This shift increases electricity prices and total emissions. However, the magnitude of these changes differs across locations. In locations (a) and (b), supply curves are steep, so ΔP is large. In locations (c) and (d), supply curves are shallow, so ΔP is small. In locations (a) and (c), emissions externalities are large. This yields large increases in climate damages, graphically illustrated as the area between S^{soc} and S^{priv} bounded by D and $D + K$.

A planner who wants to reduce the electricity cost increases and emissions associated with data center entry would prefer for the data center to choose location (d), where supply curves are shallow and emissions externalities are small. However, the firm's preferences may differ. It will choose whichever location offers the best combination of low construction costs, low electricity prices (after accounting for ΔP), and other favorable characteristics. In practice, if places with features attractive to developers also tend to have steeper slopes and larger emissions externalities (i.e., developers prefer places like Figure 1a), then we will observe that data centers have larger price and emissions impacts than we would expect by chance. If, however, the reverse is true (places attractive to developers look more like Figure 1d, with shallower supply curves and smaller emissions externalities), then we will

observe *smaller* price and emissions impacts than load growth totals alone would suggest. The correlation between location choice factors, supply curve slopes, and emissions externalities is therefore a key empirical determinant of how costly data center entry is in terms of emissions and electricity price increases.

Figure 1: Stylized impacts of data center demand on emissions and prices



(a) Steep supply, large externality

(b) Steep supply, small externality

(c) Flat supply, large externality

(d) Flat supply, small externality

Notes: Stylized electricity supply and demand in four hypothetical locations. The vertical blue lines indicate electricity demand without (D) and with ($D + K$) the addition of data centers; the added load K is the same in every panel. The red curves represent electricity supply given private costs (S^{priv}) and social costs (S^{soc}, which incorporate emissions externalities). In panel (a), ΔP marks the increase in the electricity price caused by the added demand, and MEE marks the marginal emissions externality (the vertical gap between S^{soc} and S^{priv}) at $D + K$. The top panels have steeper supply curves than the bottom panels; the left panels have larger emissions externalities than the right panels.

Data centers can affect prices through mechanisms distinct from wholesale markets too. In places with separate generation capacity markets, the scarcity created by additional load in the absence of compensating generation growth can increase the equilibrium costs of ca-

pacity. This has been particularly salient in PJM, the RTO serving much of the Mid-Atlantic region, where capacity market costs have set records since 2024.[5] However, these price impacts are also the result of PJM's particular capacity market mechanisms and interconnection backlogs, rather than the inevitable result of actual data center growth.

Other types of price impacts will be more relevant in the medium- to long-run. First, data center entry may increase the need for utilities to build or upgrade infrastructure. Some of this infrastructure will be paid for directly by data centers as part of their service contracts. Other infrastructure costs, like the transmission and distribution network serving both data centers and other customers, may be spread out across multiple regions and customer classes. Since electricity infrastructure projects take years of planning, new construction and accompanying costs may lag data center growth substantially. Before any infrastructure projects are realized, data centers could actually *decrease* infrastructure costs paid by other customers by increasing the utilization rates of existing infrastructure, effectively spreading fixed costs out over a larger rate base (Borenstein 2025). Recent policy discussions have included proposals to force data centers to pay for the incremental infrastructure made necessary by their entry. Though implementing these policies will be challenging, these measures should insulate electricity consumers from the costs of upgrading the grid to serve data centers. These policies will not, however, protect consumers from increases in generation costs due to data center electricity demand.

3.2 Data center growth and location choice factors

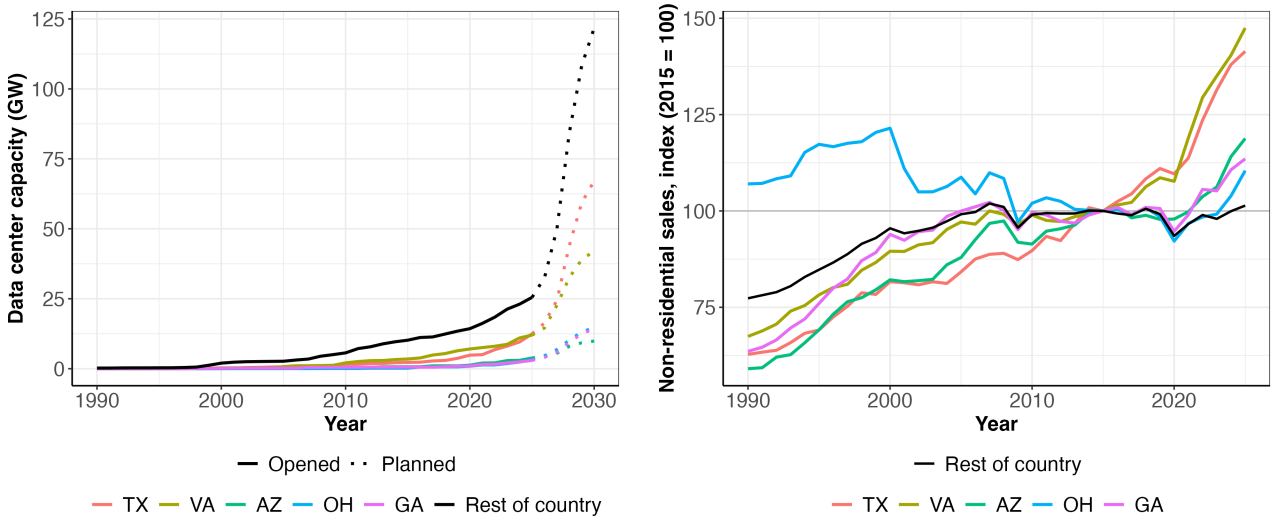
The data center buildout has taken off in recent years. In 2025 alone, 59 data center campuses opened, representing about 8.6 GW of total power capacity or the equivalent of about 5 million households.[6] This growth has been geographically concentrated. Just five states account for 58% of all data center capacity to date, and Texas and Virginia each account for about 20% of total capacity (Figure 2a). Data center growth is large relative to baseline de-

[5]The PJM market monitor found that projected data center load growth was the cause of 40% of the $16.4 billion in costs that PJM incurred in its December 2025 auction (Howland 2026). Capacity market cost increases led to wholesale price increases of about 0.9 cents per kilowatt-hour (kWh) in 2025 relative to 2024 (Wiser et al. 2026).

[6]The average U.S. household consumed 10,791 kWh of electricity in 2022 (U.S. Energy Information Administration 2024), suggesting that 800,000 households together have a power draw of about 1 GW. We assume data center power utilization at 70% of nameplate capacity.

mand. In Virginia, data centers now account for over a quarter of the state's total electricity consumption (EPRI 2024). Top data center-hosting states have seen their non-residential electricity sales increase dramatically in recent years (Figure 2b).

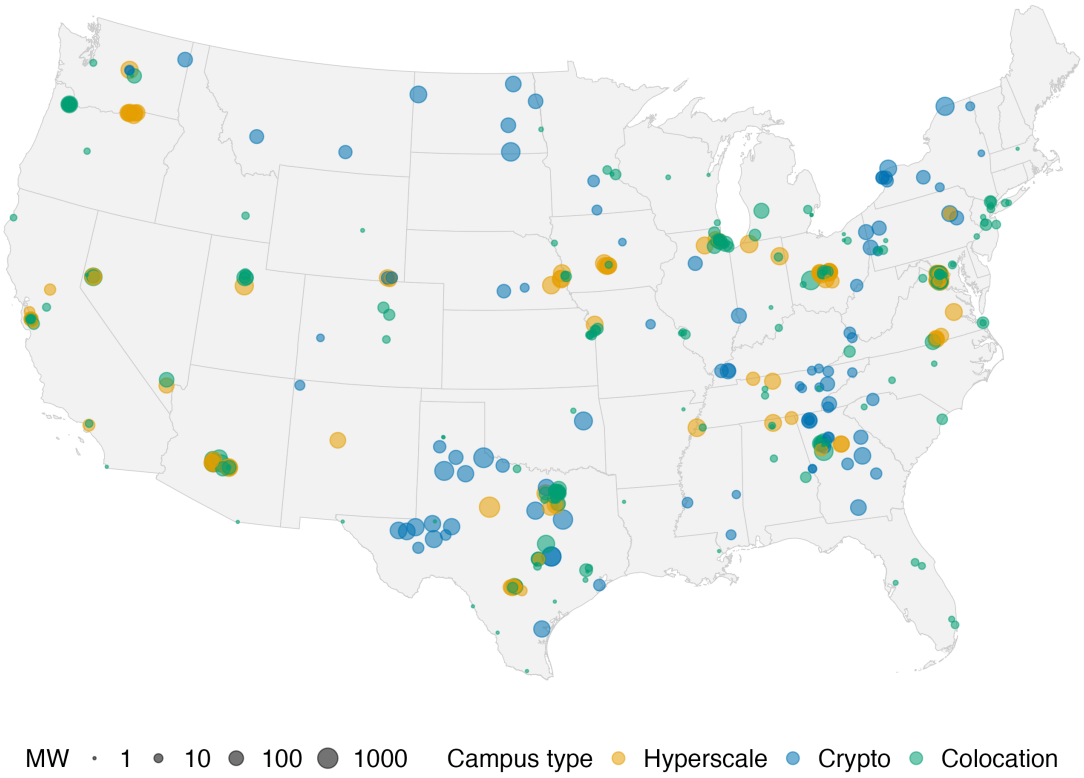
Figure 2: Data center capacity and non-residential sales for the top states by data center growth



(a) Cumulative data center capacity (GW) **(b)** Non-residential sales (indexed to 2015)

Notes: The five states with the largest 2020–2025 data center buildout (TX, VA, AZ, OH, GA); the black line is the rest-of-country aggregate (all other states summed). Panel (a): cumulative data center capacity in GW, each campus assigned to a state and capacity accumulated by reported opening year. Data on opened campuses from Schlauch et al. (2026), and data on planned campuses from Aterio. Panel (b): non-residential (all sectors except residential) retail sales indexed to 2015 = 100, from EIA-861M monthly data aggregated to calendar years (2025 preliminary); the rest-of-country line is all other states' summed non-residential sales, also indexed to 2015.

Figure 3: Data center campus locations by type, 2016–2025



Notes: Data center campuses that opened 2016–2025; data from Schlauch et al. (2026). Campuses without location, type, or power capacity information are excluded.

What makes a location attractive depends in part on the kind of data center being developed. Colocation data centers rent out space and equipment to commercial users. These users are often sensitive to latency (lags in response time), leading colocation data centers to cluster around major population centers, where they are closer to end users and high-speed fiber connections (Figure 3). Hyperscale data centers are operated by individual companies, typically large tech companies that sell internet services or operate platforms (Amazon, Microsoft, Google, Meta, etc.). Although some of their use cases are latency-sensitive, these companies can distribute tasks across their facilities so that they do not need to concentrate their development profiles in expensive, urban markets.[7] Cryptomining data centers are even less concerned with latency. Since they do not need to deliver large amounts of

[7]On this topic, an executive at real estate developer CBRE notes that "everyone is optimizing their data center footprint. Historically, they may have wanted to be in latency markets. But some of the storage files don't need to be in the most expensive real estate space in the world. We are seeing smarter and smarter optimization of workloads as AI takes hold" (Bruns 2026).

data to any customers, they can more easily locate in remote or undeveloped areas like North Dakota and West Texas (Figure 3). They also do not have the extreme reliability requirements of hyperscale and colocation facilities, which typically target uptime rates over 99.9% (M. 2026). As a result, cryptomining companies are more willing to build in places with volatile electricity prices, where demand response programs and flexible operations can yield cost savings.[8]

However, distinctions between types of data centers are blurring as more data center capacity is dedicated to artificial intelligence (AI). AI is estimated to account for over 50% of data center load by 2028 (Shehabi et al. 2024). Training AI models requires huge amounts of power, but does not depend on low-latency connections with end users. As a result, electricity availability has come to be the limiting factor in recent years. Said one industry insider, "Site selection has been converted to a location of power game. It used to be more multi-factor and based on network connectivity to the grid and access to water. But today, given the size of data center developments, power is really factors 1, 2 and 3" (Bruns 2026). Even for an AWS executive who was willing to list relevant factors other than power, electricity still literally made the list as items one through four: "We choose locations based on electricity availability; access to renewable power; redundant power; stability of the grid; proximity to an existing cluster of similar projects; and a welcoming community" (Starner 2023).

Speed-to-power is another major factor in location decisions. The AI boom has made data center construction delays extremely costly. One industry expert estimated the value of getting a 1 GW data center online one year faster at $7 billion (Roberts 2026). While data centers themselves take between 12 and 24 months to build, utility projects can take three to five years from planning to delivery (Krishnamoorthy, Westfallen and McKenna 2026). Thus, developers can save time by building in places where they can connect to the grid without needing to wait for lengthy transmission or substation upgrades. Power system modeling suggests that areas surrounding metropolitan areas and other areas of dense demand are most likely to already have the infrastructure to accommodate new large loads (Kim, Dong and Xie 2026). Outside of seeking out these kinds of sites, developers are

[8]Riot Platforms provides an extreme example of this. In August 2023, the company earned $31.7 million (more than 3.5x its earnings from selling Bitcoin) by curtailing its electricity usage and selling power back to the Texas grid (Riot Platforms 2023).

also building their own behind-the-meter (BTM) power to avoid interconnecting entirely.

Electricity is not only the limiting factor in where data centers can be built, but also these facilities' largest ongoing expense. Once they connect to the grid, data centers consume vast amounts of electricity, running up massive power bills. Training OpenAI's GPT-4, for instance, took an estimated 50 gigawatt-hours (GWh), or about as much electricity as 5,000 U.S. households consume in a year (O'Donnell and Crownhart 2025; U.S. Energy Information Administration 2023). At national median non-residential electricity prices, this electricity would cost about $5 million. However, the bill could have been either smaller or much larger: in Texas, it would have cost about $4 million, and in California, about $12 million. One study estimates that electricity accounts for about 66% of a large hyperscale data center's operating costs (Michael and Cottier 2026).

An additional location consideration is whether other data centers have already been planned or built in a given place. In some cases, data center development itself can make locations more attractive. In West Des Moines, for instance, local officials cite tax revenues from data centers as the reason they were able to invest in new roads and fiber (Moffatt 2024). Pre-existing development can also signal lower costs. One data center executive told us that locating near an existing data center provided assurance that someone else had cleared the hurdles necessary for data center construction (permitting, skilled workforce, nearby fiber connections, etc.), making it both cheaper and less risky to build. However, too much concentration raises the risk that different developers will be competing for power or skilled construction workers, raising costs.

Together, the factors above paint a coherent picture of a data center's ideal location. It has cheap electricity and some prior data center development. It already serves a large load base, so that it has the electric infrastructure to connect large loads without time-intensive upgrades. Since locations with large loads tend to have higher levels of slack generation capacity, such a location is likely to have a flatter supply curve slope, making it look like Figure 1c or 1d. We will test this hypothesis in Section 4.3.

One final factor that we will not consider in this project is local data center subsidies. These subsidies can be large, but the accompanying rules can be onerous. According to one insider we spoke to, the requirements that must be met to qualify for some subsidies are demanding enough that some developers do not pursue them even if they might be

eligible. Table D.7 lists the largest state subsidies according to Good Jobs First, which we join with our data in order to normalize subsidies by total data center capacity. Subsidies are, in most cases, smaller than the annual per-MW cost of electricity by a factor of four to five. Furthermore, subsidies are endogenous to data center interest in a location, and not an appropriate predictor of location choice without careful treatment of cross-locality competition, as in Slattery (2025).

4 Dispatch model

Our first goal is to characterize how data centers increase the price of and emissions from electricity generation. We do this by using a dispatch model, a common tool in the energy literature (Cicala 2022; Hausman 2025; Ham, Kay and Hausman 2026), to build stylized electricity supply curves across regions.

The dispatch model approach has several advantages over a reduced-form approach to measuring data center effects on prices. First, reduced-form approaches are not well-suited to this setting. As explored in Appendix A, data centers' spatial concentration in a handful of large utilities both decreases the number of treated units and makes defining an appropriate control group very challenging. Furthermore, price effects are nonlinear due to the convexity of electricity supply curves. Even if reduced-form approaches successfully recovered the past effect of data centers on prices, we would not be able to infer from these results how data centers will impact prices going forward.

A major limitation of the dispatch model approach is that it captures only changes in wholesale electricity prices, i.e., the cost of generating electricity. Retail electricity prices move not only with generation costs, but also the costs of building and maintaining electricity infrastructure, along with miscellaneous taxes and fees. However, as noted in Section 3.1, policy can change the relationship between data center growth and infrastructure costs by pinning these costs on the data centers themselves. Modeling how electricity infrastructure costs evolve as load scales and policies change is a difficult problem, and out of scope for this project.

Another limitation is that the dispatch model holds fixed the power plants that comprise each region's supply curves. In the counterfactuals section, we explore robustness to

this assumption by adding planned generators and calculating how much this changes our estimates of price and emissions effects. In future extensions, we plan to estimate versions of the dispatch model with endogenous power plant entry, using a range of plausible entry elasticities.

4.1 Setup

We model dispatch at the plant-fuel group level.[9] Since most plants consist of just one plant-fuel group, we refer to this unit as a "plant" for brevity. We restrict our attention to only fossil fuel plants, since these are the plants that respond on the margin in most region-hours.[10] In each region, we line up plants in order of ascending marginal cost. We call this the "merit order". We then "dispatch" plants until total demand is met. More formally, we construct a plant-level measure of marginal cost:

$$mc_{g,t} = \text{heat rate}_{g,t} \times \text{fuel price}_{g,t} + \text{permit costs}_{g,t} \times \text{emissions rate}_{g,t} \tag{1}$$

where g indexes plants and t indexes time. Heat rate captures the efficiency of each plant (which we measure in million British thermal units (MMBtu) per MWh from CEMS data), fuel price is the cost of the fuel used by each plant (measured in dollars per MMBtu from EIA 923), and permit costs are the cost of emissions permits through several emissions trading schemes (RGGI for CO_2, CSAPR for SO_2 and NO_x).[11] Emissions rates are calculated from CEMS data in units of pollutant weight per MWh. We observe heat rates hourly and fuel prices monthly, but since heat rates move very little within plant over time, we construct $mc_{g,t}$ as a monthly variable.[12]

We define the region-level merit order at time t by taking all the plants in a given region (regions assigned based on eGRID data and geographic location) and sorting them in in-

[9] In practice, dispatch is actually done at the generator level. The typical plant has around three generators.

[10] In future, we plan to add dispatchable hydropower plants to this set.

[11] We do not include California's carbon permit costs in our marginal cost calculations because California taxes imported electricity as well as domestically generated electricity. It would not be possible to correctly factor in these permit costs into the dispatch model without also modeling inter-state electricity imports.

[12] We make a few assumptions to correct for extreme or missing values. We cap heat rates at 30,000 Btu/kWh, and treat plant-month fuel costs outside the 1st–99th percentiles for a fuel-year as missing, then impute them from state-month and national-month fuel averages.

creasing $mc_{g,t}$ order. This forms supply curve $S_{r,t}(\cdot)$, where r indexes regions. For a given demand level $D_{r,t}$, we define the market-clearing price (MCP) to be the marginal cost of the last plant dispatched, $S_{r,t}(D_{r,t})$. This MCP is our proxy for wholesale electricity prices.[13]

We assume that non-data center demand is inelastic, allowing us to estimate the change in MCP induced by additional data center load $K_{r,t}$ to be:

$$\Delta MCP_{r,t}(K_{r,t}) = S_{r,t}(D_{r,t}) - S_{r,t}(D_{r,t} - K_{r,t}) \tag{2}$$

We can estimate marginal emissions from additional load in a similar way. Let $E_{r,t}(D)$ be marginal emissions at demand level D. Emissions from additional load $K_{r,t}$ are:

$$\Delta CO_{2,r,t}(K_{r,t}) = \int_{D_{r,t}-K_{r,t}}^{D_{r,t}} E_{r,t}(q)\, dq \tag{3}$$

As in Hausman (2025) and Ham, Kay and Hausman (2026), we define $D_{r,t}$ to be the total observed generation of fossil fuel plants in region r and hour t. This is equivalent to assuming that in-region fossil plants are always the marginal generators. Although this is true in general, it does not hold in places with significant renewable generation, or where electricity imports from other regions are important. Our model will overpredict costs and underpredict available generation capacity in these cases. Finally, we assume that there is no transmission congestion within region, meaning that any plant can respond to load changes in a given region, regardless of where that load occurs. In practice, congestion on transmission lines constrains which plants can serve new load.

We model dispatch for each of 11 regions (Figure C.10) defined based on Cambium's Generation and Emission Assessment (GEA) regions, which themselves are built from EPA's eGRID regions (Gagnon, Cowiestoll and Schwarz 2023). We preserve those GEA regions that represent the territories of RTOs, merging some subregions (e.g., SPP North and South). We then create regional aggregates from the remaining GEA regions. Our goal is to match real-life dispatch patterns where possible while approximating the integrated (though less organized) dispatch processes in non-RTO regions.

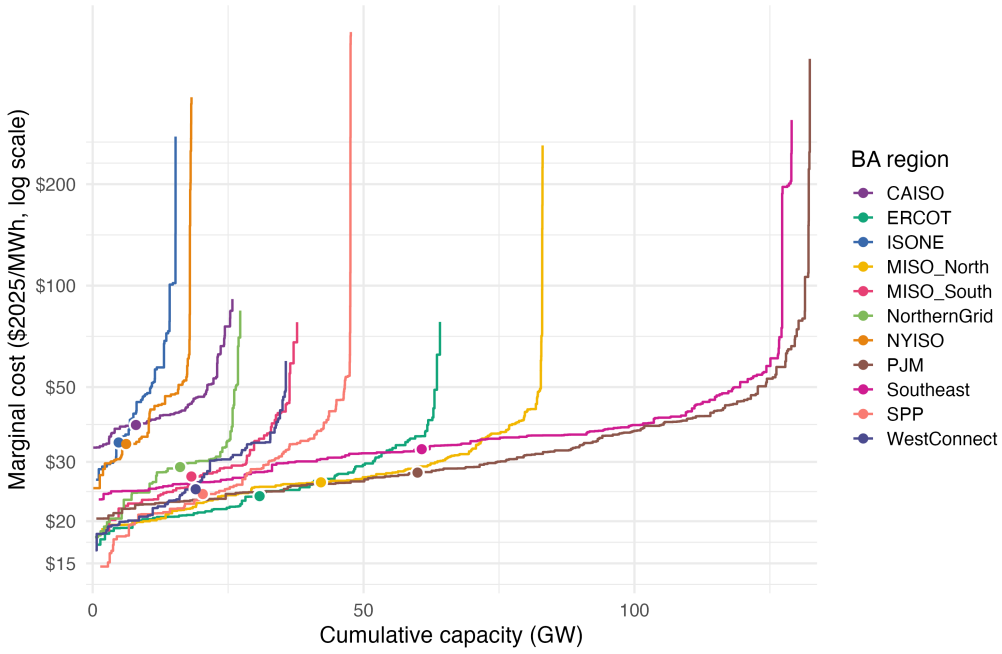
[13]This approximation does more work in areas without organized wholesale markets, where retail rates might reflect average generation costs rather than marginal. In future iterations, we will explore the impact of using average costs instead of marginal in these markets.

4.2 Supply curves

Regions vary significantly in the level, shape, and composition of their supply curves. Figure 4 overlays the 2019 supply curves for all 11 regions. CAISO, ISO-NE, and NYISO have higher prices and less capacity than the other RTOs, though all three are major importers of electricity (which we do not directly model). Generation capacity is far higher in PJM and the Southeast than in any other region. Smaller regions and regions with higher price levels tend to have steeper slopes, though this is not universal. In all regions, slopes are generally increasing in demand, particularly at the highest demand levels where peaker plants (which are less efficient and thus more costly) dominate. Cross-region differences in supply curve slopes are meaningful: in 2025, adding 1 GW of electricity demand in ISO-NE would raise prices by over $3 per MWh, while the same additional load in MISO-North would raise prices by just $0.30 per MWh (Figure C.4). The ordering of regions by marginal supply curve slope is relatively stable over time, though fuel price changes have different impacts by region.

Marginal emissions over the supply curve vary non-monotonically (Figure C.3). More expensive plants are generally less efficient, which increases emissions, but this variation is drowned out by emissions variability coming from fuel-switching. Coal generation is far more emissions-intensive than gas generation (Figure C.11), but can be comparable in cost. Regions differ in where coal falls in their dispatch: coal anchors the entire supply curve for MISO-North, is interspersed through PJM's curve, and is not used at all in CAISO (Figure C.12). In places with less or no coal usage (including NYISO and CAISO), emissions generally increase over the supply curve, but in more coal-dependent regions, emissions both rise and fall with capacity. Swings in the relative fuel prices of gas and coal reorder which fuel sits on the margin, creating large year-to-year variation in marginal emissions in the regions with both types of generation (Figure C.5). However, marginal emissions are consistently highest in coal-dependent regions like SPP and MISO-North and lowest in gas-dominant regions (ISO-NE, CAISO, NYISO). In 2025, an additional 1 GW of electricity demand would result in 3.7 megatons (million metric tons) of CO_2 in ISO-NE and 6.5 megatons in MISO-North.

Figure 4: 2019 regional dispatch model supply curves



Notes: Supply curves for each of the 11 dispatch regions, 2019. Capacity calculated based on EPA CEMS hourly data (2019). Marginal cost is on a log scale ($2025/MWh). The dot on each curve marks the region's mean hourly demand in 2019 (the marginal cost at average load).

The prices we estimate generally track annual wholesale price averages measured using Millstein, O'Shaughnessy and Wiser (2026) data (Figure C.13). Regressing measured on modeled prices yields a coefficient near one (Table D.8). Some discrepancies arise for reasons including our no-congestion assumption, the exclusion of non-fossil plants and some cost factors (such as operation and maintenance costs), and idiosyncratic regional factors (including Winter Storm Uri in 2021 and coal shortages that plagued MISO in 2022). Nevertheless, the agreement across the two series is comforting given that the dispatch model we use to estimate prices does not incorporate any wholesale price data.

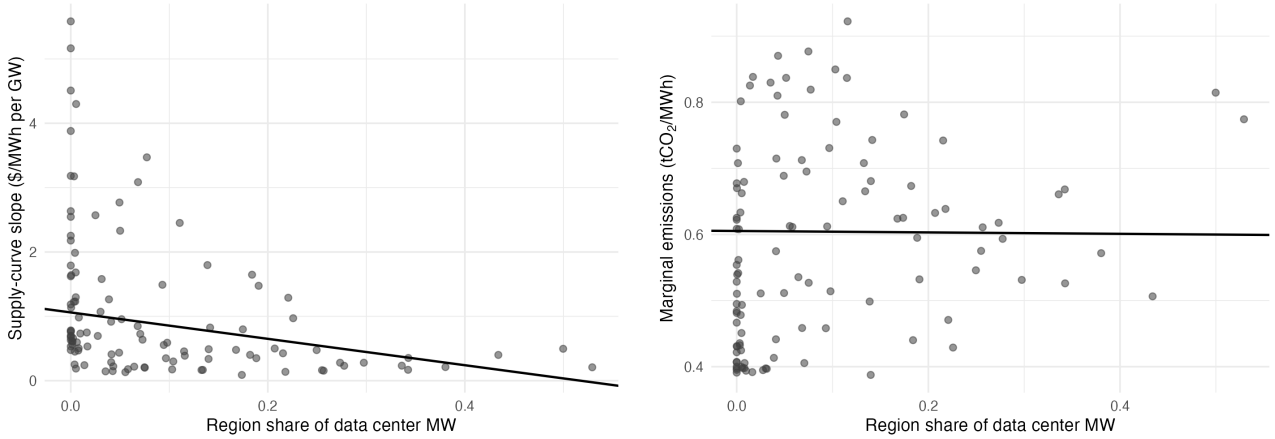
4.3 Data center location choice and grid characteristics

By looking at the correlation between data center placement and supply curve features, we can infer what kinds of grids data centers tend to prefer, in a real-world analog to the stylized model in Figure 1. We calculate region-year-level correlations between data center additions and grid characteristics between 2016 and 2025. For each year of data center entry, we calculate supply curve slopes and marginal emissions two years beforehand so that grid

characteristics are fully independent of the data center buildout.

We find that the correlation between data center additions and supply curve slope is negative and significant (Figure 5a), meaning that data centers tend to build in places with shallower slopes. Our results suggest that a 10 percentage point increase in a region's data center share is associated with a 21 cent smaller price increase from an additional 1 GW of load. This implies that data centers are disproportionately locating in places where their impact on prices will be smaller. In contrast, the relationship between data center buildout and emissions intensity is indistinguishable from zero (Figure 5b).

Figure 5: Data center concentration and grid characteristics



(a) Supply-curve slope **(b)** Marginal emissions intensity

Notes: Each point is a region-year, where region is one of the 11 dispatch regions and year is in the set 2016–2025. On the horizontal axis, we plot each region's share of the entry year's data center capacity (MW), summing to one across regions within each year. Each cohort's x-axis values are paired with y-axis values from the dispatch year two years before. In Panel (a)'s vertical axis, we report the increase in the market-clearing price from adding 1 GW to each hour's demand, averaged over the dispatch year's hours weighted by demand. Panel (b)'s vertical axis plots the mean over the dispatch year's hours of the incremental generation emissions from the same 1 GW addition (metric tCO_2/MWh). February 2021 hours are excluded for ERCOT and SPP. Lines are least-squares fits pooled across cohorts, weighted by the region-year's cohort capacity.

4.4 Realized data center impacts

We calculate data centers' aggregate effects on prices and emissions by combining each region's price and emissions curves with a region-level estimate of data center load. We construct a measure of data center electricity demand by summing the capacity of all data centers in a given region-year. We convert data center capacity to data center demand by assuming that every campus operates full-time beginning in the year we observe it opening,

and that electricity demand is equal to 70% of nameplate capacity.[14] We then apply Equations 2 and 3 to estimate $\Delta MCP_{r,t}$ and $\Delta CO_{2,r,t}$ for each region r and year t. We average region-level effects to the national level, weighting each region by observed demand ($D_{r,t}$).

We compare the actual buildout and its price and emissions effects to the geography and impacts of a buildout under random data center placement and under emissions and price-minimizing or -maximizing placements. We estimate the $\Delta MCP_{r,t}$ and $\Delta CO_{2,r,t}$ induced by these alternate data center placements by calculating the effect of their additional load on a data center-free baseline:

$$\Delta MCP_{r,t}(\tilde{K}_{r,t}) = S_{r,t}(D_{r,t} - K_{r,t} + \tilde{K}_{r,t}) - S_{r,t}(D_{r,t} - K_{r,t})$$
$$\Delta CO_{2,r,t}(\tilde{K}_{r,t}) = \int_{D_{r,t}-K_{r,t}}^{D_{r,t}-K_{r,t}+\tilde{K}_{r,t}} E_{r,t}(q)\, dq$$

for actual data center demand $K_{r,t}$ and hypothetical data center demand $\tilde{K}_{r,t}$ in region r and time t. We hold fixed the timing of entry, so that in each year, campuses that enter in the actual time series also enter in each counterfactual allocation. We construct random placements as campus-level draws from the set of 11 regions, weighting the draws by each region's area. We define emissions and price objectives that maximize or minimize $\sum_r \Delta CO_{2,r,t}$ and $\sum_r \phi_r \Delta MCP_{r,t}$ where ϕ_r is each region's demand weight. Our optimization is myopic. In each year, we make maximizing or minimizing allocations given that year's $S_{r,t}$, $D_{r,t} - K_{r,t}$, and capacity from the campuses placed in previous years, but not considering the future evolution of these variables. The counterfactuals reallocate all campuses, including those built before 2016; in the optimizations, we allocate pre-2016 campuses alongside 2016 campuses based on 2016 grid characteristics. Across counterfactual allocations, we enforce a capacity screen that prevents campuses from being placed in regions where there is insufficient generation capacity.[15]

Trends in cumulative data center capacity (Figure C.14) reveal the same pattern of geographic concentration we noted at the state level (Figure 2) and highlight where this concen-

[14] The existing literature uses assumptions of data center utilization between 50% and 80% (Norris 2025), but public data on actual utilization do not exist. True utilization rates have likely changed over time and vary with data center type (Shehabi et al. 2024).

[15] We reject data center placements in region-years where the additional load would push total demand above total thermal generation capacity in the 99.9th percentile hour for that year.

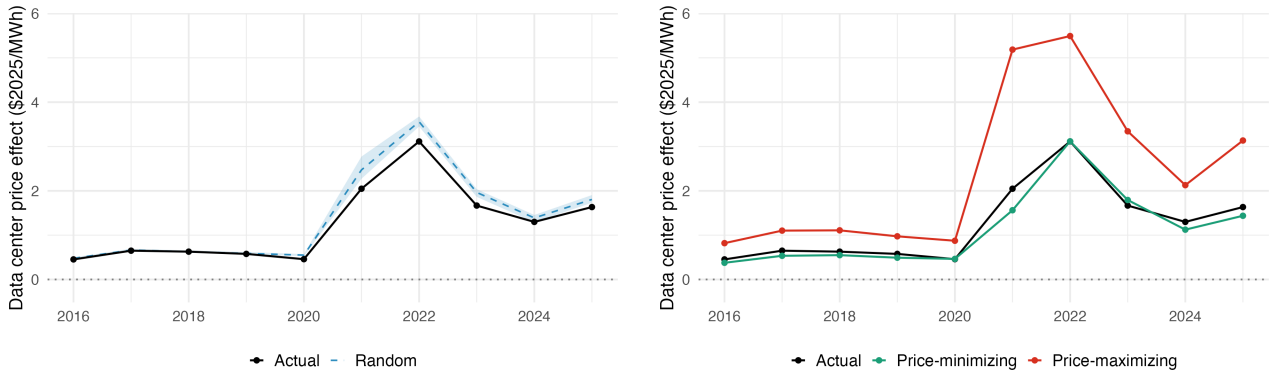
tration diverges from random. In ERCOT and PJM, which encompass Texas and Virginia, actual data center growth far exceeds growth under the random allocation. In PJM alone, data center capacity implies total data center electricity consumption of nearly 10 terawatt-hours (TWh) per month, about as much as 11 million households and over four times the consumption that would have occurred under a random allocation. Conversely, data center construction in California, New England, and most of the Midwest is below what we would expect under random growth.

We find that, depending on its geography, the cumulative data center buildout could have increased average wholesale prices in 2025 by as much as $3.13 per MWh or as little as $1.44 per MWh (Figure 6b). Scaling by total residential electricity consumption in 2019 and assuming full cost pass-through, these figures represent either $2.1 billion or $4.5 billion in extra annual costs for residential customers. The actual buildout's price effect in 2025 was $1.63, far closer to the minimum than the maximum. This corresponds to an annual cost to residential customers of $2.3 billion, and an average price effect of $0.33–0.41/MWh per GW of nameplate data center capacity.[16]. Price effects vary by region (Figure C.15), but are on the order of a few dollars per MWh even in regions like ERCOT and PJM that have substantial data center capacity.[17] The national average of the data center price effect is smaller than the median under area-weighted randomization ($1.80), consistent with data centers preferring shallow-slope locations.

[16] We calculate this average by dividing the dispatch model's total price wedge by total data center capacity, summed across region-years. The lower bound excludes the NorthernGrid, where our thermal-only dispatch model overstates price effects.

[17] Although estimated impacts in the NorthernGrid are larger (surpassing $10 in some years), we do not believe we are accurately capturing price dynamics in this region, where hydropower makes up over half of total generation. We plan to correct this in future iterations of our model.

Figure 6: Cumulative data center wholesale price effects



(a) Actual vs. random

(b) Actual vs. min and max

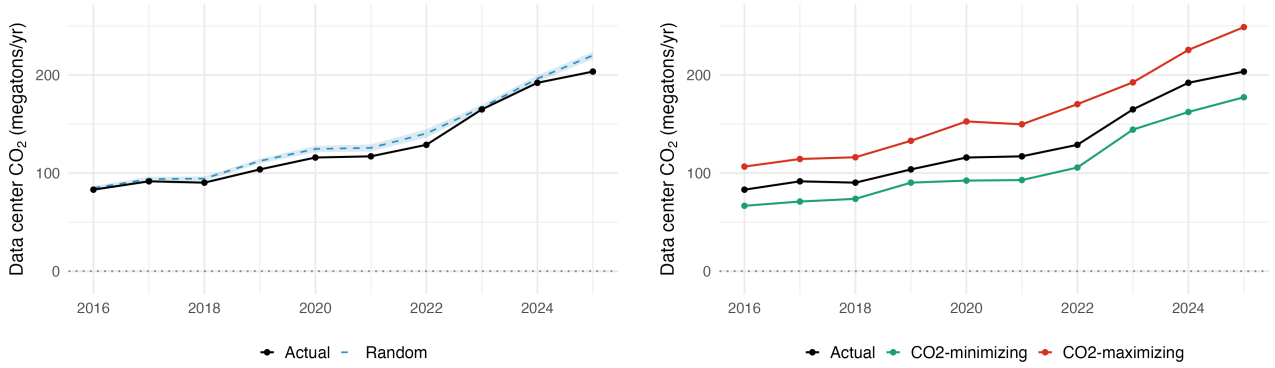
Notes: Annual wholesale price effects of cumulative data center load (Schlauch et al. 2026), evaluated year by year on each year's own dispatch grid. Price effects are demand-weighted across hours, months, and regions. Actual: nets the additions' load out of observed regional demand. Random: all campuses assigned randomly to regions with probability proportional to region geographic area, entering in their actual opening years; bands are 90% intervals over 2,000 draws, line indicates median across draws. Min/max: all data center load is placed to minimize/maximize the cumulative demand-weighted price effect. All hypothetical allocations are subject to a generation capacity screen. Data center load = nameplate $\times$ 0.70. February 2021 is excluded for ERCOT and SPP (Winter Storm Uri).

Assuming a social cost of carbon of $185 per metric ton (Rennert et al. 2022), we find that the emissions costs of data centers are an order of magnitude larger than their price effects. In 2025, total CO_2 emissions from data centers equaled 203.5 megatons, equivalent to $37.6 billion in climate damages. Again, if data centers were distributed differently, costs could have been far higher: the climate damage-maximizing allocation created emissions of 249 megatons and climate damages of $46 billion in 2025. The climate damage-minimizing allocation reduces data centers' climate cost, but still yields climate damages over $30 billion. Although data centers do not systematically prefer locations with lower marginal emissions (Figure 5b), their emissions are nearer to the climate damage-minimizing allocation than the maximizing allocation, and significantly below the emissions under the random scenario. This pattern is driven by data centers' extreme preference for ERCOT and PJM, which both have relatively low emissions marginal generation. Under random allocation, data center capacity is shifted out of these regions and into MISO-North and SPP (Figure C.14), the two regions with the highest marginal emissions intensity in 2025 (Figure C.5).

Together, our results indicate that data center location choice matters. Reallocating the past few decades of data center growth across space could have, in the worst case scenarios,

increased residential customer costs by 92% or increased climate damages by 22% relative to their actual effects. Over the same period, reallocation could not have significantly shrunk data center price effects through relocation because, consistent with our results in Figure 5a, facilities were already choosing relatively low-impact locations. Although it would have been possible for the data center buildout to emit 13% less CO_2, such an allocation would have come at the cost of 39% larger electricity price effects (Figure C.22).

Figure 7: Cumulative data center emissions effects



(a) Actual vs. random

(b) Actual vs. min and max

Notes: CO_2 effects of cumulative data center load, evaluated year by year on each year's own dispatch grid. Actual: nets the additions' load out of observed regional demand. Random: campuses assigned randomly to regions with probability proportional to region geographic area, entering in their actual opening years, subject to a capacity screen; bands are 90% intervals over 2,000 draws, line indicates median across draws. Min/max: all data center load is placed to minimize/maximize total CO_2 emissions. All hypothetical allocations are subject to a generation capacity screen. Data center load = nameplate $\times$ 0.70. February 2021 is excluded for ERCOT and SPP (Winter Storm Uri).

5 Location choice model

Our dispatch model results indicate that data center effects on prices and emissions vary with where those data centers are located. In order to explore how policy can shape data center geography, and by extension their impacts on prices and emissions, we model the data center location choice. This is a complicated problem because data centers are long-lived assets, data center costs and revenues have changed rapidly over just the past few years, and the industry is concentrated in the hands of a few major players.

We make several abstractions to make the problem more tractable and allow us to focus more narrowly on tradeoffs across locations. First, we model the intensive margin decision

of where to construct data centers to minimize costs, holding fixed the extensive margin of how much total data center capacity is built. Statements from major tech firms with specific compute targets and timelines (e.g., Musk 2025) suggest that the cost-minimization framing is reasonable. In a future extension, we plan to model the extensive margin by considering how policies affect total investment in data center capacity.

Second, we model entry as a static decision within a series of five-year periods. Entrants decide where to locate by inferring lifetime costs using contemporaneous information (including the choice strategies of same-period peers) but do not anticipate future changes in costs. This is a less consequential assumption than it would be in other industries because of tremendous uncertainty and high discount rates: forecasts of future data center capacity are unstable, sometimes revised dramatically over the space of a few months,[18] and development-stage discount rates for data center projects are around 17–20% (Houlihan Lokey 2024). If data centers take two years to build, then this discount rate implies that firms value costs five years after opening at only about 30% of face value. In a future extension, we will quantify the gap between the static version of the location choice problem and a full dynamic version.

Finally, we model each campus as independently located, conditional on a firm's pre-existing footprint. Modeling within-firm interdependence of siting decisions, as in Castro-Vincenzi (2024)'s treatment of the auto industry, requires significantly more structure. However, data center operators often rely on third-party development firms to find and prepare sites by securing power, permits, and necessary infrastructure (Grant 2024). Large tech companies then either buy or lease sites from the original developer. This decentralizes the location choice decision, making campus-level location choice more plausible.

5.1 Firm problem

We follow existing work in the literature by modeling firm location choice as a discrete choice problem (Guimarães, Figueirdo and Woodward 2003; Devereux, Griffith and Simpson 2007; Brülhart, Jametti and Schmidheiny 2012; Suárez Serrato and Zidar 2016; Fajgel-

[18] For example, BloombergNEF revised its estimate of 2035 data center capacity upward 36% between its April 2025 and December 2025 forecasts, then upward another 83% for its July 2026 forecast (BloombergNEF 2026 2025).

baum et al. 2019). We assume that there exists a fixed stock of data center campuses $\mathcal{C}$ indexed by c choosing between locations $i \in \mathcal{I}$. Each campus has exogenous characteristics: it is built in period $t(c)$, has size K_c, and is operated by firm $n(c)$. Each location i belongs to dispatch region $r(i)$. Campuses choose where to locate based on features of each location, plus an idiosyncratic campus-location taste shock, as observed prior to entry period $t(c)$. Campuses with the same $t(c)$ make location decisions simultaneously.

The value of entering in i depends on local electricity prices, which depend on total entry in $r(i)$ because wholesale markets clear at the dispatch region-level. For $\mathbf{L}_t$ the length-$|\mathcal{I}|$ vector of data center capacity entry across locations in period t, we define price function $p_{i,t}(\mathbf{L}_t)$ that maps entry into the retail electricity price faced by firms entering in location i and period t. We construct this function offline based on the dispatch model approach from Section 4. We assume that campuses entering in period t expect to face $p_{i,t}(\mathbf{L}_t)$ for their entire lifetime, which we set at 10 years.[19] We convert $p_{i,t}(\mathbf{L}_t)$, a per-MWh price, into a per-MW lifetime electricity cost $\tilde{p}_{i,t}(\mathbf{L}_t)$:

$$\tilde{p}_{i,t}(\mathbf{L}_t) = v \cdot \sum_{m=1}^{120} \delta^m p_{i,t}(\mathbf{L}_t) \tag{4}$$

for monthly discount factor δ, defining v to be the monthly MWh electricity usage for a 1 MW data center.

We define the utility from building c in location i to be (suppressing the c-dependence of t and n for conciseness):

$$V_{ci} = \beta_0 K_c \tilde{p}_{i,t}(\mathbf{L}_t) + \boldsymbol{\beta}' \mathbf{X}_{i,t} + \rho H_{ni,t} \tag{5}$$

$$U_{ci} = V_{ci} + \varepsilon_{ci} \tag{6}$$

where $K_c \tilde{p}_{i,t}(\mathbf{L}_t)$ captures lifetime electricity costs for a K_c-sized campus, $H_{ni,t} \in \{0, 1\}$ is an indicator for whether firm n already has a campus in i, $\mathbf{X}_{i,t}$ is a vector of other location-time characteristics, and ε_{ci} is a Type I extreme value error that is i.i.d. across campuses and locations.

[19] Estimates from Michael and Cottier (2026) place the lifespan of an AI data center's IT equipment at three to seven years and the facility lifespan at 14 years. IT equipment accounts for around 60% of upfront capex, compared to about 18% for the physical facility.

Because of our error distribution assumption, the choice probability of campus c in location i and period t takes the conditional logit form:

$$Pr_{ci} = \frac{\exp(\beta_0 K_c \tilde{p}_{i,t}(\mathbf{L}_t) + \boldsymbol{\beta}' \mathbf{X}_{i,t} + \rho H_{ni,t})}{\sum_{i' \in \mathcal{I}} \exp(\beta_0 K_c \tilde{p}_{i',t}(\mathbf{L}_t) + \boldsymbol{\beta}' \mathbf{X}_{i',t} + \rho H_{ni',t})} \tag{7}$$

Importantly, $\mathbf{L}_t$ and therefore $\tilde{p}_{i,t}(\mathbf{L}_t)$ are both endogenous. We follow Seim (2006) in using a Bayesian Nash solution concept. We define the set of period-t entrants as $\mathcal{C}_t = \{c \quad \text{s.t.} \quad t(c) = t\}$ and $\mathbf{K}_t$ a vector of campus sizes K_c for $c \in \mathcal{C}_t$. Let $\mathbf{Pr}_t^*$ be a $|\mathcal{C}_t| \times |\mathcal{I}|$ matrix, where each entry is the probability that campus $c \in \mathcal{C}_t$ will enter in location i. Then, $\mathbf{Pr}_t^{*'} \mathbf{K}_t$ is a length $|\mathcal{I}|$ vector of the expected data center capacity added in each location i in period t. Firms optimize given shared equilibrium conjecture $\mathbf{Pr}_t^*$:

$$Pr_{ci}^* = \frac{\exp(\beta_0 K_c \tilde{p}_{i,t}(\mathbf{Pr}_t^{*'} \mathbf{K}_t) + \boldsymbol{\beta}' \mathbf{X}_{i,t} + \rho H_{ni,t})}{\sum_{i' \in \mathcal{I}} \exp(\beta_0 K_c \tilde{p}_{i',t}(\mathbf{Pr}_t^{*'} \mathbf{K}_t) + \boldsymbol{\beta}' \mathbf{X}_{i',t} + \rho H_{ni',t})} \tag{8}$$

This equilibrium is Bayesian Nash and symmetric in observables: campuses with identical observable characteristics (size, type, incumbency) share the same equilibrium choice probabilities, and campuses optimize for expected profits given knowledge of all observables and beliefs about competitors' entry behavior.[20]

We also construct a variant of this model that allows different types of data centers to have unique location preferences. Let $\tau(c) \in \mathcal{T} = \{\text{hyperscale}, \text{colocation}, \text{crypto}\}$ denote campus c's exogenous type. We then write utility with type-specific preference parameters:

$$V_{ci} = \beta_0^{\tau(c)} K_c \tilde{p}_{i,t}(\mathbf{L}_t) + \boldsymbol{\beta}^{\tau(c)\prime} \mathbf{X}_{i,t} + \rho^{\tau(c)} H_{ni,t} \tag{9}$$

Again, this yields a conditional-logit choice probability:

$$Pr_{ci} = \frac{\exp(V_{ci})}{\sum_{i' \in \mathcal{I}} \exp(V_{ci'})}$$

with the same equilibrium structure as the pooled variant.

[20] For tractability, we assume that campuses evaluate prices at the expected allocation rather than integrating over allocation realizations. These two values coincide for locally linear price functions.

5.2 Consumer surplus

We assume consumers in each location i are immobile and have inelastic electricity demand. They consume an exogenous amount of electricity $Q_{i,t}$ up to a choke price $\bar{p}$. Like data centers, consumers face electricity prices $p_{i,t}(\mathbf{L}_t)$ that depend on data center load. Consumers also bear climate damages from generation across all locations $i = 1, \ldots, \mathcal{I}$. Although consumers value data center services, we assume that value is location-invariant. Since we also hold constant the total amount of data center capacity, this allows us to exclude data center location choice from the consumer welfare function.

We write total consumer surplus across locations as the difference between consumer surplus from electricity consumption and consumer harm from climate damages created by electricity generation, plus government transfers:

$$CS_t = \sum_{i \in \mathcal{I}} \Big(\underbrace{[\bar{p} - p_{i,t}(\mathbf{L}_t)]\, Q_{i,t}}_{\text{Electricity surplus}} \Big) - \underbrace{\mathcal{D}(\mathbf{L}_t)}_{\text{Climate damages}} + \underbrace{G_t}_{\text{Government transfers}} \tag{10}$$

This specification implies that the harms from data center entry in location i are increasing in the slope of the supply curve in region $r(i)$, total residential electricity demand in region $r(i)$, and the marginal emissions intensity of generation in location i:

$$\frac{\Delta CS_t}{\Delta L_{i,t}} = \sum_{j \in \mathcal{I}} \left(-\frac{\Delta p_{j,t}(\mathbf{L}_t)}{\Delta L_{i,t}} Q_{j,t} \right) - \frac{\Delta \mathcal{D}(\mathbf{L}_t)}{\Delta L_{i,t}} + \frac{\Delta G_t}{\Delta L_{i,t}} \tag{11}$$

where the first term is the total change in residential electricity costs due to data center load.

5.3 Model estimation

To estimate the model presented in Section 5, we define three periods t: 2016–2020, 2021–2025, and 2026–2030. We estimate parameters for the first two periods using observed data center entry from Schlauch et al. (2026) and the last period with recent and planned data center entry from Aterio.[21] Industry reporting suggests that the location choice problem may have changed recently, due to popular opposition to data centers and decreasing power

[21] Aterio assigns data centers a likelihood value, indicating their judgment of the probability that a project will be executed. We restrict attention to high-likelihood projects, which we reconcile against the Schlauch et al. (2026) dataset to avoid double-counting data center capacity.

availability, so we stack the first two periods into a single estimation but separately estimate location choice parameters for the forward-looking period. We assign campuses to a single utility-state based on campus location and utility-state service area shapefiles, dropping campuses with unknown opening date, size, or firm.[22]

We estimate both pooled and type-specific parameters. For the backward-looking estimation, we define types based on campus function: colocation, hyperscale, and crypto.[23] In the forward-looking estimation, function definition becomes complicated: the company associated with each project might be a developer rather than the eventual buyer or end user, and some firms are in the process of transitioning between data center categories due to the AI boom. For this period, we instead define types based on whether or not a campus is flagged as having behind-the-meter power.

We choose $\mathbf{X}_{i,t}$ covariates to capture factors that industry publications highlight as important. As discussed in Section 3.2, the most important factors are power availability and speed-to-power. We proxy for these using a utility's total non-residential electricity sales (measured in 2010, so that it is not contaminated by data center load growth) and metrics capturing the interconnectedness of the transmission network.[24] We also control for utility area, so that the sales measure is essentially sales density. We include covariates that measure an area's history of data center development, another factor that industry insiders describe as predictive of location suitability. We control for the number of campuses built in an area through 2010 as well as whether a given firm has built in that area before. These controls account for unobservables that make certain locations more appealing in general (low construction costs, fiber connections, etc.) and for specific companies (existing relationships with the local utility or municipal government, cross-campus latency). Finally, we control for whether a state has restructured ("deregulated") its electric utilities, a process that introduced more competition into electricity generation. In restructured states, it may be easier for large customers like data centers to participate in wholesale markets and

[22] Since many utility-state service territories overlap, we must remove duplicates. For campuses with multiple utility-state matches, we assign campuses to the utility with the largest non-residential sales in 2016.

[23] We assign the few campuses without a type label, or with a label outside of these three, to the colocation category.

[24] We measure the mean degree of both high-voltage ($\geq$220 kV) and mid-voltage (69–161 kV) transmission networks within a 10 kilometer buffer of a utility-state's service territory. "Degree" captures the number of transmission lines meeting at each substation or junction; a higher value indicates a denser network.

physical power purchase agreements, potentially lowering costs. Analyzing this relationship is outside of the scope of this work.[25] We anchor utility characteristics to a pre-data center growth year where possible.[26] Starting price values are measured from a reference year two years prior to each period start (e.g., 2014 for the 2016–2020 period). We z-score continuous covariates for easier interpretation.

We define price function $p_{i,t}(\mathbf{L}_t)$ based on our dispatch model from Section 4. We take as given supply curve $S_{r,\bar{t}}$ and baseline demand $D_{r,\bar{t}}$, where $\bar{t}$ is the reference year for period t and r is the region to which utility-state i belongs. We then calculate the region-level change in market-clearing price (MCP) from adding load $L_{r,t} = \sum_{i \in r} L_{i,t}$:

$$\Delta MCP_{r,t} = S_{r,\bar{t}}(D_{r,\bar{t}} + L_{r,t}) - S_{r,\bar{t}}(D_{r,\bar{t}})$$

For tractability, we approximate $S_{r,\bar{t}}$ as a single spline function for each region-reference year, rather than allowing merit order to vary month-to-month. Next, since $\Delta MCP_{r,t}$ is region-level but choice is utility-state-level, we regress local MCPs on regional MCPs and scale $\Delta MCP_{r,t}$ by the resulting coefficient, ζ: $\Delta MCP_{i,t} = \zeta \Delta MCP_{r,t}$. Refer to Appendix B for more details. We then add $\Delta MCP_{i,t}$ to the reference year non-residential retail rate in utility-state i to yield entry-adjusted $p_{i,t}$. This assumes that changes in wholesale prices are passed through completely and instantaneously to retail rates.[27]

We estimate model parameters θ using maximum likelihood with log-likelihood function:

$$\mathcal{L}(\theta) = \sum_{c \in \mathcal{C}} \left[V_{ci(c)} - \log \left(\sum_{i' \in \mathcal{I}} e^{V_{ci'}} \right) \right]$$

Note that the static version of this model can equivalently be estimated using PPML with

[25] See Borenstein and Bushnell (2015) for an overview of restructuring to date.

[26] Our transmission network is based on 2022 data, which introduces the possibility that some lines may have been built in response to data center location decisions. However, this endogeneity is limited because recent transmission construction has been limited (only 56 miles of high-voltage transmission were built in 2023) and recent projects typically take 10 or more years to complete, far longer than the development timeline for a data center (Wiser et al. 2024).

[27] Due to rate regulations that allow utilities to fully pass on generation costs, we expect pass-through rates to be 100% in the medium run. Since we calculate future discounted retail prices over a data center's lifetime (assumed to be 10 years, with a monthly discount rate of 0.008), the timing of pass-through ends up being a relatively minor point.

campus fixed effects, as described in Guimarães, Figueirdo and Woodward (2003). We insert an inner loop for each θ, wherein we iterate to find expected $\mathbf{L}_t$ and $\tilde{p}_{i,t}(\cdot)$ in equilibrium. In practice, this inner loop moves price coefficients by about 1% in the realized period, and by around 13% in the future period.

5.4 Results

Our location choice parameter estimates are displayed in Table 1. Across specifications, both pooled and type-specific, we estimate a negative coefficient on the price-size interaction term, consistent with firms preferring cheaper utilities. The one exception is behind-the-meter campuses, which show no significant preference for cheaper electricity. This is intuitive, since campuses with behind-the-meter power do not generally purchase electricity from the grid, and suggests that our negative estimates for other types are not driven by omitted variable bias. Coefficient magnitudes by type line up with observed concentration in low-price areas (Figure C.20b), and imply price elasticities of -2.0 for hyperscale facilities, -0.2 for colocation, and -0.4 for crypto at the median size for each type of campus (164 MW, 18 MW, and 59 MW respectively). The colocation and crypto estimates are in the same range as estimates from Kahn and Mansur (2013) on the electricity price elasticities of electricity-intensive manufacturing industries, while the hyperscale estimate exceeds the largest elasticity they calculate (-1.2, for the primary metal industry).

Behind-the-meter campuses' preferences also diverge from the other campus types' in disfavoring locations with an early history of data center development. However, all types prefer to build in locations where the firm already has some presence. The pooled estimate suggests that firms are 5.3 times more likely to place a campus in a location where they have built previously. Behind-the-meter and crypto campuses are the only types with point estimates indicating a preference for areas with less valuable land, indicative of these types being less dependent on the networking infrastructure (electricity transmission or fiber) important for traditional data center campuses.

We find positive coefficients on both sets of variables designed to proxy for power availability and speed-to-power. 2010 log non-residential electricity sales predict more data center presence for all types, though this preference is attenuated for planned behind-the-meter

campuses. This pattern is consistent with an industry preference for locations with better developed electricity infrastructure. The degree of the high-voltage grid, which captures interconnectedness, is also positively correlated with entry, with the largest coefficient on hyperscale campuses. Hyperscale facilities are the largest type on average, and therefore the most likely to connect directly to very high voltage lines. Mid-voltage degree is positively correlated with entry as well, though the coefficient is significant only for crypto and behind-the-meter facilities.

Table 1: Campus location choice estimates: realized and announced buildouts

	2016–2025				2026–2030	
	Pooled	Hyperscale	Colocation	Crypto	Grid	BTM
β_0: price × size ($K_c\tilde{p}$, std.)	−1.536	−1.731	−1.951	−0.958	−1.442	−0.482
	(0.396)	(0.466)	(0.637)	(0.419)	(0.472)	(0.355)
Log 2010 non-residential sales (std.)	+1.080	+0.810	+1.405	+1.048	+1.067	+0.532
	(0.181)	(0.273)	(0.270)	(0.221)	(0.277)	(0.354)
Log utility-state area (std.)	+0.639	+0.818	+0.595	+0.611	+0.935	+1.862
	(0.280)	(0.466)	(0.303)	(0.507)	(0.366)	(0.703)
Pre-2010 campus count (std.)	+0.125	+0.140	+0.133	+0.075	+0.096	−0.091
	(0.069)	(0.109)	(0.078)	(0.066)	(0.079)	(0.061)
Log land value (std.)	+0.034	+0.318	+0.267	−1.354	+0.030	−0.765
	(0.241)	(0.423)	(0.221)	(0.384)	(0.257)	(0.293)
Deregulated state	+0.309	+0.483	+0.185	+0.348	+0.032	+1.646
	(0.543)	(0.806)	(0.498)	(0.742)	(0.608)	(0.855)
High-voltage (≥220 kV) mean degree (std.)	+1.036	+1.552	+0.894	+0.533	+0.627	+0.328
	(0.341)	(0.225)	(0.444)	(0.463)	(0.382)	(0.492)
Mid-voltage (69–161 kV) mean degree (std.)	+0.368	+0.256	+0.441	+0.750	+0.176	+1.023
	(0.281)	(0.353)	(0.315)	(0.190)	(0.415)	(0.255)
ρ: firm incumbency H	+1.657	+2.018	+0.913	+2.388	+2.168	+2.383
	(0.442)	(0.955)	(0.375)	(0.291)	(0.199)	(0.460)
Campuses (choosers)	417	122	203	92	465	65
of which $H = 1$ at chosen location	99	49	41	9	197	12

Notes: Campus-level conditional logit over the choice set of utility-states in the contiguous U.S.; utility-states missing reference-year prices are dropped. Coefficients represent the change in a campus's location utility per one-standard-deviation increase in the covariate, except "Deregulated state" and firm incumbency H, which are indicators. Standard errors in parentheses are two-way clustered on firm and destination state. Prices are evaluated at the price-feedback equilibrium; behind-the-meter campuses face the baseline price. "Pooled" stacks all campus types; the remaining columns split by type. β_0 is the coefficient on the product of expected lifetime electricity price per MW and campus size K_c in MW, z-scored over the estimation panel. Baseline prices are measured at each cohort's reference year: 2014 and 2019 for the realized periods, 2024 for the announced buildout. Mean degree is the average number of transmission lines meeting at each substation or junction within 10 km of the utility-state, for high-voltage (≥220 kV) and mid-voltage (69–161 kV) lines separately. The realized-buildout specifications also include an unreported indicator for no high-voltage node within 10 km, with one coefficient common to all campus types. $H = 1$ where the campus's firm already operates in the utility-state.

6 Counterfactuals

We use the parameters $\hat{\theta}$ estimated above to model price and emissions effects of data center load under different counterfactual scenarios. In each scenario, we re-solve the inner fixed point at $\hat{\beta}$ to find the equilibrium data center allocation, then trace through the price and emissions effects implied by this data center allocation. We calculate counterfactuals across two time horizons, backward- and forward-looking. The backward-looking scenarios use 2019 region-level power plants and baseline demand, then add 2021–2025 data center entrants. The forward-looking scenarios use 2024 power plants and demand and add 2026–2030 data center entrants. In both, we construct supply curves assuming fuel prices equal to their average values between 2019 and 2025.

We model four different types of data center allocation. First, we model data center placement according to the location preferences we estimate in Section 5.4. We assign campuses to utility-states drawing from our modeled choice probabilities. Second, we model a counterfactual that internalizes the climate damages associated with electricity generation. We add a carbon tax to each power plant's marginal generation cost, based on our estimates of its carbon-intensity and a social cost of carbon (SCC) of \$185/tCO_2 (Rennert et al. 2022). We recalculate merit order based on these new costs, then calculate model-implied location choice probabilities based on the carbon-inclusive prices. Third, we model an \$11/MWh surcharge on data center electricity consumption in Virginia, a policy enacted by the state in mid-2026. Finally, we model a set of single-state data center moratoria. We remove one state at a time from the choice set, then assign campuses according to their modeled choice probabilities. This last counterfactual is inspired by the many states and localities proposing pauses or outright bans on data center development (National Conference of State Legislatures 2026).

In each scenario, we calculate the change in consumer surplus. As in our model from Section 5.2, consumer surplus is surplus from electricity consumption, net of climate damages plus government transfers. Data center entry decreases electricity consumption surplus by increasing local residential rates and increases climate damages by inducing additional electricity generation, which emits CO_2. In the carbon tax scenario, we assume that tax revenues are rebated lump-sum to consumers so that changes in climate damages from

data centers are perfectly offset by additional government transfers.

Across scenarios, we hold fixed the number of campus entrants and their characteristics (including size and type). We also fix location characteristics *except* for price. As in our model estimation, we iterate to find prices and choice probabilities in equilibrium. We cap regional data center entry so that it cannot exceed available generation. In the scenarios composed of sets of random draws, we implement this cap by rejecting campus placements that put a region over its thermal generation capacity in more than 0.1% of hours. If data center capacity does push a region past thermal generation capacity for some share ($< 0.1\%$) of hours, we assume that data center load is curtailed but that data center location choices are made assuming that prices ramp up to a scarcity value of $9,000 per MWh (capturing the value to the firm of lost load). We exclude behind-the-meter campuses from our counterfactuals, since their location choices have no effect on consumer welfare in our model.

6.1 Backward-looking reallocation

Price-based policies induce some data center relocation, but this relocation does not dramatically shift price or emissions effects. A carbon tax reduces total climate damages by about $100 billion per year (Table 2), but the effect is largely unchanged by whether we allow data centers to relocate: the benefits are coming through a reordering of electricity dispatch so that cleaner plants are dispatched first. 17% of data center load reallocates in response to the tax, but this only cuts incremental carbon emissions by 8% and total climate damages by less than 1%. The Virginia electricity surcharge has negligible effects on topline prices or emissions. This $11 per MWh surcharge, which amounts to an 11% electricity price increase for data centers, causes 820 MW to relocate, or about 17% of the data center capacity that would have entered Virginia in this period without the surcharge. This saves Virginia's residential electricity consumers about $4 million but increases costs in other states, leading to a net increase in residential electricity costs of about $5 million. The surcharge yields $278 million in tax revenues from new data centers, plus revenues from the data centers that chose Virginia prior to 2021, which we assume cannot move.

Table 2: Counterfactual data center allocations (2021–2025 buildout)

		Carbon tax		Virginia surcharge	
Outcome	Baseline	Locations fixed	Firms relocate	Locations fixed	Firms relocate
Price effect ($/MWh)	0.97	5.71	5.91	0.97	0.97
Incremental CO_2 (Mt)	106.44	117.28	107.36	106.44	106.76
Total climate damages ($B/yr)	$333.6	$233.1	$231.3	$333.6	$333.7
Δ residential electricity costs ($B/yr)	$1.39	$8.19	$8.47	$1.39	$1.40
Δ climate damages ($B/yr)	$19.69	$21.70	$19.86	$19.69	$19.75
Δ government transfers ($B/yr)	$0.00	$21.70	$19.86	$0.33	$0.28
Δ consumer surplus ($B/yr)	-$21.09	-$8.19	-$8.47	-$20.75	-$20.87
Capacity in Virginia (GW)	4.93	4.93	4.40	4.93	4.12
Capacity relocated vs. baseline (GW)	0.00	0.00	5.27	0.00	0.82

Notes: Data center campuses built 2021–2025 reallocated across utility-states using the type-specific preference estimates in Table 1. *Baseline* and *firms relocate* columns report the median of each row over 2,000 campus-level draws from the equilibrium choice probabilities, subject to a generation capacity screen; *locations fixed* columns re-evaluate the *Baseline* draws under the policy. Because rows are medians taken separately, the consumer surplus row need not equal the sum of its components. *Carbon tax*: $185/t$CO_2$ added to generator marginal costs, with dispatch re-sorted and revenue on data center–induced generation rebated to consumers. *Virginia surcharge*: $11/MWh on data center consumption in Virginia, paid by data centers only, with revenue rebated to consumers. Price effect is the residential-load-weighted retail price change from data center load. Incremental CO_2 is total-grid emissions with minus without data center load; total climate damages value total-grid emissions at the SCC. Residential electricity costs are positive for an increase. Consumer surplus $=$ $-$costs $-$ climate damages $+$ transfers. Capacity relocated is measured on the expected allocation. Location choice is evaluated with load beyond thermal capacity priced on a scarcity ramp rising linearly to $9,000/MWh over a band equal to 5% of thermal capacity; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC $185/t (Rennert et al. 2022); data center load $=$ nameplate $\times$ 0.70; 2019 dispatch.

In our moratorium counterfactual results (Table 3), we find that states that ban data centers can either impose significant costs on their neighbors or induce large cost savings, depending on the grid characteristics of the state enacting the ban. A ban in California, which has a grid with a steep supply curve but clean marginal generation, would have increased total climate damages by $153 million per year, while increasing total (in-state plus out-of-state) residential electricity costs by just $6 million. These large climate damages make California's ban the most costly per MW of displaced capacity. A ban in Virginia, which is part of a region (PJM) with a relatively flat supply curve, would have decreased annual in-state electricity costs by $25 million but increased out-of-state residential costs by $80 million. This is despite the ban's creating cost-savings in the other states in PJM (Figure 8). A Virginia ban would also have increased total climate damages by $349 million because

displaced data centers move to regions with dirtier grids, such as the Southeast and MISO-North. In contrast, a ban in Michigan would have net benefits because Michigan, as part of MISO-North, has a coal-intensive grid. Removing Michigan from data centers' choice set would have caused campuses to move to regions with less carbon-intensive generation, decreasing climate damages by $156 million (0.8 megatons of CO_2) while increasing total residential electricity costs by just $9 million. The average proposed ban would have increased net residential electricity costs by $1 million and increased climate damages by $58 million.

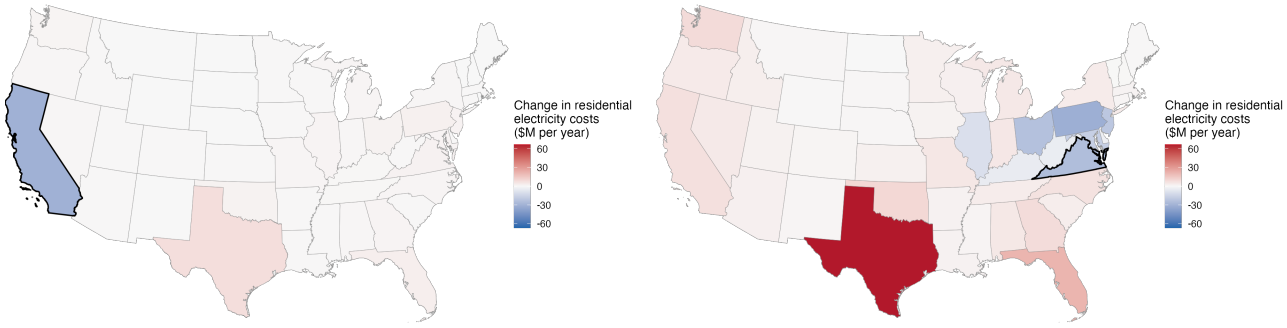
These results highlight the importance of local grid characteristics. If bans force data centers to relocate to regions with steeper supply curves and dirtier generation, the aggregate effect will be an increase in emissions and climate damages. In contrast, if bans push data centers toward regions with flatter supply curves and cleaner generation, the aggregate effect can be a decrease in emissions and climate damages. Understanding firm location choice is therefore key to predicting the welfare tradeoffs involved in bans and other policies.

Table 3: State data center moratoria ranked by Δ consumer surplus (2021–2025 buildout)

		Δ residential electricity costs ($M)		Total relocation effects	
	Capacity displaced (GW)	In-state	Out-of-state	Climate damages ($M)	Consumer surplus ($M)
Texas*	5.85	-281.4	236.1	+$1014.6	-$969.3
Virginia*	4.96	-25.0	80.2	+$348.8	-$404.0
California	0.75	-30.8	36.9	+$153.1	-$159.2
Ohio	1.40	-5.9	18.6	+$82.0	-$94.7
New York*	0.51	-22.7	25.5	+$85.3	-$88.1
⋮					
Iowa	0.45	-1.1	8.1	-$129.8	+$122.7
Oklahoma*	0.80	-23.1	9.7	-$122.4	+$135.8
Georgia*	1.75	-13.6	10.8	-$140.6	+$143.4
Indiana	0.54	-4.0	12.8	-$152.2	+$143.4
Michigan*	0.58	-4.7	14.0	-$156.2	+$146.9
Avg, proposed ban*	1.04	-23.8	24.9	+$58.0	-$59.1
Avg, no ban proposed	0.44	-3.6	6.2	-$14.0	+$11.3

Notes: Per-state moratorium on the modeled 2021–2025 buildout, using the type-specific preference estimates in Table 1: the state is dropped from every campus's choice set and its capacity redistributed at the price-feedback equilibrium. Capacity displaced is the baseline modeled capacity in the banned state. Residential electricity costs are positive for an increase, so an in-state saving is negative. Climate damages are the signed change in total-grid CO_2 valued at the SCC. Consumer surplus $= -$in-state costs $-$ out-of-state costs $-$ climate damages. The five largest consumer-surplus decreases and increases are shown; $*$ marks the 16 states with a proposed or enacted moratorium as of September 2026 (NCSL tracker; Office of the Texas Governor 2026). Group rows are unweighted averages within each group (16 proposed-ban states, 32 other states in the contiguous U.S.). Moratorium allocations are expected allocations and are not subject to the generation capacity screen. Location choice is evaluated with load beyond thermal capacity priced on a scarcity ramp rising linearly to $9,000/MWh over a band equal to 5% of thermal capacity; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC $185/t (Rennert et al. 2022); data center load $=$ nameplate $\times$ 0.70; 2019 dispatch.

Figure 8: Change in residential electricity costs by state under a California or Virginia data center moratorium (2021–2025 buildout)



(a) California moratorium **(b)** Virginia moratorium

Notes: Change in annual residential electricity costs ($M per year) in each state in the contiguous U.S. when the outlined state bans data centers, on the modeled 2021–2025 buildout: the banned state is dropped from every campus's choice set and its capacity redistributed at the price-feedback equilibrium. A state's value is the sum, over its utility-states, of residential consumption times the retail price change the reallocation induces. Warm shading is a cost increase, cool a saving; the two panels share one color scale. Data center load = nameplate × 0.70; 2019 dispatch.

6.2 Forward-looking reallocation

To project forward how future data center growth will impact prices and emissions, we model the addition of data center campuses that are planned to open between 2026 and 2030 to the 2024 grid. We immediately find that new planned capacity (146 GW by 2030) is large enough to exhaust the available generation capacity on the existing grid. These constraints bind differentially across regions. While the Southeast has enough headroom to accommodate all new campuses planned for that region, the nameplate capacity of ERCOT's planned data centers is four times larger than the headroom available on its grid in 2024 (Table D.16). To accommodate the entire planned buildout, the grid would need to add over 20 GW of new generation, more than New England's entire existing fossil fleet.

To make this counterfactual tractable, we assume that only 50% of the planned capacity will be realized by 2030. This assumption aligns our total capacity prediction with projections from S&P Global (Hering and Dlin 2025), and is compatible with quoted capacity numbers in our data being projections of eventual buildout rather than total buildout over the next few years.[28] In some specifications, we also add power plants planned to come

[28]Our data indicate 61 GW of data center capacity built in or before 2025. In our data, 50% of planned,

online in the next few years to the 2024 merit order, increasing the total amount of generation capacity available for data centers to use. We assume that planned fossil plants have the average marginal costs and capacity factors of recently built plants of the same technology. For planned renewable plants, we impose region-technology-hour capacity factors calculated from OEDI and EIA 930 data.[29]

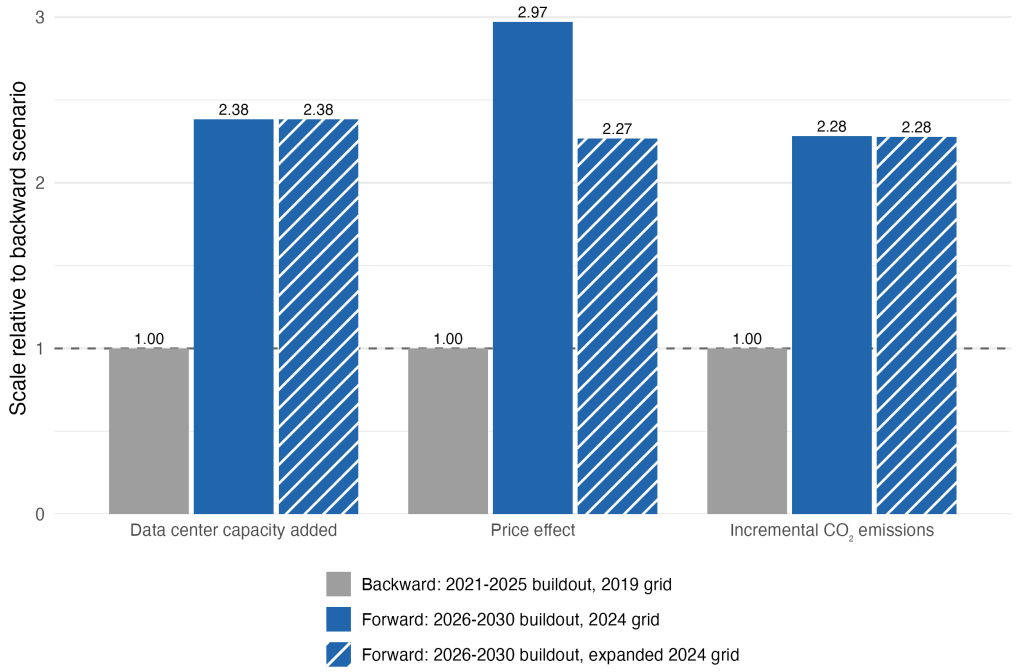
Even at 50% realization, the buildout planned for 2026–2030 is still 2.4 times larger than the buildout that occurred 2021–2025 (Figure 9). We find that price effects scale superlinearly with buildout. The 2026–2030 buildout increases average prices by $2.87 per MWh, three times larger than the price effect of the 2021–2025 data center additions. Although emissions do scale roughly linearly with data center capacity, the magnitude of these emissions is still huge. On the 2024 grid, the climate damage-minimizing allocation leads to 219 megatons of CO_2 emissions per year, equivalent to about $40 billion in climate damages per year. Allowing the grid to grow through the addition of new power plants reduces the overall price and emissions impacts of the data center buildout (Figure 9). Under the augmented version of the 2024 grid, data center price effects are 24% smaller than price effects estimated using the baseline 2024 grid. Incremental CO_2 emissions are essentially unchanged by grid expansion, though the grid's average CO_2 intensity drops by about 10%.

Because the amount and type of new generation added varies across regions, grid growth changes the relative costs of placing data centers in different regions (Figure C.25). Generation additions cause marginal emissions rates in MISO-South and PJM to drop 11% and 10%. Marginal emissions intensity actually increases in some regions (including a 10% increase in ERCOT) because renewables growth shifts net demand down into the coal region of the supply curve in more hours. Marginal supply curve slopes also become flatter, especially in PJM, ERCOT, and MISO-South. The costs of individual state moratoria change due to grid growth (Tables D.3 and D.4). On average, proposed moratoria become less socially costly on an augmented grid (-$64 million) than on the 2024 base grid (-$129 million). This shift is largely driven by a decrease in climate damages from a moratorium in Texas.

grid-connected data center capacity (2026–2030) amounts to 73 GW, summing to 134 GW of total data center capacity in 2030. S&P's corresponding estimates (Hering and Dlin 2025) are 62 GW in 2025 and a total of 134 GW of data center capacity in 2030.

[29]Within our thermal-only dispatch model, adding renewable capacity is achieved by decreasing observed fossil demand by the amount of predicted renewables generation in a given hour. We include planned wind and solar plants from EIA 860.

Figure 9: Counterfactual outcomes relative to the 2021–2025 buildout



Notes: Capacity, price effect, and incremental emissions under each counterfactual scenario, divided by the same outcome in the backward scenario (dashed line at 1). Backward: data center campuses built 2021–2025, allocated across utility-states using the pooled 2016–2025 preference estimates in Table 1, on the 2019 merit order. Forward: the announced 2026–2030 grid buildout at 50% realization, allocated using the Grid (2026–2030) estimates in Table 1, on the 2024 merit order. The expanded grid adds all EIA-860 proposed thermal units to the 2024 merit order and nets planned wind and solar from thermal demand. Price effect is the residential-load-weighted retail price change. Incremental CO_2 is total-grid emissions with minus without data center load. Price and emissions outcomes correspond to the median over 2,000 campus-level draws. In the backward scenario, nameplate data center capacity is 30.7 GW, the price effect is \$0.97/MWh, and incremental emissions are 107 megatons. Data center load = nameplate $\times$ 0.70.

7 Conclusion

In this paper, we develop a theoretical framework for understanding how location choice shapes the costs of data center development. We build a dispatch model to show that where data centers locate matters for their overall effects on electricity costs and emissions. Depending on the chosen region, a new data center with power draw of 1 GW could raise electricity prices by over \$3 or closer to \$0.30, and could create either 3.7 megatons of annual carbon emissions or as much as 6.5 megatons. We find that, so far, data centers have only modestly increased electricity generation costs, in part because they have tended towards locations with flatter electricity supply curves. In contrast, data centers' electricity

demand has already led to large climate damages, totaling nearly $40 billion in 2025.

In order to estimate how policy proposals could change data centers' cost impacts, we build and estimate a model of data center location choice, which embeds our dispatch model so that electricity prices can move endogenously with data center load. We find that, by taking unilateral action to block local data center development, state policymakers can reduce data center-induced electricity costs for their constituents. However, depending on the grid characteristics of the implementing state, these bans could either increase or decrease data centers' total impacts on electricity costs and carbon emissions.

As spare capacity on the grid runs out, data centers may have to locate in places with steeper supply curves or simply build their own generation. Our model suggests that as data center load continues to grow, generation costs will increase superlinearly with additional load. Furthermore, regardless of location choice, we project that the climate costs of new data centers will be on the order of $50 billion per year by 2030, even if only half of the announced buildout is realized. Building new generation flattens electricity supply curves, reducing data centers' impacts on prices, but does not necessarily reduce the carbon intensity of marginal generation. In future work, we hope to explore how policy combining new clean generation with data center development could leverage data center developers' willingness to pay for power to advance decarbonization goals.

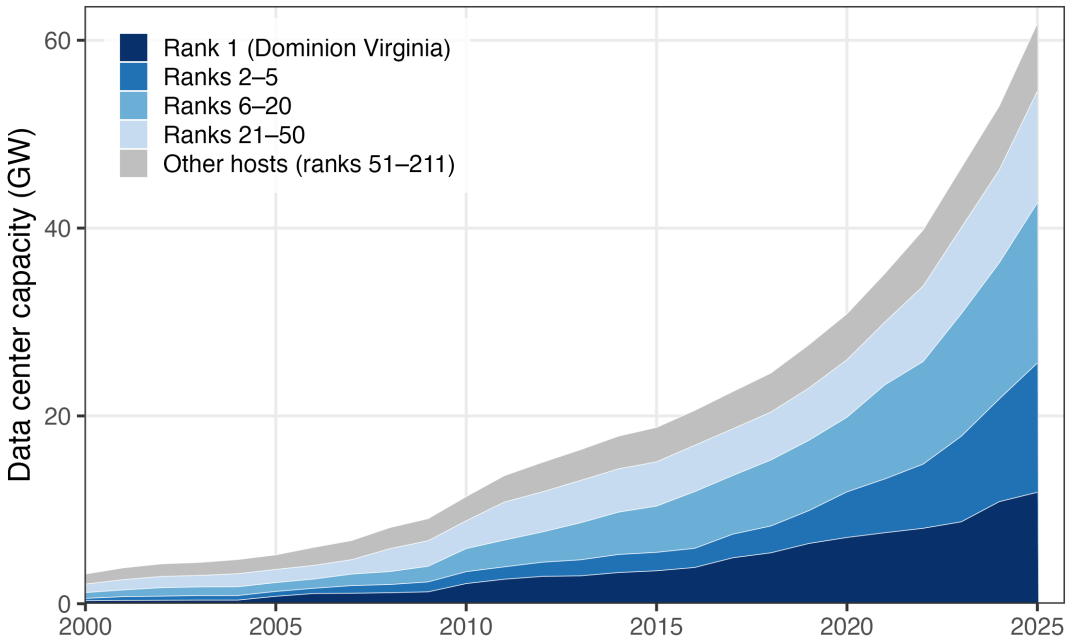
8 Appendix

A Estimating reduced-form effects of data centers on prices

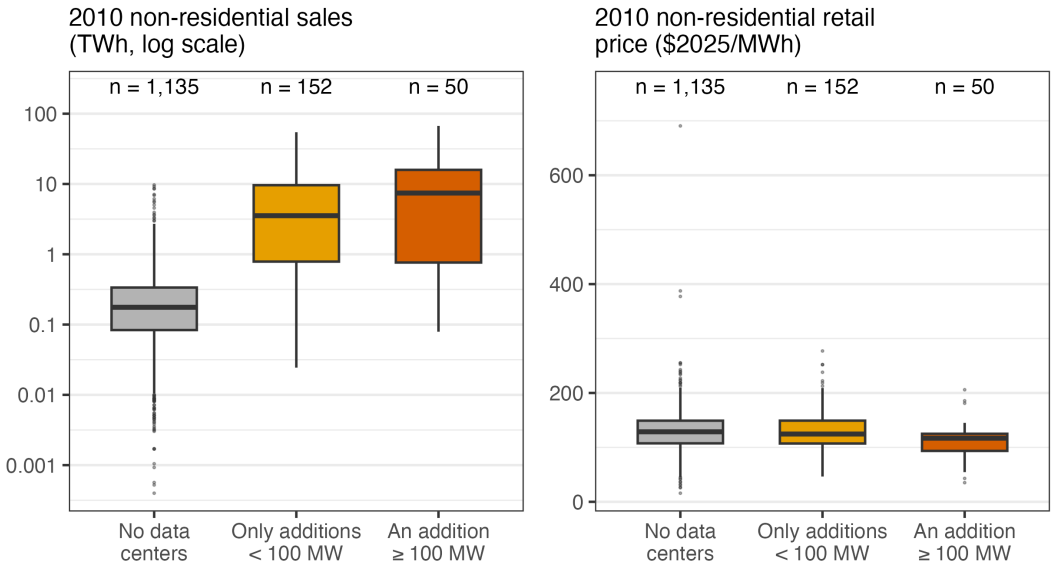
Our dispatch model describes how data center growth affects wholesale electricity prices, but the electricity prices customers pay are not solely determined by wholesale markets. As outlined in Section 3.1, data centers can also affect electricity prices through other channels, including capacity markets and infrastructure costs. These other channels are challenging to model.

It is tempting to use realized price data and causal identification strategies to estimate the net effects of data centers on retail prices, but there are several reasons why these techniques may not be well-suited to this setting. First, electricity prices are convex in load, so backward-looking estimates of price effects will tend to underpredict future effects. This is true for wholesale electricity prices (as shown in Figure 1), but also for infrastructure costs, where additional electricity demand can actually decrease the costs paid by other customers up until the point when new infrastructure must be built (Borenstein 2025). Second, the locations where data centers are built are not easily comparable to locations without data centers. Data centers are heavily concentrated in just a few utilities (Figure A.1a), and their preferred utilities are both larger and cheaper than utilities with less or no data center capacity (Figure A.1b). As a result, we have few treated units and it is difficult to argue that non-treated units form a reasonable control group.

Figure A.1: Concentration of data center capacity and characteristics of host utility-states



(a) Cumulative data center capacity by utility-state rank



(b) Utility-state characteristics by data center exposure

Notes: Panel (a): cumulative data center capacity (GW), 2000–2025, for groups of utility-states ranked by their 2025 capacity. A campus matched to several utility-states is split among them in proportion to 2016 non-residential sales. Capacity is cumulated by reported opening year; campuses without an opening year are excluded, and capacity opened before 2000 is in the 2000 stock. 1,149 of 1,360 utility-states host no data center capacity. Panel (b): 2010 non-residential electricity sales (TWh, log scale) and average non-residential retail price ($2025/MWh) for three groups of utility-states: no data center capacity in 2000–2024; some capacity but no annual addition of 100 MW or more; and at least one annual addition of 100 MW or more. Boxes span the interquartile range, with the median marked; whiskers extend to 1.5 times the interquartile range. Numbers above the boxes are the utility-states in each group with data. Excludes two utility-states whose first annual addition of 100 MW or more was before 2000. Campus data from Schlauch et al. (2026); sales and prices from EIA-861.

We explore these issues through two different identification strategies. First, we use the timing of campus openings to estimate local projections difference-in-differences, with which we attempt to trace out the dynamic effects of data center construction on prices (Section A.1). We next construct a firm Bartik instrument for data center expansion that exploits persistence in firms' location preferences (Section A.2).

Because wholesale electricity markets are regional, data centers can move electricity prices in all the utilities that share a region with the utility actually hosting the facility. We consider two types of analysis: within-region and across-region. In the first, we compare utilities to same-region neighbors. Our dispatch model predicts a null effect: a data center recipient's generation costs should evolve in the same way as its neighbors'. If there is within-region congestion, or if data centers increase non-generation costs (infrastructure costs, capacity market costs, etc.) for other customers, then we will recover a positive price effect of data center entry. If data centers decrease non-generation costs for other customers, then we will recover a negative within-region effect. In contrast, the across-region analysis compares treated utilities to out-of-region controls, eliminating wholesale market spillovers. This analysis captures the average net effect of both wholesale price effects and other price mechanisms.

A.1 Local projections difference-in-differences

We use an event-stacked implementation of the LP-DiD approach from Dube et al. (2025) with a continuous treatment dose. The event-stacked approach allows us to define distinct clean control groups for each event, making possible the within-region vs. across-region distinction. We exclude utilities that have never had any data center additions from the set of possible controls, since these tend to be far smaller (in both area and sales) than the utilities in either our preferred control group or our treated group (Table D.6). We identify the effect of data center growth on electricity prices by exploiting the precise timing of data center opening, which depends on factors such as local regulations and equipment delays.

We define event $e = (i, t_0)$ to be utility i's first $\geq$ 100 MW addition. We include only events that occurred in or after 2001, leaving us with 44 events. We define "clean controls" C_e as utilities either in the same region as i (for the within-region analysis) or outside of i's

region (for the across-region analysis) that have non-zero data center capacity by period-end but do not receive a data center addition ≥ 100 MW through t_0+5. This category includes both never-treated units, which only ever experience small data center additions, and not-yet-treated units.[30]

We stack i with these clean controls and estimate long differences relative to $h = -1$:

$$y_{u,t_0+h} - y_{u,t_0-1} = \mu_e^h + \beta^h \, \text{dose}_e \cdot \mathbf{1}\{u = i\} + \varepsilon_{e,u,h}, \qquad u \in \{i\} \cup C_e \tag{A.1}$$

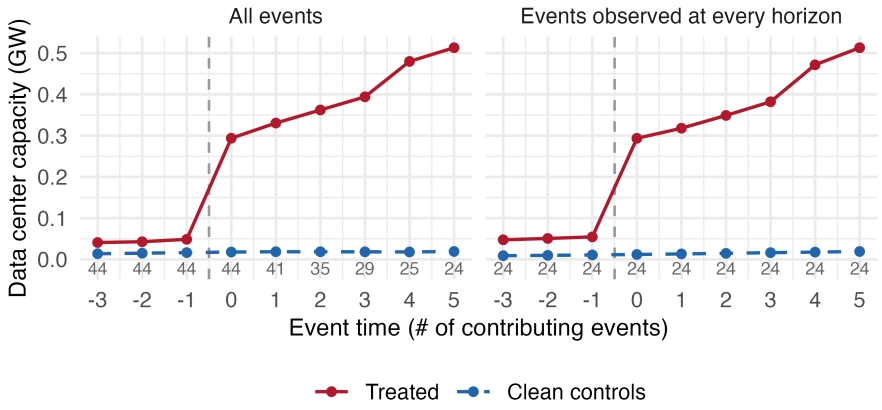
where μ_e^h is an event fixed effect and dose_e is the capacity added in the crossing year t_0. We cluster errors at the state level. β^h is then the effect per GW of the triggering addition. Because our data indicate selection even with our filter on control units (Figure A.1b), we estimate results with and without entropy balancing (Hainmueller 2012). We balance unit-years on deregulated market status as well as non-residential price and log(1 + non-residential sales) trends in t_{-3} through t_{-1}.

Looking at the raw data underlying the LP-DiD design, we find that control and treated utilities differ substantially. In Figure A.2, we plot mean data center capacity, non-residential electricity sales, and residential electricity prices for treated and control groups, stacked by event time. We see that treated units have more non-residential sales and lower electricity prices than control units, even though we cut the smallest utilities (which never receive any data center capacity) from the control set. We also notice that event composition matters: in the price series, though 43 utilities are treated, only 23 are observed for the full post-period because many utilities received their first large data center only in the past few years. Event composition drives a large price increase after event year two, and an implausibly large increase in non-residential sales by $t + 5$.[31] The full-horizon price series is not vulnerable to composition effects, but rests on just 23 events occurring in utilities that are highly unrepresentative of the full set of utilities.

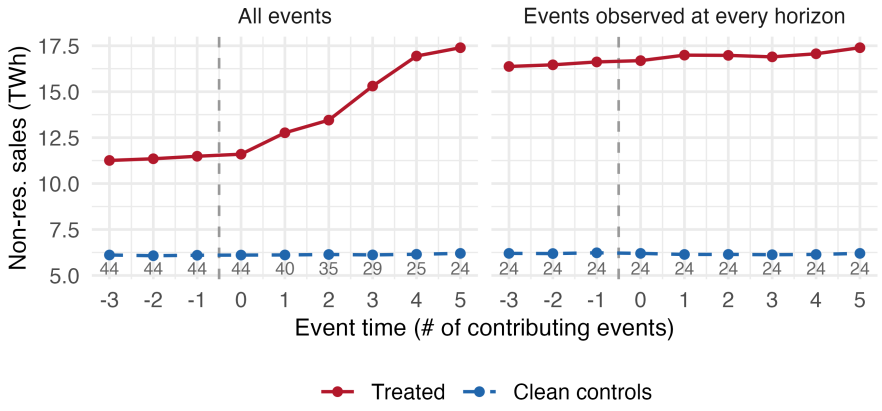
[30] The median never-treated unit has just 14 MW of total data center capacity by period end, small enough that we would not expect significant price effects.

[31] The average treated utility sees the addition of about 0.5 GW of data center capacity by $t + 5$. At 70% utilization, this would correspond to about 3 TWh in additional non-residential electricity demand. The all-event non-residential sales series implies the addition of 6 TWh by $t + 5$, while the full-horizon series indicates +0.8 TWh by $t + 5$.

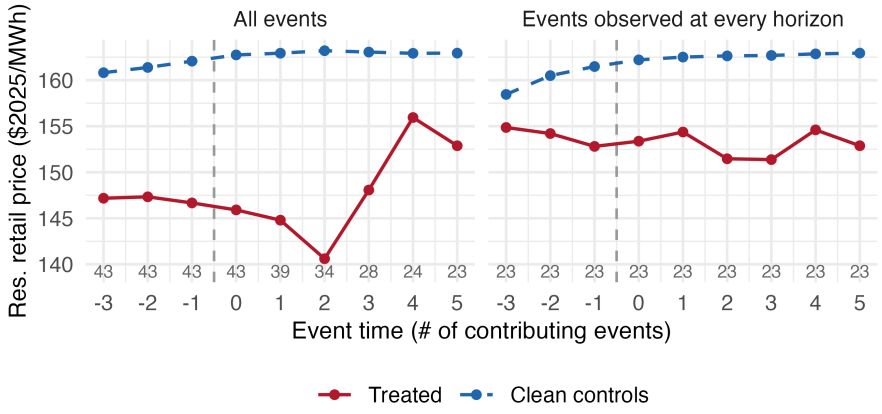
Figure A.2: Treated and control utility-states by event time, across-region controls



(a) Data center capacity (GW)



(b) Non-residential electricity sales (TWh)



(c) Residential retail electricity price ($2025/MWh)

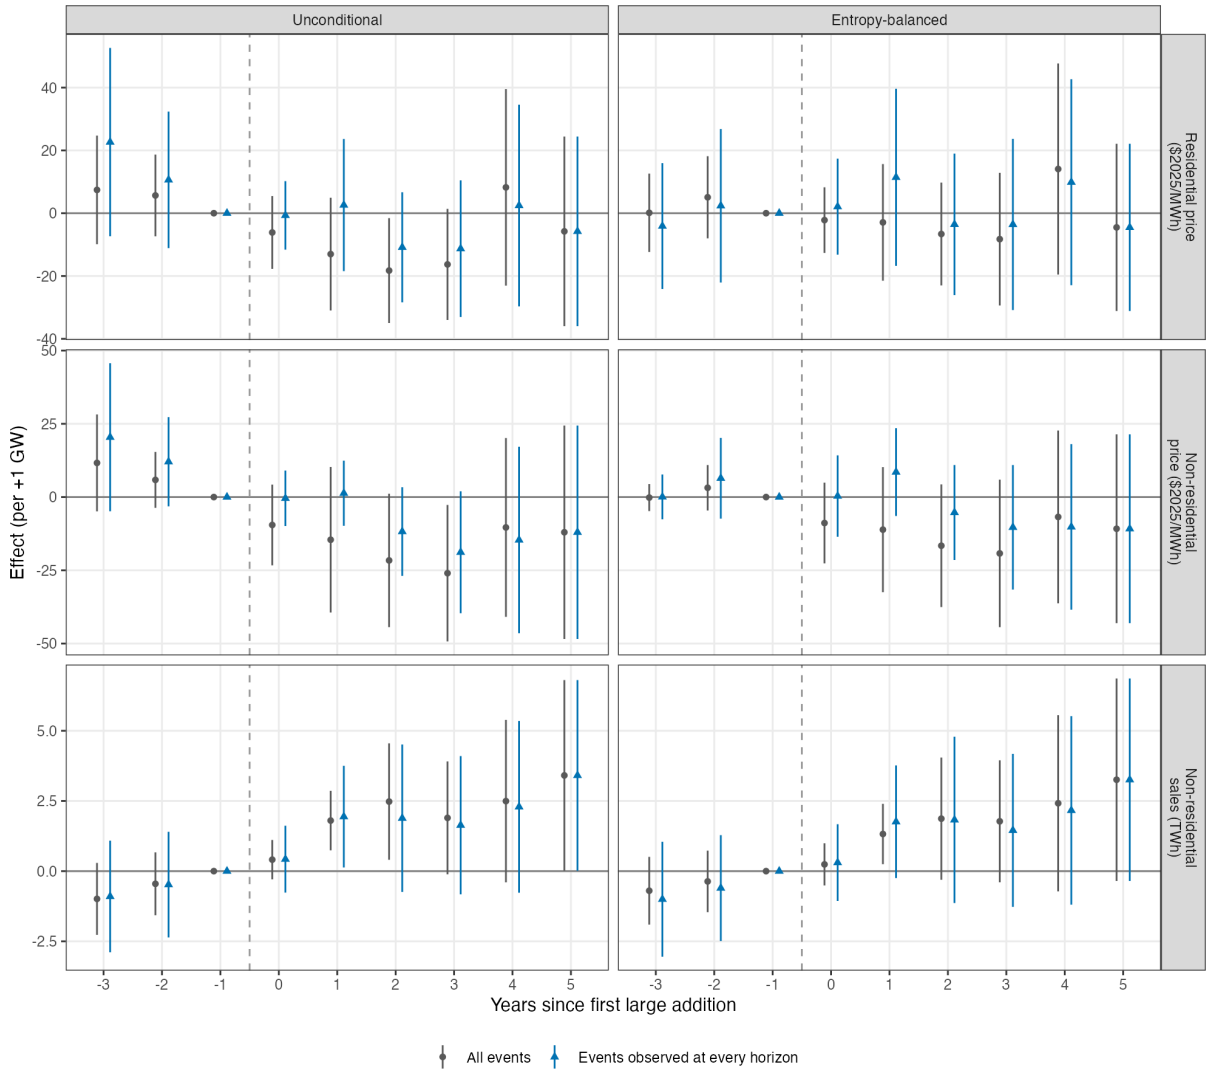
Notes: Means by event time on a utility-state annual panel (2000–2024), with no regression or weighting. An event is the first year a utility-state adds 100 MW of data center capacity or more. Each treated utility-state is paired with its clean controls: utility-states outside its dispatch region that have non-zero data center capacity by 2024 but no $\geq$100 MW addition through t_0+5. Left panels use all events; right panels use only events observed at every horizon from t_0-3 to t_0+5. Numbers above the horizontal axis are the events contributing at each horizon.

Across-region results are in Figure A.3. Both unconditional and entropy-balanced estimates yield almost entirely insignificant estimates of the effect of data center entry on residential and non-residential electricity prices. Most point estimates are negative, but the bounds are extremely wide– $30 or more, relative to retail price levels that are generally between $100 and $200. Our dispatch model predicts wholesale price increases of just $0.33–0.41/MWh per GW. The non-residential sales results indicate a positive effect of data center entry, though again the confidence intervals are very wide and most include zero. A 1 GW data center that is operating full-time at 70% capacity should consume 6.1 TWh per year, which is larger than all point estimates but within the confidence intervals of the year five estimates.

Within-region results (Figure C.19) are similarly noisy and generally share the same signs as the across-region results. Our dispatch model predicts that within-region estimates should be $0.33–0.41/MWh per GW below the across-region values. Our estimation is not precise enough to detect this level of difference.

As further evidence of our results' fragility, we present the results of leave-one-out testing for the across-region specification (Table A.1). With only about 20 events in the analysis, the estimates are sensitive to the exclusion of single events. Only the non-residential sales result keeps its sign through all specifications; the price point estimates vary dramatically in magnitude and sometimes flip signs.

Figure A.3: Local projections DiD: retail price and load effects of data center entry, across-region controls



Notes: Event-stacked local projections difference-in-differences on a utility-state annual panel (2000–2024), with event fixed effects. An event is the first year a utility-state's annual data center capacity addition reaches 100 MW, restricted to events in or after 2001; the dose is that crossing-year addition, rescaled so estimates are per $+1$ GW. Each treated utility-state is stacked with its own clean controls: utility-states *outside* its dispatch region that have non-zero data center capacity by 2024 but no $\geq$100 MW addition through $t_0{+}5$. Rows are the residential retail price and the non-residential retail price ($2025/MWh) and non-residential sales (TWh); horizons -3 to 5, with $h = -1$ the reference. Unconditional estimates weight events by $1/\text{dose}$; entropy-balanced estimates additionally weight controls to match treated units on deregulated-market status and on the changes in non-residential price and log(1 + non-residential sales) from $t_0{-}3$ to $t_0{-}1$, and drop events whose controls cannot be balanced. Grey circles use all events (43 for the residential price, 44 for the non-residential outcomes); blue triangles use only events whose treated utility-state is observed at every horizon from $t_0{-}3$ to $t_0{+}5$ (23 and 24). Vertical lines are 95% confidence intervals from state-clustered standard errors.

Table A.1: Leave-one-event-out sensitivity of LP-DiD estimates, across-region controls

Outcome	Series	Estimate	Range, one event dropped	Sign flips	Most influential event
Residential price	Unconditional	-5.1	[-10.0, 2.6]	1 of 23	Oncor (TX)
	Entropy-balanced	1.6	[-6.4, 7.5]	3 of 21	PG&E (CA)
Non-residential price	Unconditional	-7.4	[-11.4, -2.6]	0 of 24	Northern Wasco (OR)
	Entropy-balanced	-1.7	[-7.1, 2.2]	3 of 22	PG&E (CA)
Non-residential sales	Unconditional	1.47	[0.75, 1.67]	0 of 24	Oncor (TX)
	Entropy-balanced	1.34	[0.58, 1.54]	0 of 22	Oncor (TX)

Notes: Leave-one-event-out re-estimates of the across-region event-stacked LP-DiD (Figure A.3) on events whose treated utility-state is observed at every horizon from t_0-3 to t_0+5. Estimate is the average effect per +1 GW over $h = 0$–3 ($2025/MWh for prices, TWh for sales). Range is the minimum and maximum of that average across re-estimates that each drop one event. Sign flips counts the drops that reverse the sign of the average, out of the events in the estimate. Most influential event is the drop that moves the average the most. Unconditional estimates weight events by 1/dose; entropy-balanced estimates additionally weight controls to match treated units on pre-event trends and deregulated-market status, and drop events whose controls cannot be balanced.

A.2 Bartik instrument

Several large data center companies show a preference for developing more data centers in utilities where they already have a presence (Figure C.18). We make use of this pattern to construct a firm Bartik instrument (Bartik 1991; Goldsmith-Pinkham, Sorkin and Swift 2020).

For utility-state i in region $r(i)$, we define the instrument as:

$$B_i^{\text{firm}} = \sum_f \underbrace{s_{fi}^{\text{pre}}}_{\text{Firm } f\text{'s share}} \times \underbrace{\widetilde{G}_{f,-i}^{2019\text{–}24}}_{\text{Firm } f\text{'s shift}}$$

where s_{fi}^{pre} is firm f's share of pre-2015 MW located in utility-state i and $\widetilde{G}_{f,-i}$ is the total MW of data center capacity built by firm f in the period 2019–2024 in all locations except i. We focus on the six largest firms by total MW of data center capacity built 2019–2024 (Amazon, Microsoft, Google, Meta, Digital Realty, and QTS) that also have pre-2015 presence in at least three utility-states. These firms account for 37% of pre-2015 MW and 42% of MW built 2019–2024 (Table D.9).

The first stage of this instrument is not strong enough to permit across-region analysis, since data center buildout is concentrated in a few regions (around half of data center MWs in this period are built in ERCOT and PJM). However, it does have enough predictive power to allow for weak-IV-robust within-region inference. We define Δy_i to be the 2018–

2024 change in y_i (residential price, non-residential price, or non-residential sales), M_i to be 2019–2024 data center additions in MW, and a vector of controls X_i that includes 2016 non-residential price, 2016 non-residential sales, and land values. We estimate:

$$\text{OLS:} \quad \Delta y_i = \beta\, M_i + \gamma' X_i + \mu_{r(i)} + \varepsilon_i,$$
$$\text{IV:} \quad \Delta y_i = \beta\, \widehat{M}_i + \gamma' X_i \; + \; \mu_{r(i)} + \varepsilon_i, \quad \widehat{M}_i = \alpha\, B_i^{\text{firm}} + \delta' X_i \; (+\, \mu_{r(i)}) + u_i,$$

with region FE $\mu_{r(i)}$ and standard errors clustered on state. Because of a low first-stage F, we use Anderson-Rubin confidence sets for the IV estimates (Anderson and Rubin 1949).

The results are presented in Table A.2. IV confidence intervals are wide, but suggest a negative effect of data center capacity on non-residential retail prices. The estimated effect on residential retail prices is negative but insignificant. Again, we can compare the sales result to the expected impact from a 1 GW data center operating full-time at 70% capacity: the point estimate of 3.13 TWh is smaller than the predicted 6.1 TWh per year, though the confidence interval almost includes this larger number.

Table A.2: Firm Bartik estimates: data center capacity, retail prices, and load

	OLS	IV
Res. price ($/MWh)	+1.48	-2.81
	[-3.64, 6.60]	[-15.01, 0.31]
Nonres. price ($/MWh)	-3.14	-2.40
	[-5.52, -0.75]	[-9.35, 0.64]
Nonres. sales (TWh)	+4.18	+3.13
	[1.66, 6.71]	[2.66, 5.46]
First-stage F (state-clust.)		7.7
Observations	1,253	1,253

Notes: OLS and firm Bartik instrumental-variables estimates of the effect of a +1 GW own data center capacity addition on the 2018–2024 change in each outcome: residential retail price, non-residential retail price ($2025/MWh), and non-residential sales (TWh). Observations are utility-states in the contiguous U.S.; specifications include dispatch-region fixed effects and controls (2016 non-residential retail price, log 2016 non-residential sales, and log land value), with standard errors clustered on state. OLS brackets are conventional 95% confidence intervals; IV brackets are Anderson and Rubin (1949) 95% confidence sets, robust to weak identification. The first-stage F is state-clustered.

When we assess which firms and utilities are driving this result, we find that the instrument relies entirely on just a handful of utilities. 29 utilities have nonzero instrument values,

and one of these (Dominion Virginia, home to 20% of all data center capacity built to date) is responsible for 80% of the instrument's first stage covariance. We can also decompose the Bartik 2SLS estimates into per-firm just-identified estimates and calculate Rotemberg weights to determine which firms are driving our results (Table D.11), following Goldsmith-Pinkham, Sorkin and Swift (2020). We find that Amazon carries by far the largest weight, followed by Microsoft, QTS, then Digital Realty. Prior to 2015, Amazon had data centers in just three utility-states: Dominion Virginia, plus two utilities in Oregon with which Amazon has built long-term relationships. Excluding Dominion shifts point estimates and widens confidence intervals (Table D.13).

A.3 Takeaways

Both the Bartik and LP-DiD approaches suggest negative retail price impacts from data centers, with point estimates at least a few times larger than the wholesale price increases we estimate using our dispatch model. However, neither design is credible. Because of the extreme geographic concentration of data center development, the shift-share result is identified primarily on a single utility, Dominion Virginia. The local projections approach compares treated and control units that differ in fundamental ways, and again relies on a tiny fraction of all utilities. Parallel trends are unlikely to hold.

The failure of standard reduced form approaches motivates our use of a dispatch model. Our dispatch model does not assume that treated units are comparable to untreated units. It is also able to capture price effects that scale nonlinearly with dose and are heterogeneous across regions.

B Mapping load changes to local retail price changes

B.1 Regional dispatch

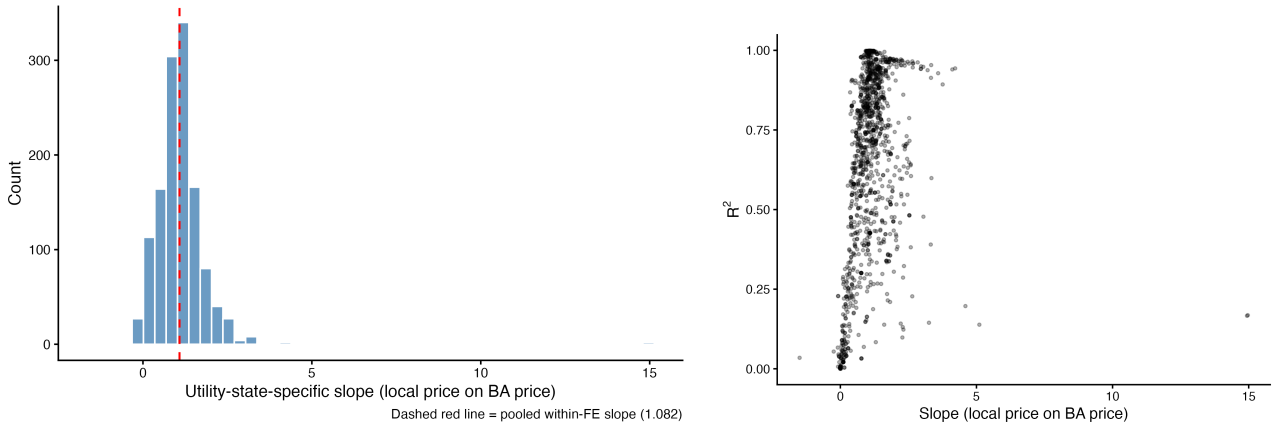
First, we build a set of splines that approximate the dispatch model described in Section 4 and can simulate region-level wholesale prices given different load scenarios. We construct supply curve splines based on data from each reference year (2014, 2019, 2024). For each region (Figure C.10), we estimate a piecewise linear interpolation over the 2019 merit order.

We define plant capacity as the 99th percentile of gross load within the year, and plant marginal cost as the generation-weighted mean of marginal costs across operating hours. We assume that in the baseline, before data center entry, demand exactly matches the hourly demand observed in the 2019 data. Although we assume that the supply curve is constant for all hours, variation in hourly demand generates heterogeneity in price impacts across hours and months. As in our dispatch model analysis in Section 4, we assume that only thermal generation responds to marginal load changes, and that there are no transmission constraints within regions. A natural enrichment of this model would be to separately estimate supply curve splines for different months in the year, to account for variation over time in fuel costs.

B.2 Relationship between local and regional wholesale price changes

We estimate the relationship between regional wholesale prices and local wholesale prices using a monthly panel (2016–2024) of dispatch model estimates at both geographies. On average, we find that a \$1/MWh change in regional market-clearing price (MCP) maps almost one-for-one to changes in local MCPs: a pooled regression yields a coefficient of 1.08. However, utility-state-specific regressions reveal substantial heterogeneity. Though the median slope is 1.04, the interquartile range spans 0.72 to 1.38 (Figure B.1). Differences in slope reflect differences in nearby generation assets: in utility-states where local MCPs move more than regional MCPs (slope $>$ 1), the local supply curve must be steeper. Although this is a rough proxy for how local wholesale prices should be affected by regional wholesale price variation, and one that may well be endogenous as load changes, we believe it is a reasonable first approximation. We truncate slopes to lie between 0 and 2 to avoid using implausible or weakly identified values.

Figure B.1: Utility-state-specific fit of local MCPs on regional MCPs

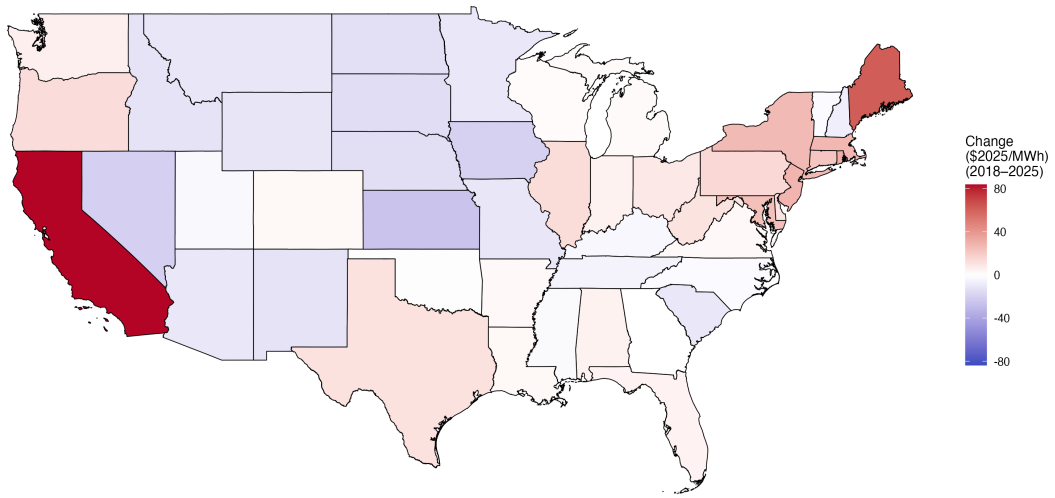


(a) Distribution of utility-state-specific slopes **(b)** Slope vs. within-utility-state R^2

Notes: Utility-state-specific regressions of local (50-mile buffer) MCPs on regional MCPs, monthly panel 2016–2024. Each utility-state's local MCP is regressed on the MCP estimated for its corresponding region. The dashed red line in (a) marks the pooled within-FE slope (1.082).

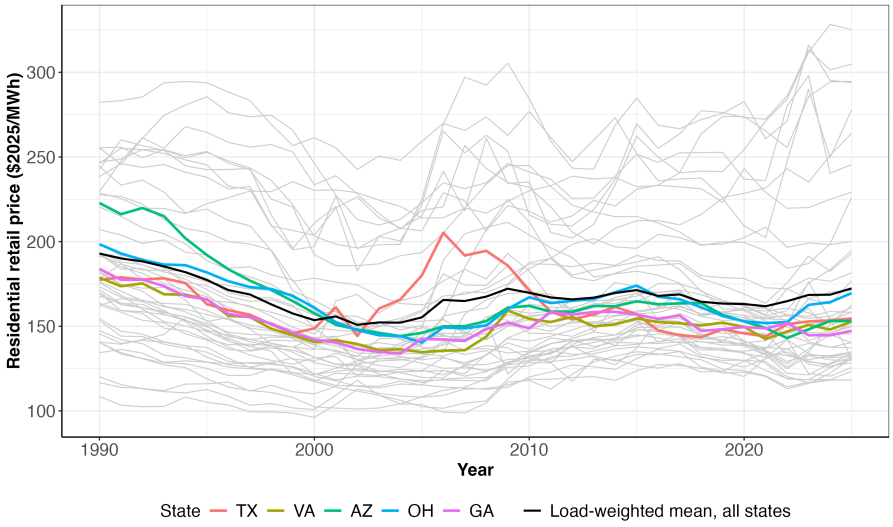
C Figures

Figure C.1: Map of state residential retail price changes, 2018–2025

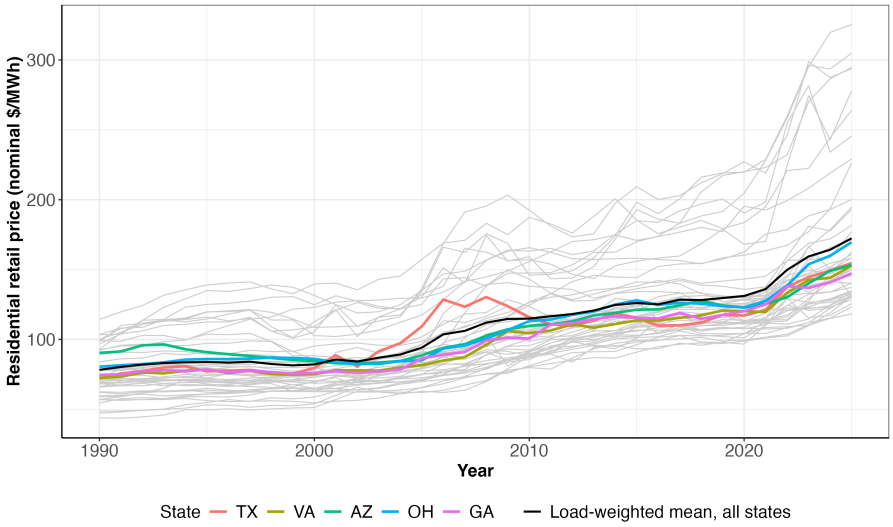


Notes: We plot state-level changes in residential retail prices between 2018 and 2025 ($2025/MWh). Source: EIA-861M monthly; CPI-U deflation.

Figure C.2: Residential retail prices by state, real and nominal



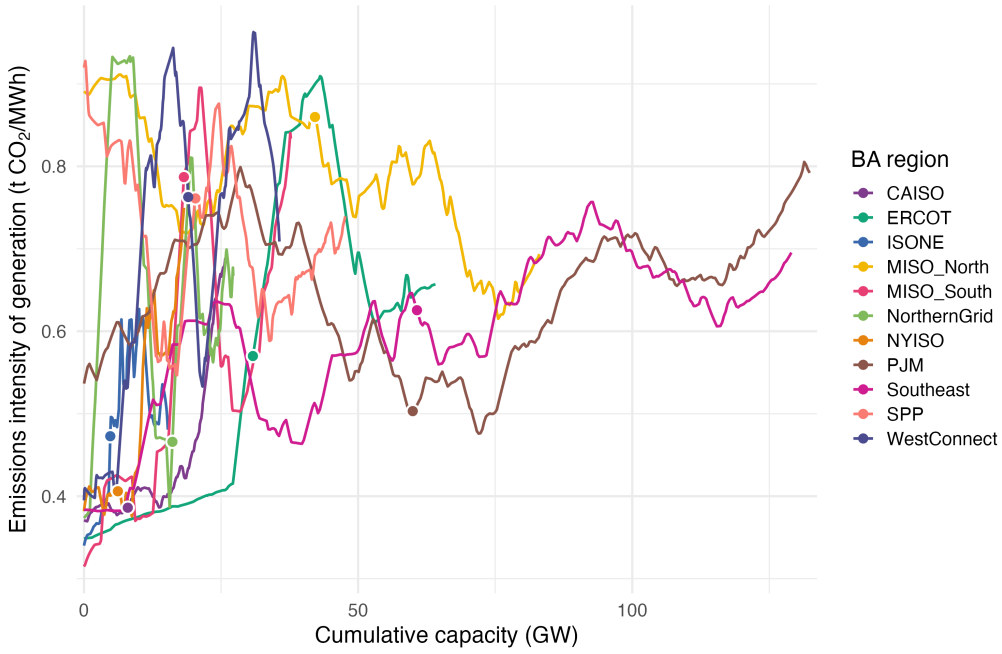
(a) Real ($2025/MWh)



(b) Nominal ($/MWh)

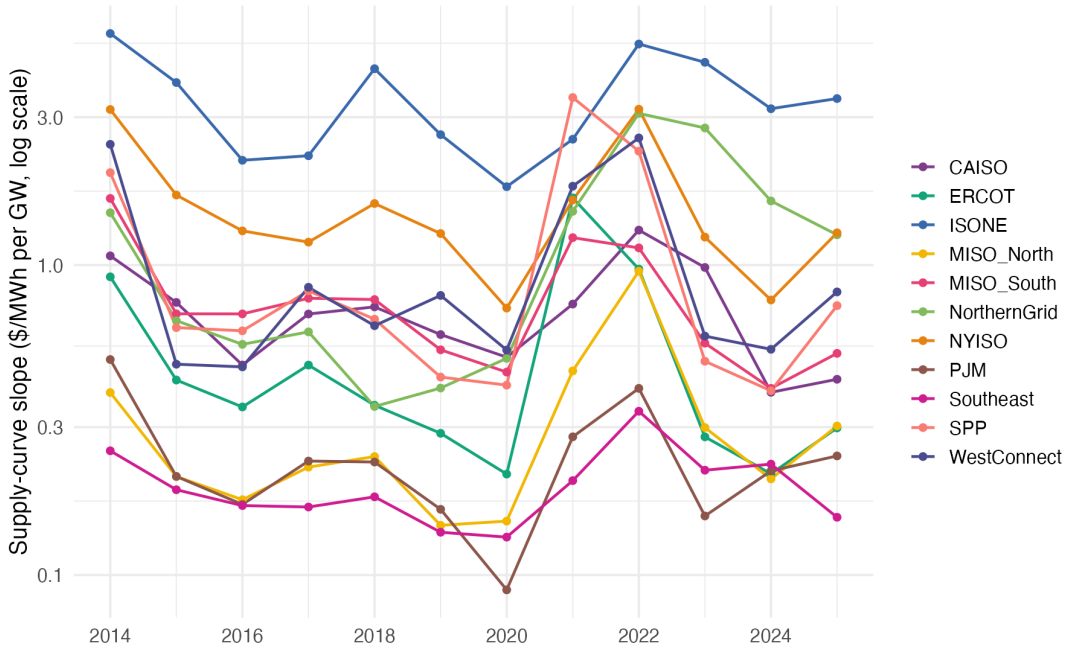
Notes: Residential retail price by state, highlighting the five states with the largest 2020–2025 data center buildout (TX, VA, AZ, OH, GA). Annual state price is summed residential revenue divided by summed residential sales (EIA-861M monthly data aggregated to calendar years with all 12 months reported; 2025 preliminary). Panel (a) is deflated to $2025/MWh; panel (b) is in nominal dollars. The black line is the load-weighted average across all other states (their summed residential revenue over summed residential sales); each grey line is a state outside the top five by data center buildout. 1990–2025.

Figure C.3: 2019 regional dispatch supply curves: marginal emissions intensity



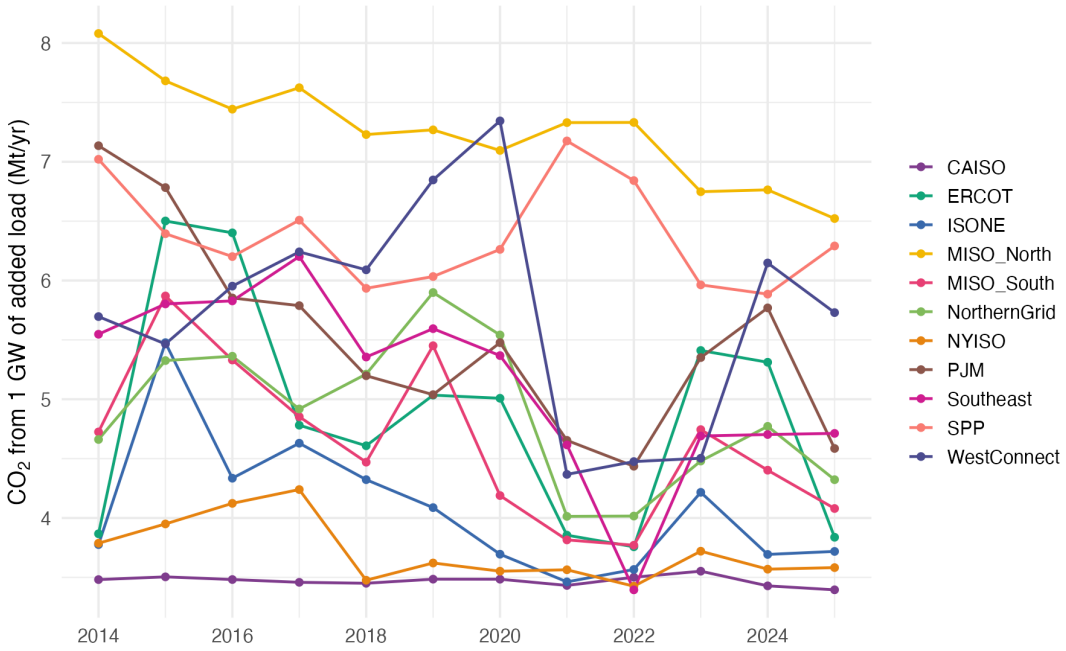
Notes: Merit-order supply curves for each of the 11 dispatch regions, 2019, with plants ordered cheapest to most expensive by marginal cost. The horizontal axis is cumulative capacity in GW, so each region's curve ends at its total capacity, where each plant's capacity is the 99th percentile of its hourly gross load over the year. EPA CEMS hourly data (2019). The marginal plant's CO_2 emissions intensity (metric tCO_2/MWh) is each plant's gross-load-weighted annual CO_2 mass per MWh winsorized at the 99th percentile, shown as a capacity-weighted rolling mean over a window of 15% of each region's total capacity. The dot on each curve marks the region's mean hourly demand in 2019.

Figure C.4: Supply-curve slope by region and year



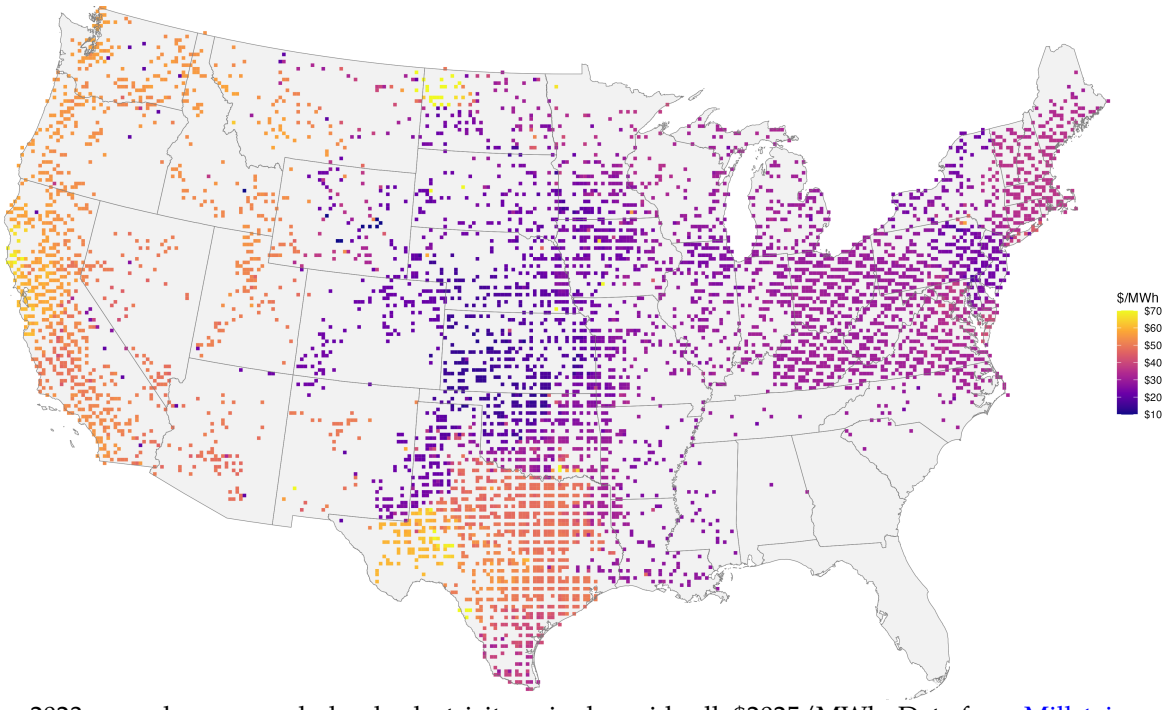
Notes: For each of the 11 dispatch regions and each dispatch year (2014–2025), the demand-weighted annual mean incremental market-clearing price of a +1 GW load band added to that year's annual merit order ($2025/MWh per GW; log scale). Hourly price increments are winsorized at the region–year 99th percentile. February 2021 ERCOT and SPP (Winter Storm Uri) hours are excluded. EPA CEMS hourly data.

Figure C.5: Marginal emissions by region and year



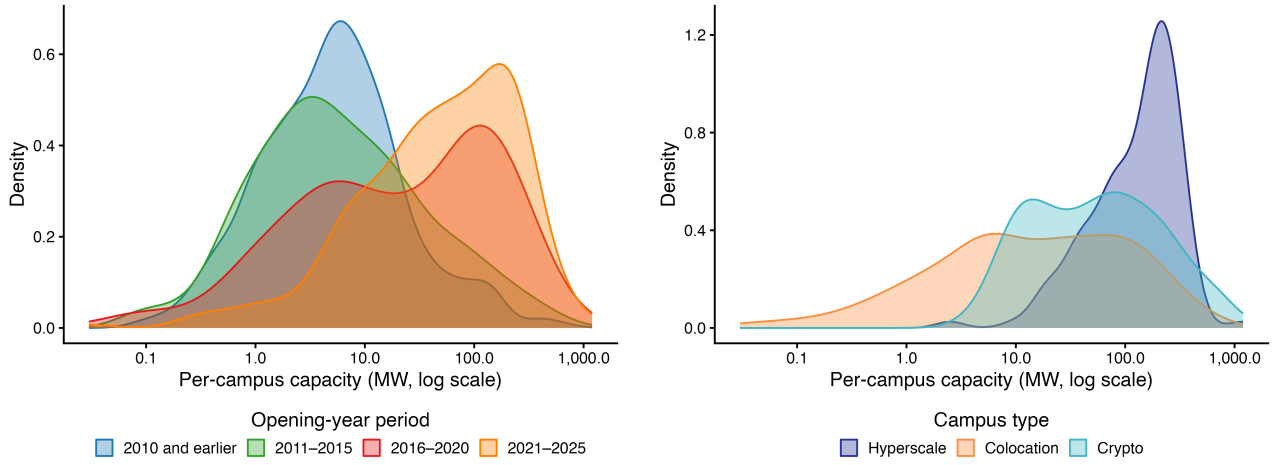
Notes: For each of the 11 dispatch regions and each dispatch year (2014–2025), the annual CO_2 emissions from adding 1 GW of constant load to that year's annual merit order, in megatons (Mt) per year. We calculate this as the sum over hours in each year of the incremental CO_2 emissions from adding 1 GW of load to observed hourly fossil generation. February 2021 ERCOT and SPP (Winter Storm Uri) hours are excluded.

Figure C.6: LBNL gridcell wholesale prices, 2023



Notes: 2023 annual average wholesale electricity price by grid cell, $2025/MWh. Data from Millstein, O'Shaughnessy and Wiser (2026).

Figure C.7: Campus size distributions by period and type

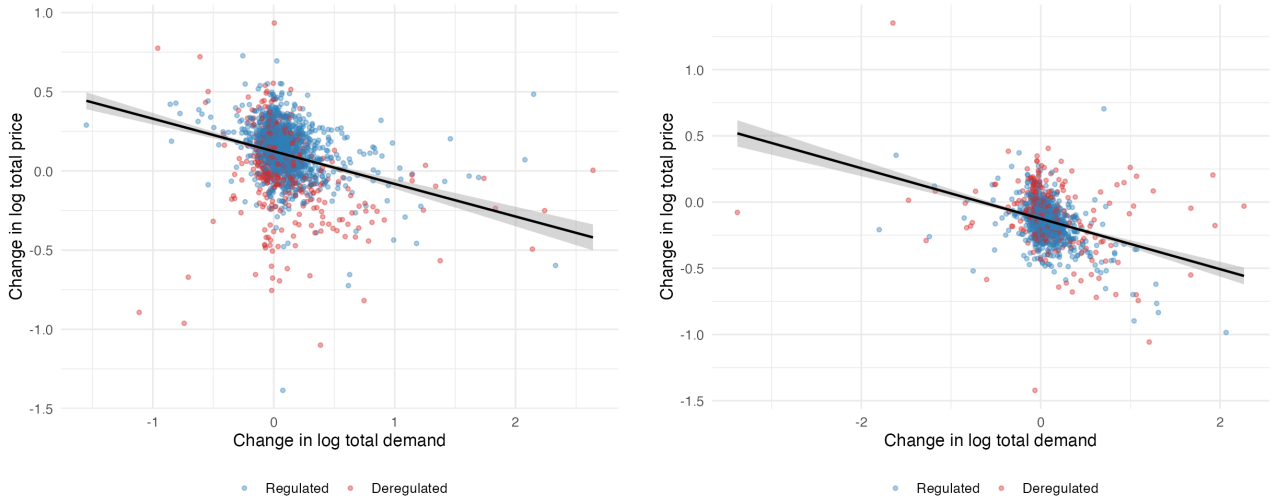


(a) By opening-year period

(b) By campus type (2016–2025)

Notes: Kernel densities of per-campus capacity (MW); each curve is estimated in $\log_{10}$ space and area-normalized to one, so shapes are comparable despite unequal campus counts, and the horizontal axis is on a log scale. Campus data from Schlauch et al. (2026); campuses with no reported capacity are excluded. Panel (a): one curve per opening-year period (2010 and earlier; 2011–2015; 2016–2020; 2021–2025), further excluding campuses with no opening year or opening after 2025. Panel (b): campuses opening 2016–2025, one curve per campus type (hyperscale, colocation, crypto).

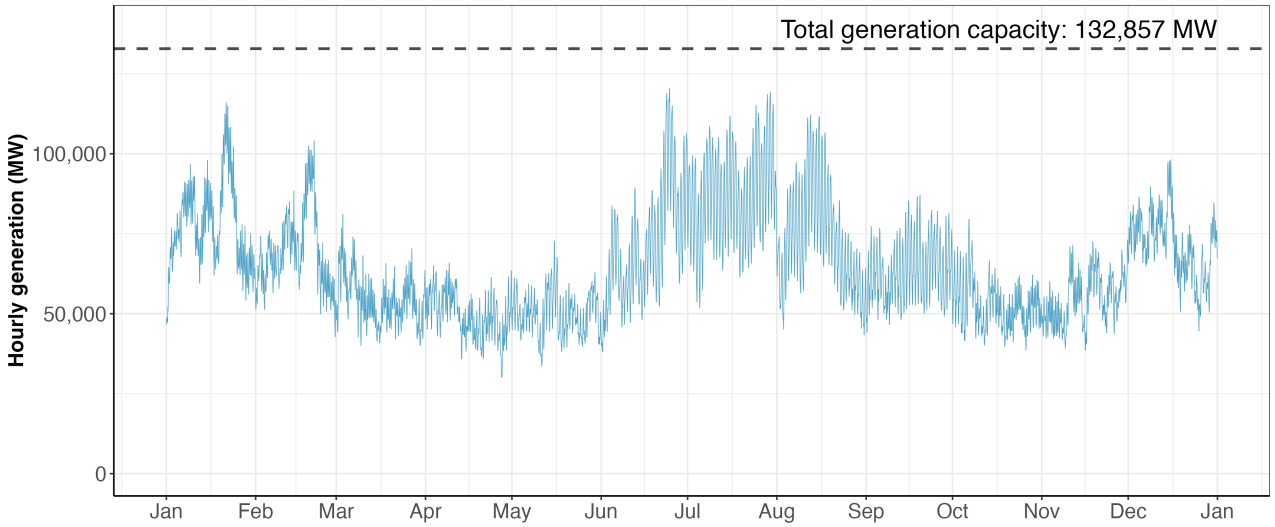
Figure C.8: Total demand and price changes across utility-states



(a) 2005–2015 **(b)** 2015–2024

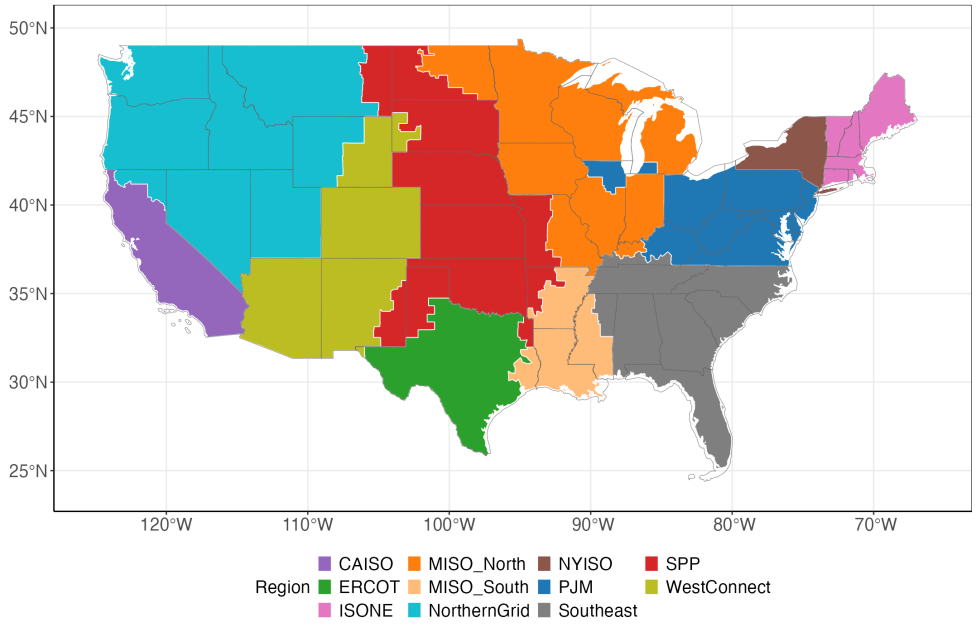
Notes: Each point is a U.S. utility-state; both panels plot the change in log total (all-sector) retail price against the change in log total electricity sales over the stated period, from the annual EIA-861 utility-state panel (prices deflated to \$2025/MWh). Sample: utility-states with total sales of at least 50,000 MWh in both end-point years—dropping a sparse small-utility tail—and with total prices in $(0, \$1{,}000]$/MWh. Points are colored by whether the state has retail-choice ("deregulated") electricity markets; the black line is an ordinary-least-squares fit to all points, with a 95% confidence band. Panel (a) covers 2005–2015 ($n = 1{,}697$); panel (b) covers 2015–2024 ($n = 1{,}267$).

Figure C.9: PJM's hourly generation and total capacity, 2025



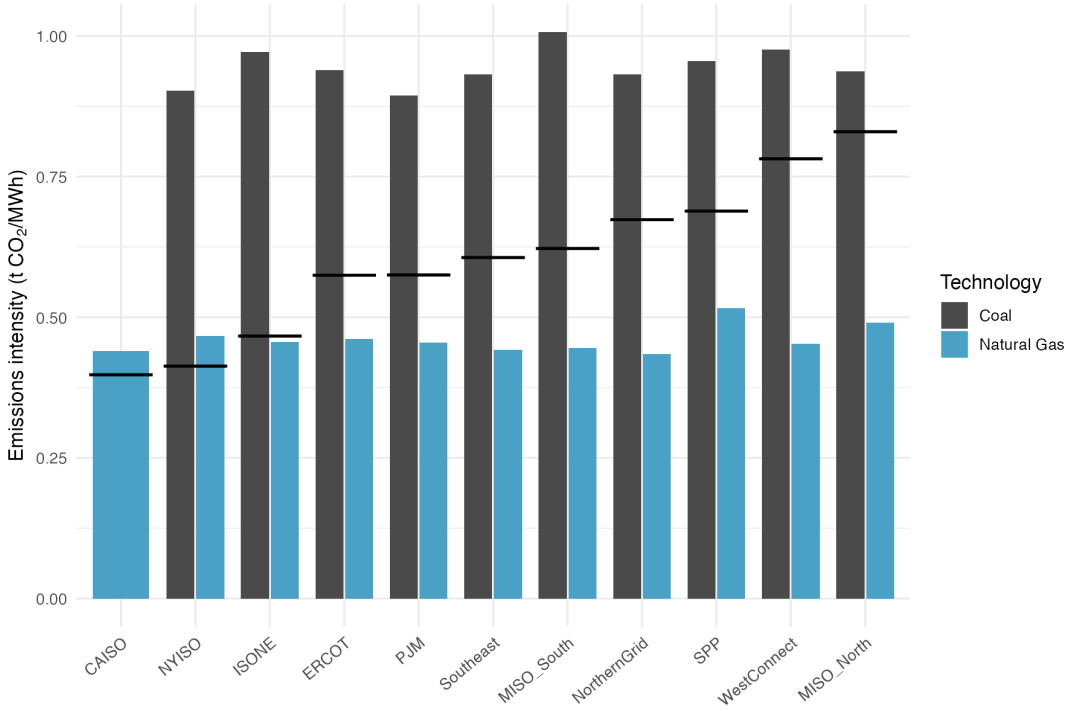
Notes: Hourly generation in the PJM region over 2025, summed across the region's EPA CEMS plants (the fossil-fuel fleet). The dashed horizontal line is total capacity: the sum across plants of each plant's 99th-percentile hourly gross load. Plants are assigned to regions by their eGRID balancing-authority code for the RTO markets (MISO split at 36.5°N) and by county location otherwise. Hours are as reported in CEMS.

Figure C.10: Dispatch model regions



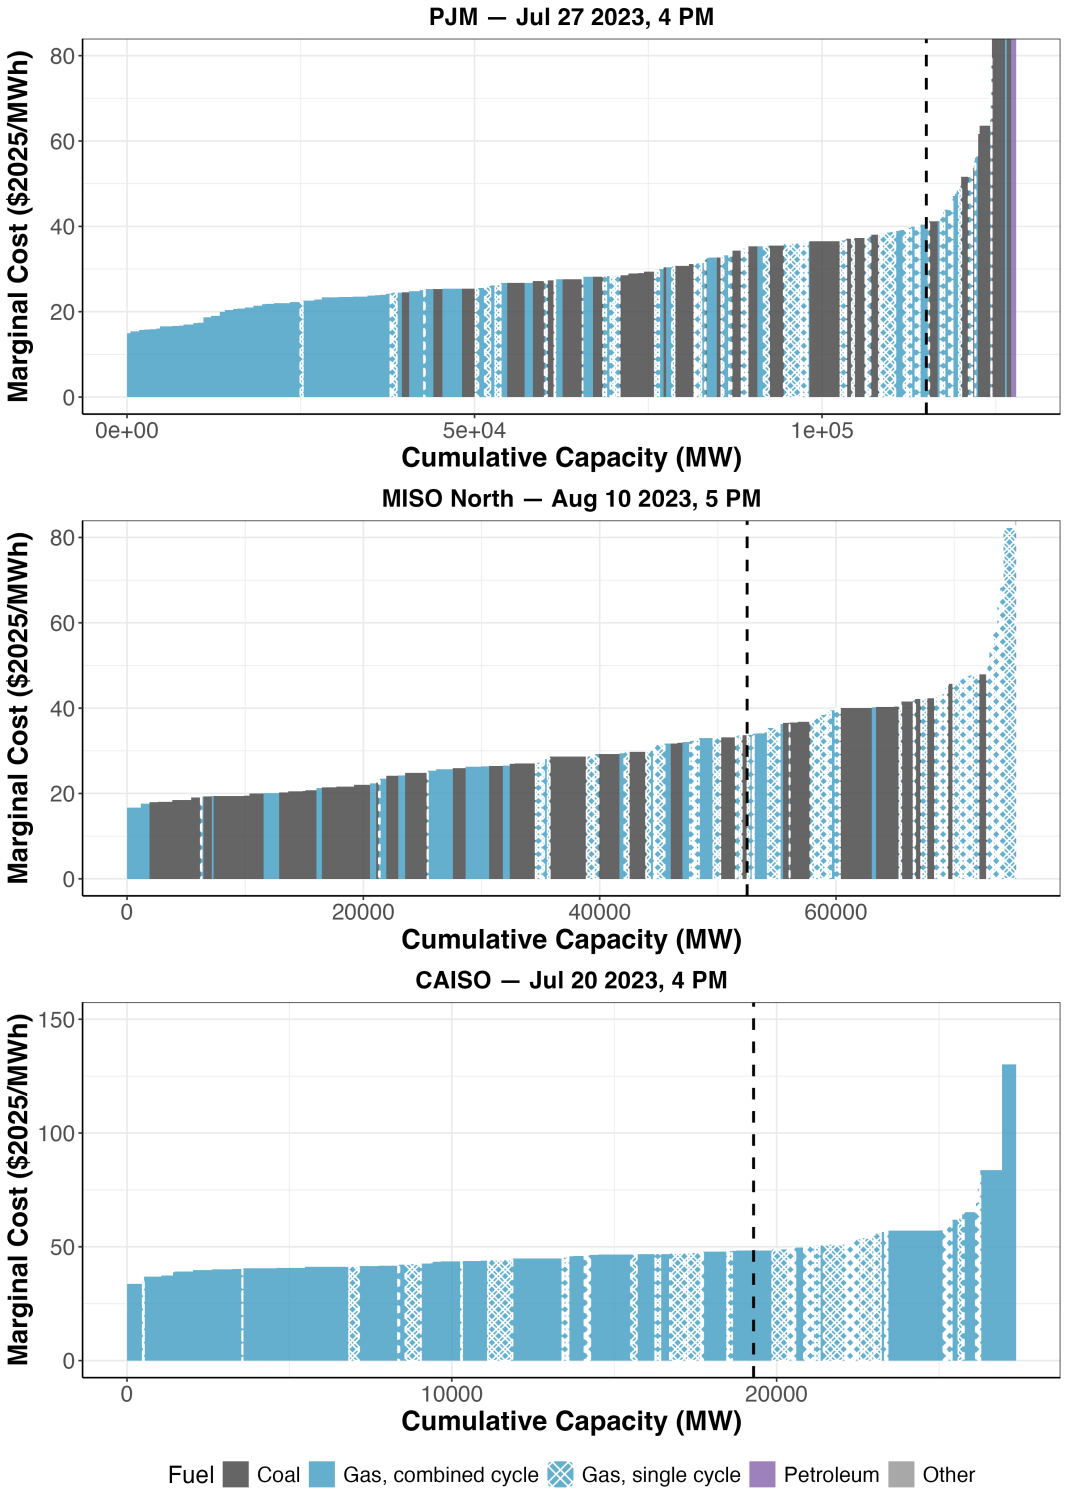
Notes: The 11 dispatch regions on which the model clears wholesale prices, formed as unions of Cambium's 18 GEA regions: the organized RTO markets (CAISO, ERCOT, ISO-NE, NYISO, PJM, SPP, and MISO split into North and South) plus three non-RTO residual regions (NorthernGrid, Southeast, WestConnect). Regions are shaded by color; grey lines are state boundaries. Contiguous U.S.

Figure C.11: Marginal emissions intensity vs. coal and gas fleet intensity, by region



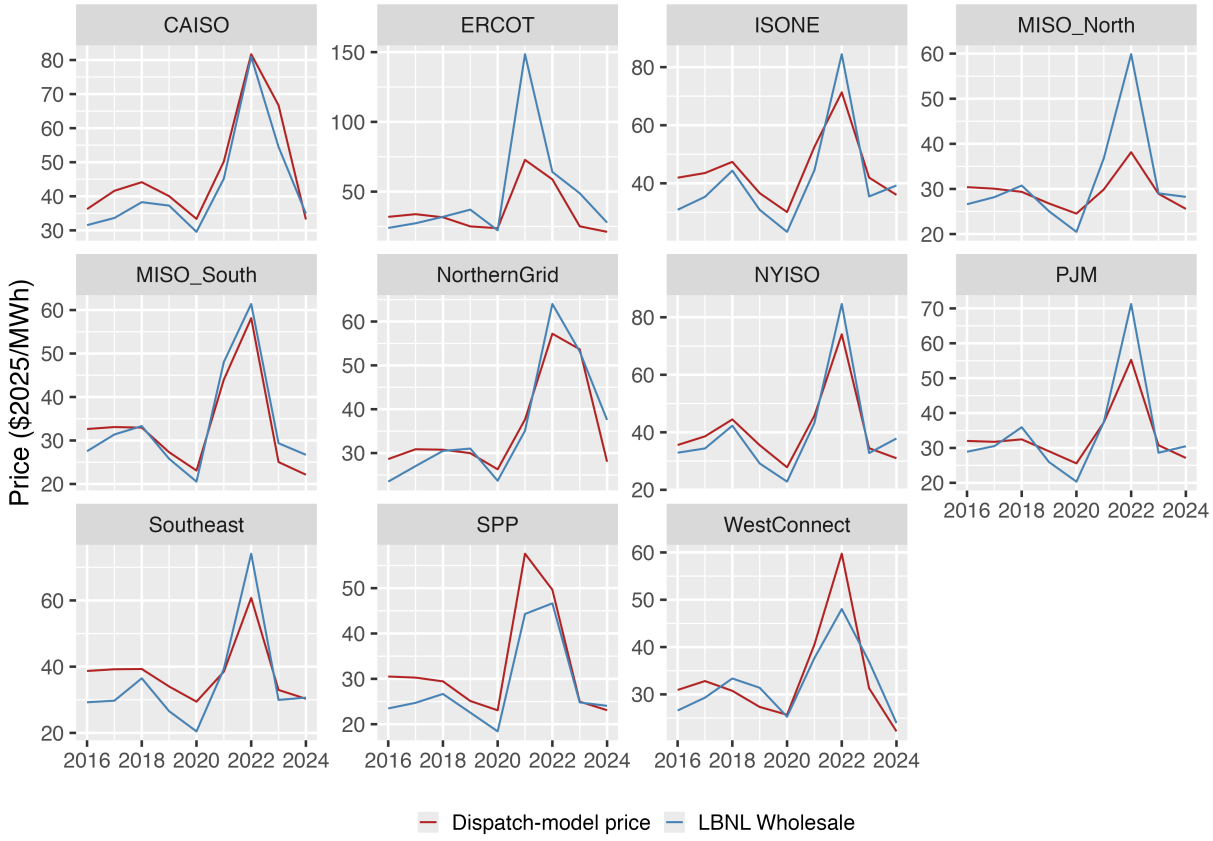
Notes: Each of the 11 dispatch regions, ordered by marginal emissions intensity in 2019. The two bars are the capacity-weighted mean CO_2 emissions intensity of the region's coal and natural-gas fleets (metric tCO_2/MWh), with each plant weighted by its capacity (the 99th percentile of its hourly gross load over the year). The black tick is the region's marginal emissions intensity, measured as the mean over the year's hours of the incremental generation emissions from adding 1 GW to that hour's demand (metric tCO_2/MWh). Per-plant CO_2 intensities are winsorized at the pooled 99th percentile (1.27 t/MWh) before both the fleet means and the marginal emissions are calculated. EPA CEMS hourly data (2019).

Figure C.12: Example supply curves



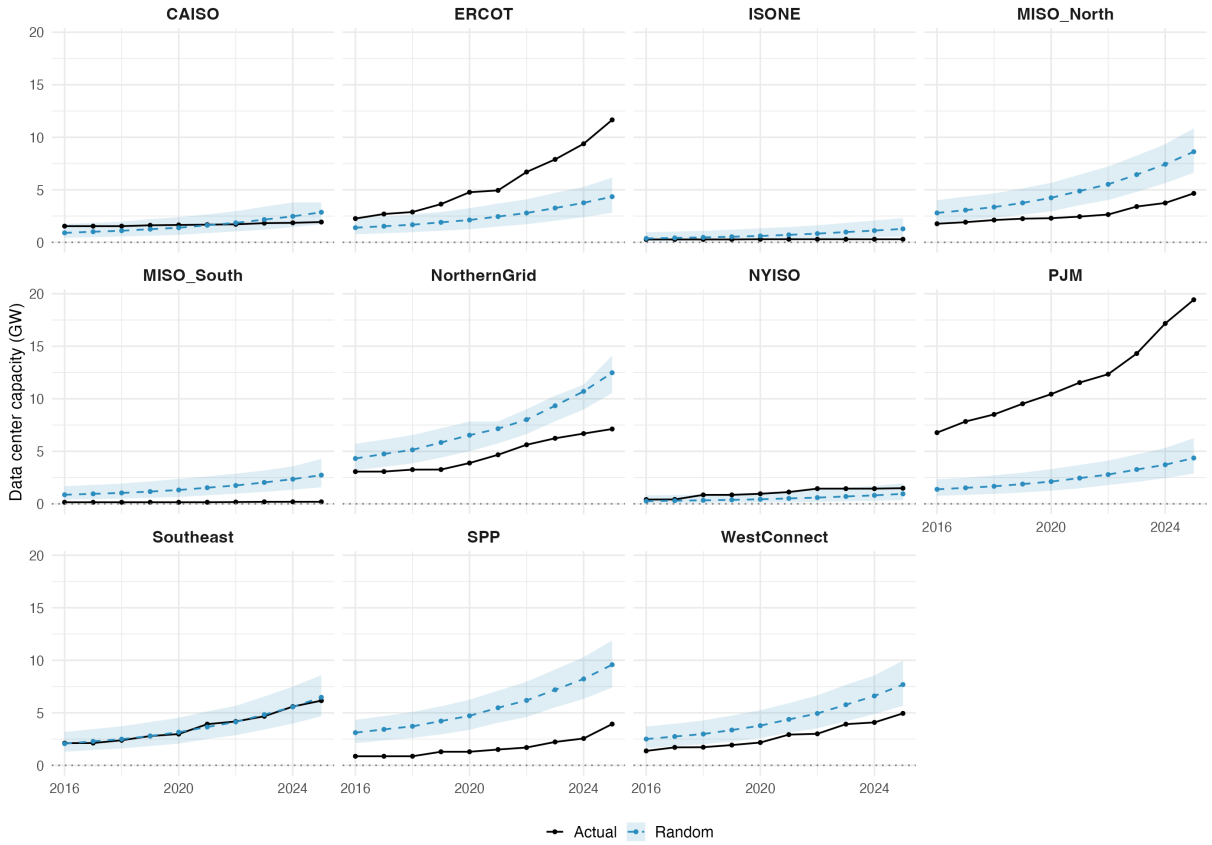
Notes: Plant-fuel-level marginal cost ($/MWh, deflated to $2025) plotted against cumulative capacity for three different region-hours. In each, plants are sorted in ascending order of marginal cost. We define plant capacity as the 99th percentile of hourly gross load in the 2023 CEMS; marginal cost is the generation-weighted mean of plant-hour CEMS marginal cost. Bars are colored by fuel type (coal, natural gas, petroleum, other). The dashed vertical line indicates realized demand (total gross load across dispatched plants) at that hour.

Figure C.13: LBNL wholesale prices vs. modeled prices by region



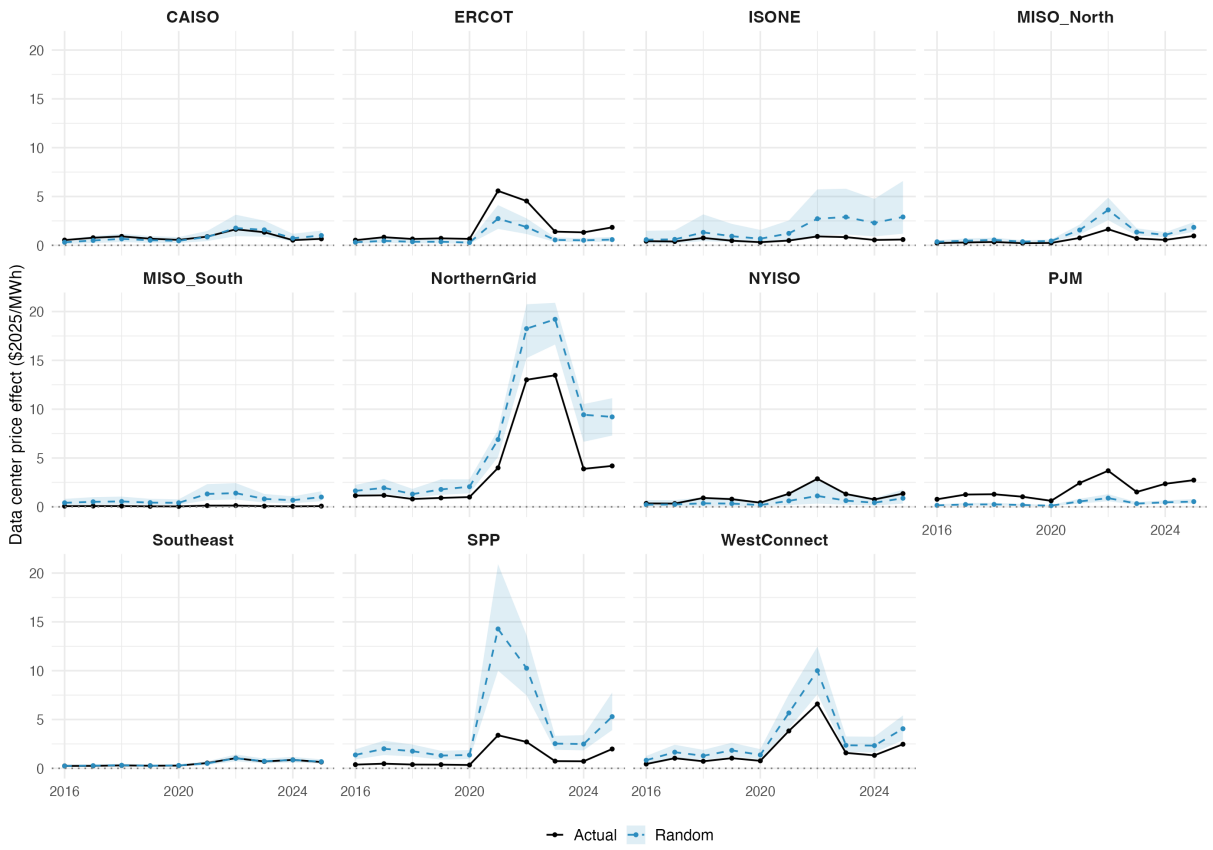
Notes: Annual time series of observed LBNL wholesale prices ($2025/MWh) and prices calculated from the dispatch model, by region, 2016–2024.

Figure C.14: Cumulative data center capacity by region



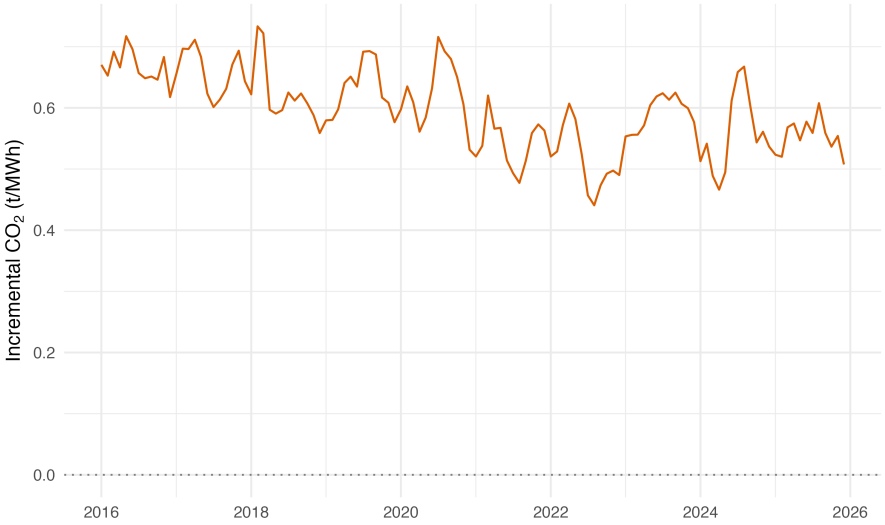
Notes: Cumulative data center capacity (GW) by region over time, actual and random. Constructed from campus-level nameplate MW for campuses with reported open years. Randomization is campus-level.

Figure C.15: Data center-induced wholesale price effect by region



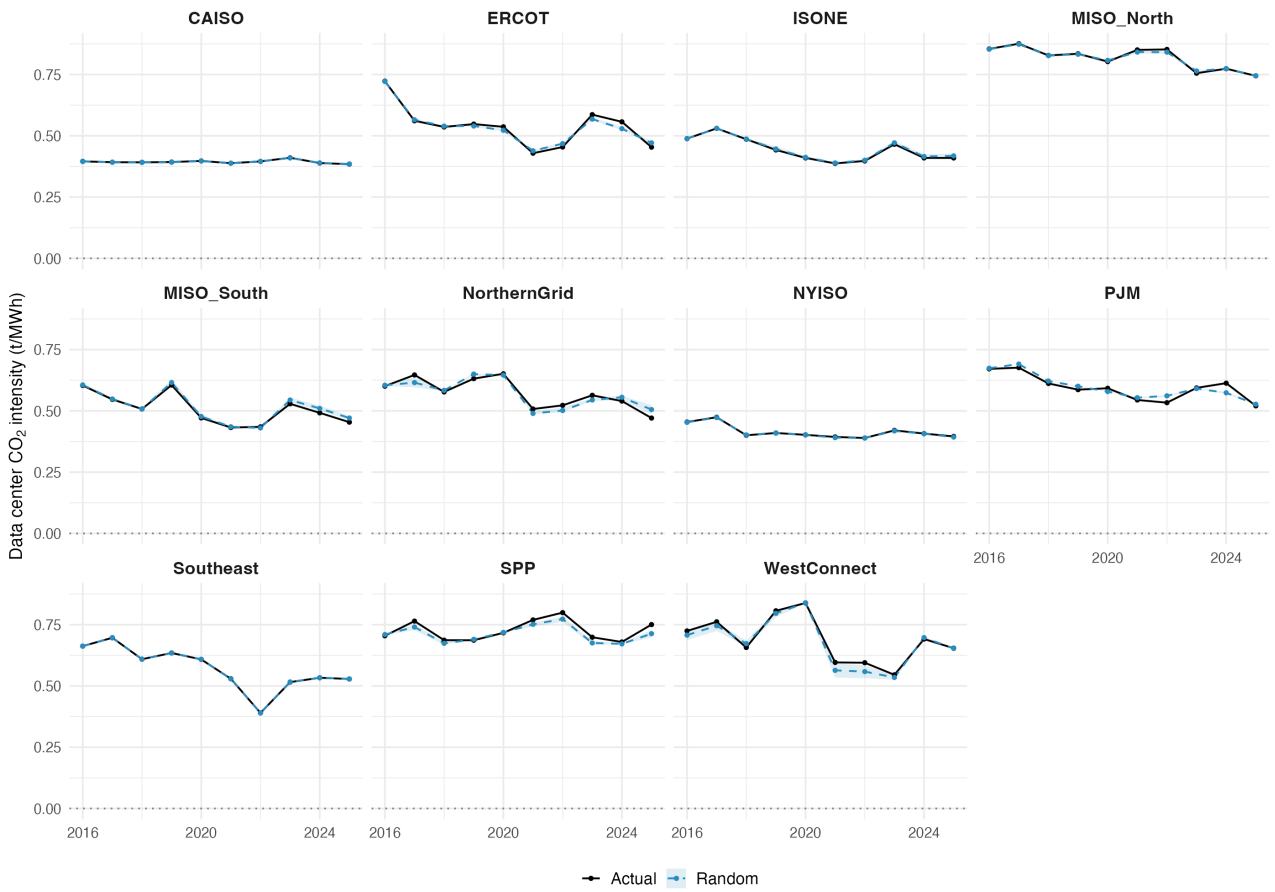
Notes: Annual wholesale price effect of cumulative data center load (Schlauch et al. 2026) in each dispatch region, evaluated year by year on each year's own dispatch grid. The price effect is the dispatch-model marginal price at demand including data center load minus the marginal price at demand net of that load, averaged over hours and weighted across months by regional demand; load beyond thermal capacity is curtailed rather than priced. Actual: nets the additions' load out of observed regional demand. Random: all campuses assigned randomly to regions with probability proportional to region geographic area, entering in their actual opening years, subject to a generation capacity screen; bands are 90% intervals over 2,000 draws, line indicates median across draws. Data center load = nameplate × 0.70. February 2021 is excluded for ERCOT and SPP (Winter Storm Uri).

Figure C.16: National marginal emissions intensity of data center load



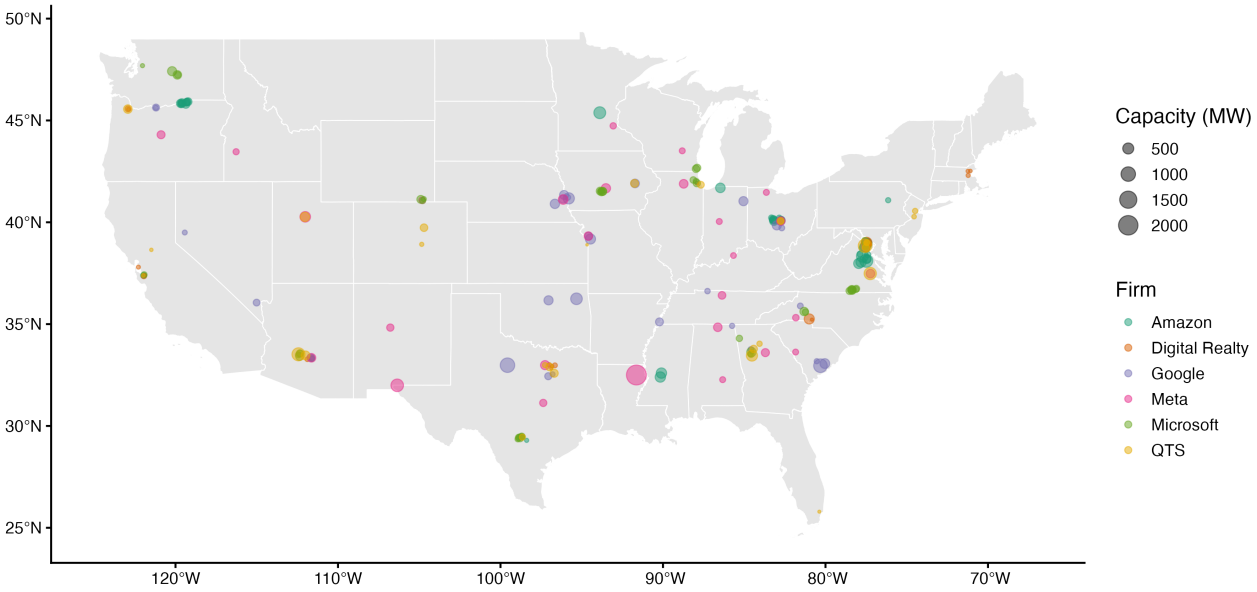
Notes: Monthly national marginal emissions intensity of data center load (metric tCO_2 per MWh): the incremental grid CO_2 attributable to data center load divided by that load. Incremental CO_2 is the dispatch-model emissions at observed region-level demand minus emissions at demand less realized data center load, summed across region-hours to the month; it is divided by the summed data center load (MWh) that month. We assume that data center load is equal to 70% of nameplate capacity, and that data centers operate continuously. February 2021 is excluded for ERCOT and SPP (Winter Storm Uri).

Figure C.17: Marginal emissions intensity of data center load, by region



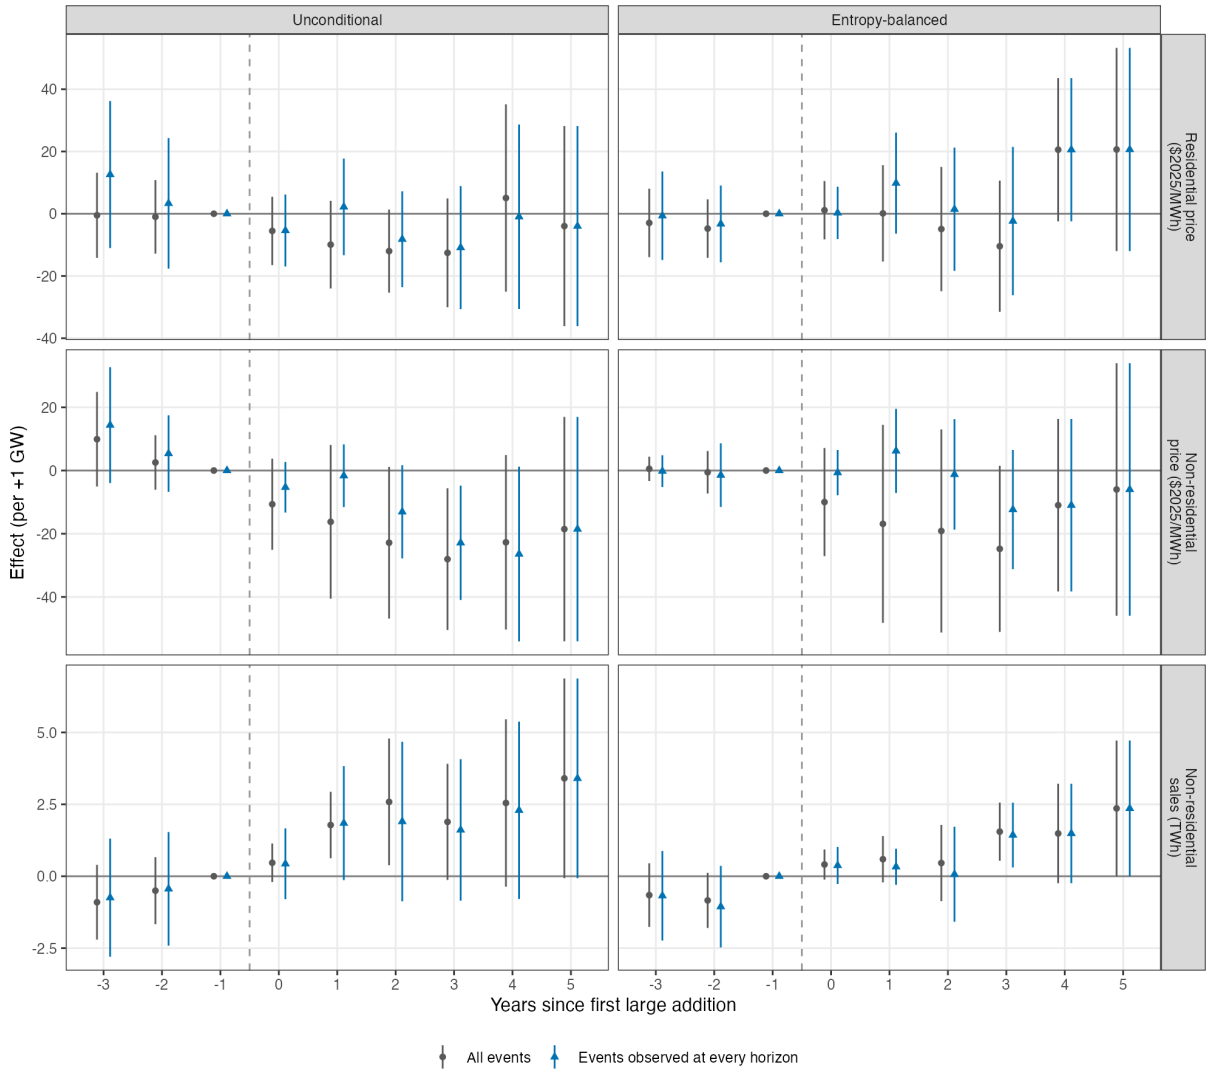
Notes: Annual marginal emissions intensity of cumulative data center load (Schlauch et al. 2026) in each dispatch region (metric tCO_2 per MWh), evaluated year by year on each year's own dispatch grid: the incremental grid CO_2 attributable to data center load divided by that load. Incremental CO_2 is the dispatch-model emissions at demand including data center load minus emissions at demand net of that load, summed over the year's hours. Actual: nets the additions' load out of observed regional demand. Random: all campuses assigned randomly to regions with probability proportional to region geographic area, entering in their actual opening years, subject to a generation capacity screen; intensity is computed draw by draw, bands are 90% intervals over the 2,000 draws that place load in the region, line indicates median across draws. We assume that data center load is equal to 70% of nameplate capacity, and that data centers operate continuously. February 2021 is excluded for ERCOT and SPP (Winter Storm Uri).

Figure C.18: Focus firms' campuses



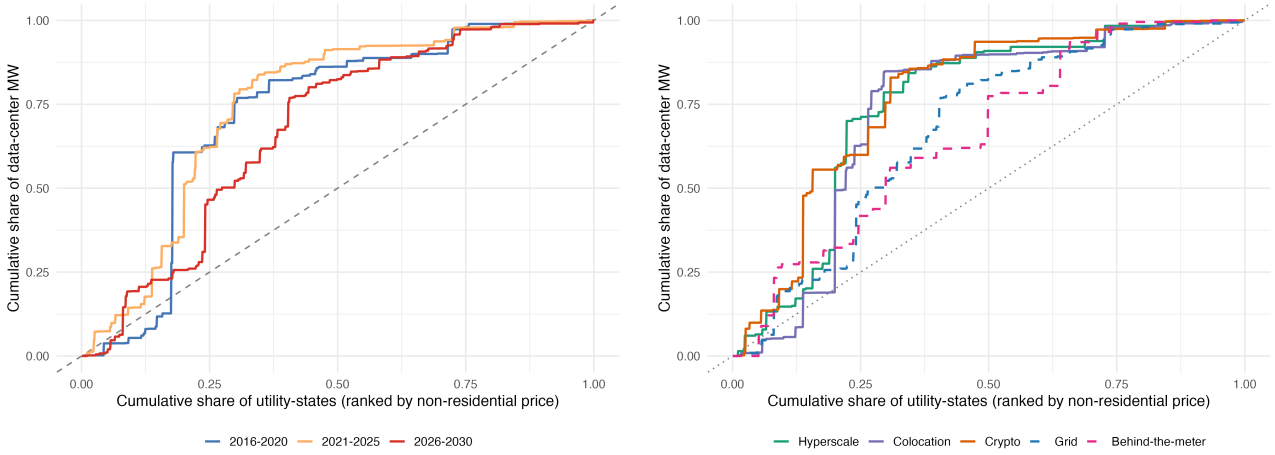
Notes: Data center campuses of the six focus firms used to construct the firm Bartik instrument, colored by firm and sized by capacity (MW). Contiguous U.S.; campuses with a reported opening year and capacity. Source: campus data from Schlauch et al. (2026).

Figure C.19: Local projections DiD: retail price and load effects of data center entry, within-region controls



Notes: Event-stacked local projections difference-in-differences on a utility-state annual panel (2000–2024), with event fixed effects. An event is the first year a utility-state's annual data center capacity addition reaches 100 MW, restricted to events in or after 2001; the dose is that crossing-year addition, rescaled so estimates are per $+1$ GW. Each treated utility-state is stacked with clean controls from *inside* its own dispatch region: utility-states with no $\geq$100 MW addition through $t_0{+}5$, so incidental sub-threshold data center capacity is allowed. Region-wide price movements common to the host and its neighbors difference out, so the coefficient is the within-region differential. Rows are the residential retail price and the non-residential retail price ($2025/MWh) and non-residential sales (TWh); horizons -3 to 5, with $h = -1$ the reference. Unconditional estimates weight events by $1/\text{dose}$; entropy-balanced estimates additionally weight controls to match treated units on deregulated-market status and on the changes in non-residential price and log(1 + non-residential sales) from $t_0{-}3$ to $t_0{-}1$, and drop events whose controls cannot be balanced. Grey circles use all events (43 for the residential price, 44 for the non-residential outcomes); blue triangles use only events whose treated utility-state is observed at every horizon from $t_0{-}3$ to $t_0{+}5$ (23 and 24). Vertical lines are 95% confidence intervals from state-clustered standard errors.

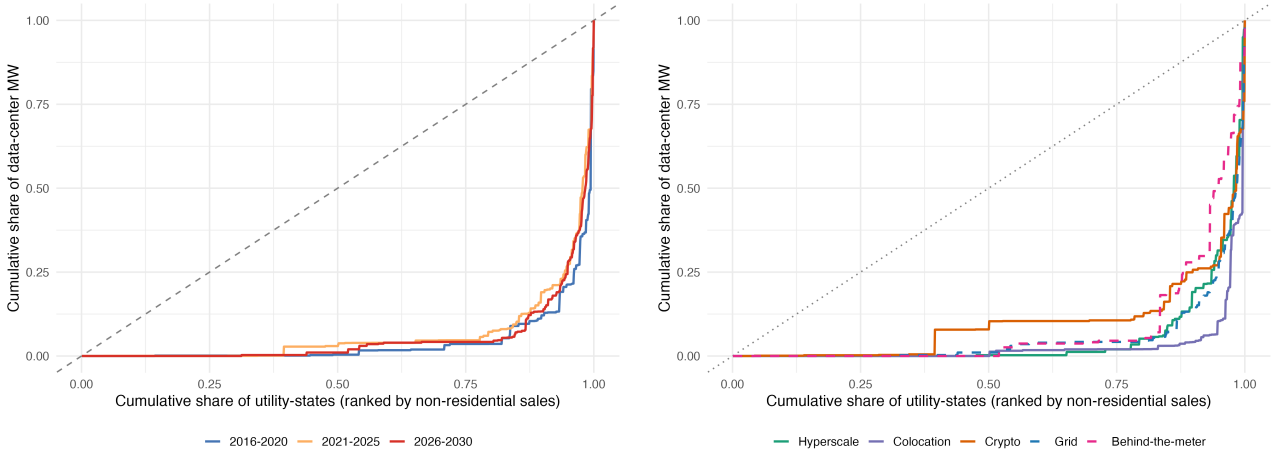
Figure C.20: Data center capacity concentration by non-residential price



(a) By buildout period **(b)** By campus type

Notes: Concentration curve of cumulative data center capacity (cumulative share of MW, y-axis) against utility-states ranked by non-residential retail price (cumulative share of utility-states, x-axis). Panel (a): one line per buildout period, each ranked by its reference-year price (2014 for 2016–2020, 2019 for 2021–2025, 2024 for 2026–2030); the 2026–2030 line is the announced Aterio high-likelihood grid-connected net-new buildout. Panel (b): one line per campus type—hyperscale, colocation, and crypto over realized 2016–2025 openings (solid, ranked by 2019 price) and grid and behind-the-meter for the 2026–2030 Aterio buildout (dashed, ranked by 2024 price). Realized campuses are from Schlauch et al. (2026). Sample is the 1,274 utility-states in the contiguous U.S. with valid non-residential price and sales in all three reference years. The 45-degree line corresponds to a uniform allocation across the price rank.

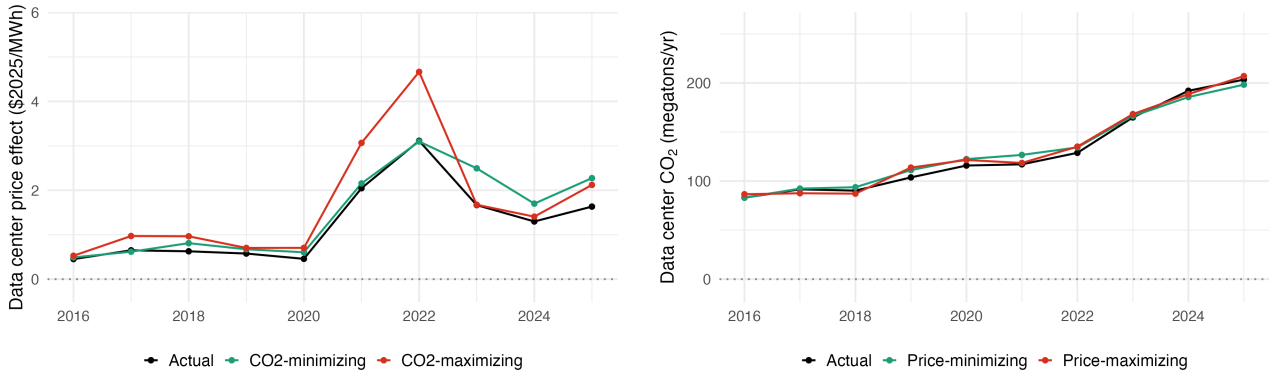
Figure C.21: Data center capacity concentration by non-residential sales



(a) By buildout period **(b)** By campus type

Notes: Concentration curve of cumulative data center capacity (cumulative share of MW, y-axis) against utility-states ranked by non-residential electricity sales (cumulative share of utility-states, x-axis). Panel (a): one line per buildout period, each ranked by its reference-year sales (2014 for 2016–2020, 2019 for 2021–2025, 2024 for 2026–2030); the 2026–2030 line is the announced Aterio high-likelihood grid-connected net-new buildout. Panel (b): one line per campus type—hyperscale, colocation, and crypto over realized 2016–2025 openings (solid, ranked by 2019 sales) and grid and behind-the-meter for the 2026–2030 Aterio build-out (dashed, ranked by 2024 sales). Realized campuses are from Schlauch et al. (2026). Sample is the 1,274 utility-states in the contiguous U.S. with valid non-residential price and sales in all three reference years. The 45-degree line corresponds to a uniform allocation across the sales rank.

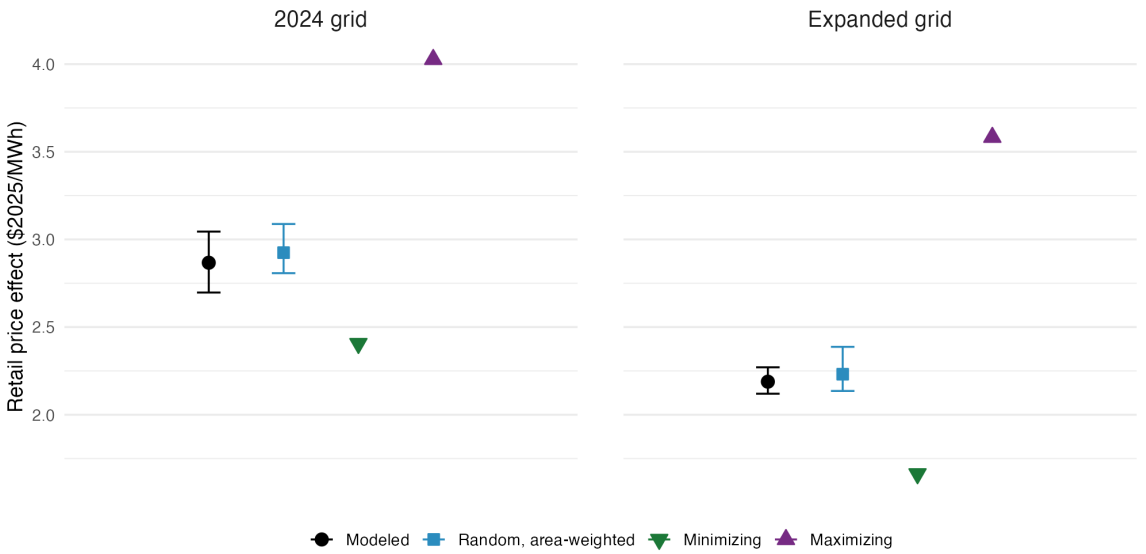
Figure C.22: Cross-objective effects of the optimizing allocations



(a) Price effect of CO_2-optimizing allocations **(b)** CO_2 of price-optimizing allocations

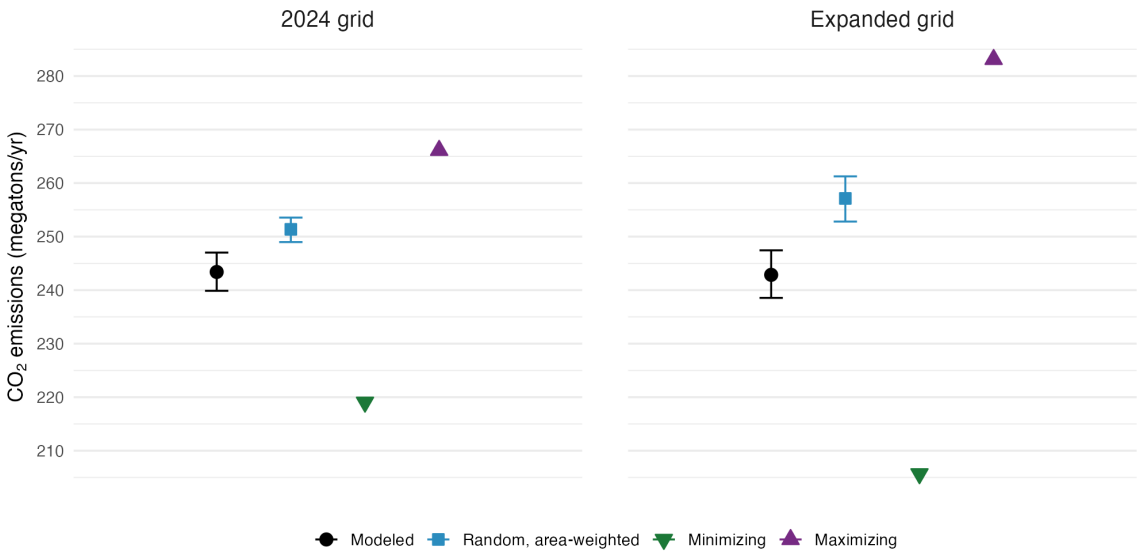
Notes: Annual outcomes of cumulative data center load under the sequential optimizing allocations, evaluated year by year on each year's own dispatch grid. In each panel the plotted outcome is the one the allocation does not optimize: panel (a) reports the demand-weighted wholesale price effect ($2025/MWh) of the climate damage-minimizing and climate damage-maximizing allocations, and panel (b) reports incremental grid CO_2 (megatons per year) of the price-minimizing and price-maximizing allocations. All campuses are reallocated, subject to a generation capacity screen. Price effects are demand-weighted across hours, months, and regions. Data center load = nameplate $\times$ 0.70. February 2021 is excluded for ERCOT and SPP (Winter Storm Uri).

Figure C.23: Retail price effect by data center allocation, 2024 grid and augmented grid (2026–2030 buildout, 50% realization)



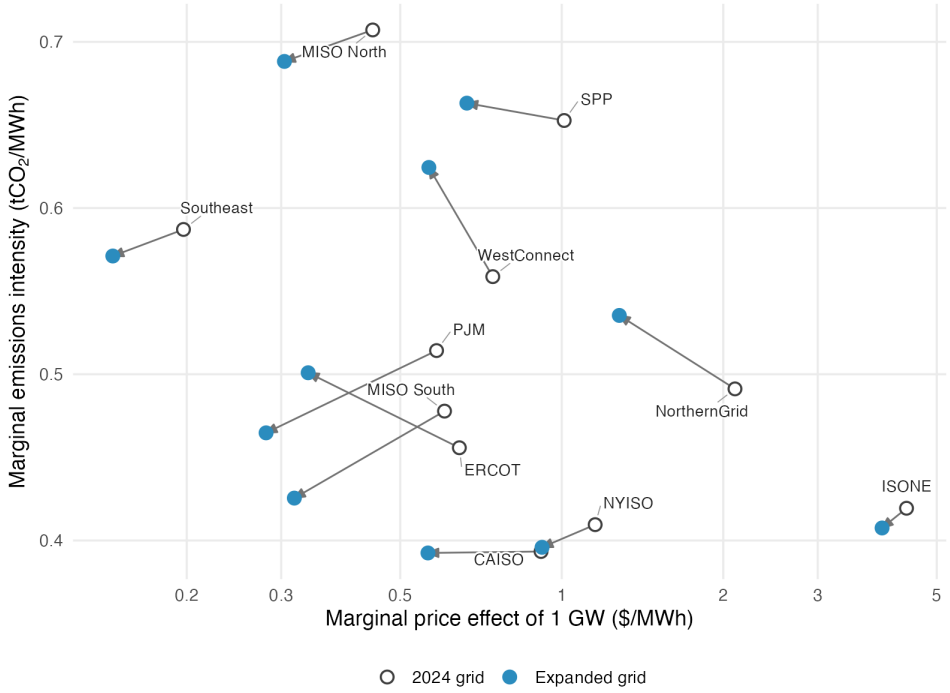
Notes: Residential-load-weighted retail price change ($2025/MWh) from the announced 2026–2030 grid buildout (Aterio high-likelihood pipeline; behind-the-meter excluded) at 50% realization, under four allocations across utility-states. *Modeled*: campuses drawn from the equilibrium choice probabilities under the Grid (2026–2030) preference estimates in Table 1. *Random*: campuses drawn in proportion to land area. For these two, markers are medians and bars are 5th–95th percentile ranges over 2,000 campus-level draws subject to a generation capacity screen. *Minimizing* and *Maximizing*: the allocations that minimize and maximize residential electricity costs under the same screen, found by a greedy placement followed by local search over single-campus moves and pairwise swaps. Left panel: 2024 merit order. Right panel: 2024 merit order plus all EIA-860 proposed thermal units, with planned wind and solar netted from thermal demand. The panels share one vertical axis. Location choice is evaluated with load beyond thermal capacity priced on a scarcity ramp rising linearly to $9,000/MWh over a band equal to 5% of thermal capacity; household costs are evaluated with that load unserved, so prices are bounded by the most expensive generator. Data center load = nameplate $\times$ 0.70.

Figure C.24: Incremental emissions by data center allocation, 2024 grid and augmented grid (2026–2030 buildout, 50% realization)



Notes: Incremental CO_2 (megatons per year), total-grid emissions with minus without data center load, from the announced 2026–2030 grid buildout (Aterio high-likelihood pipeline; behind-the-meter excluded) at 50% realization, under four allocations across utility-states. *Modeled*: campuses drawn from the equilibrium choice probabilities under the Grid (2026–2030) preference estimates in Table 1. *Random*: campuses drawn in proportion to land area. For these two, markers are medians and bars are 5th–95th percentile ranges over 2,000 campus-level draws subject to a generation capacity screen. *Minimizing* and *Maximizing*: the allocations that minimize and maximize climate damages under the same screen, found by a greedy placement followed by local search over single-campus moves and pairwise swaps. Left panel: 2024 merit order. Right panel: 2024 merit order plus all EIA-860 proposed thermal units, with planned wind and solar netted from thermal demand. The panels share one vertical axis. Location choice is evaluated with load beyond thermal capacity priced on a scarcity ramp rising linearly to \$9,000/MWh over a band equal to 5% of thermal capacity; emissions are evaluated with that load unserved. Data center load = nameplate $\times$ 0.70.

Figure C.25: Marginal price and emissions effects of added load by dispatch region, 2024 grid and augmented grid



Notes: Each point is one of the 11 dispatch regions under one grid; arrows run from the 2024 merit order to the augmented grid, which adds all EIA-860 proposed thermal units and nets planned wind and solar from thermal demand. Horizontal axis: change in the wholesale market-clearing price from adding 1 GW to each hour's demand, demand-weighted across hours ($2025/MWh, log scale). Vertical axis: incremental generation CO_2 from the same 1 GW addition, averaged across hours (metric tCO_2/MWh). Both are evaluated at one data center allocation held fixed across the two grids: the modeled expected allocation of the announced 2026–2030 grid buildout at 50% realization on the 2024 grid. Load beyond a region's thermal capacity is unserved. Data center load = nameplate $\times$ 0.70.

D Tables

Table D.1: Policy counterfactuals with data center locations fixed and with firms relocating (2026–2030 buildout, 50% realization)

		Carbon tax		Virginia surcharge	
Outcome	Baseline	Locations fixed	Firms relocate	Locations fixed	Firms relocate
Price effect ($/MWh)	2.87	16.17	16.46	2.87	2.82
Incremental CO_2 (Mt)	243.40	260.62	248.10	243.40	243.98
Total climate damages ($B/yr)	$326.1	$240.3	$238.0	$326.1	$326.2
Δ residential electricity costs ($B/yr)	$4.12	$23.25	$23.67	$4.12	$4.05
Δ climate damages ($B/yr)	$45.03	$48.21	$45.90	$45.03	$45.14
Δ government transfers ($B/yr)	$0.00	$48.21	$45.90	$0.73	$0.64
Δ consumer surplus ($B/yr)	-$49.15	-$23.25	-$23.67	-$48.41	-$48.56
Capacity in Virginia (GW)	10.84	10.84	10.28	10.84	9.48
Capacity relocated vs. baseline (GW)	0.00	0.00	7.78	0.00	1.31

Notes: Announced 2026–2030 grid buildout (Aterio high-likelihood pipeline; behind-the-meter excluded) at 50% realization, reallocated across utility-states using the Grid (2026–2030) preference estimates in Table 1. *Baseline* and *firms relocate* columns report the median of each row over 2,000 campus-level draws from the equilibrium choice probabilities, subject to a generation capacity screen; *locations fixed* columns re-evaluate the *Baseline* draws under the policy. Because rows are medians taken separately, the consumer surplus row need not equal the sum of its components. *Carbon tax*: $185/t$CO_2$ added to generator marginal costs, with dispatch re-sorted and revenue on data center–induced generation rebated to consumers. *Virginia surcharge*: $11/MWh on data center consumption in Virginia, paid by data centers only, with revenue rebated to consumers. Price effect is the residential-load-weighted retail price change from data center load. Incremental CO_2 is total-grid emissions with minus without data center load; total climate damages value total-grid emissions at the SCC. Residential electricity costs are positive for an increase. Consumer surplus $= -$costs $-$ climate damages $+$ transfers. Capacity relocated is measured on the expected allocation, which is not subject to the capacity screen. Location choice is evaluated with load beyond thermal capacity priced on a scarcity ramp rising linearly to $9,000/MWh over a band equal to 5% of thermal capacity; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC $185/t (Rennert et al. 2022); data center load $=$ nameplate $\times$ 0.70; 2024 dispatch.

Table D.2: Policy counterfactuals with data center locations fixed and with firms relocating (2026–2030 buildout, 50% realization, augmented grid)

		Carbon tax		Virginia surcharge	
Outcome	Baseline	Locations fixed	Firms relocate	Locations fixed	Firms relocate
Price effect ($/MWh)	2.19	13.16	13.04	2.19	2.18
Incremental CO_2 (Mt)	242.86	220.91	212.37	242.86	243.63
Total climate damages ($B/yr)	$287.7	$197.7	$196.1	$287.7	$287.9
Δ residential electricity costs ($B/yr)	$3.15	$18.93	$18.75	$3.15	$3.14
Δ climate damages ($B/yr)	$44.93	$40.87	$39.29	$44.93	$45.07
Δ government transfers ($B/yr)	$0.00	$40.87	$39.29	$0.67	$0.58
Δ consumer surplus ($B/yr)	-$48.07	-$18.93	-$18.75	-$47.39	-$47.61
Capacity in Virginia (GW)	9.98	9.98	9.46	9.98	8.61
Capacity relocated vs. baseline (GW)	0.00	0.00	6.51	0.00	1.32

Notes: Announced 2026–2030 grid buildout (Aterio high-likelihood pipeline; behind-the-meter excluded) at 50% realization, reallocated across utility-states using the Grid (2026–2030) preference estimates in Table 1. *Baseline* and *firms relocate* columns report the median of each row over 2,000 campus-level draws from the equilibrium choice probabilities, subject to a generation capacity screen; *locations fixed* columns re-evaluate the *Baseline* draws under the policy. Because rows are medians taken separately, the consumer surplus row need not equal the sum of its components. *Carbon tax*: $185/t$CO_2$ added to generator marginal costs, with dispatch re-sorted and revenue on data center–induced generation rebated to consumers. *Virginia surcharge*: $11/MWh on data center consumption in Virginia, paid by data centers only, with revenue rebated to consumers. Price effect is the residential-load-weighted retail price change from data center load. Incremental CO_2 is total-grid emissions with minus without data center load; total climate damages value total-grid emissions at the SCC. Residential electricity costs are positive for an increase. Consumer surplus $=$ $-$costs $-$ climate damages $+$ transfers. Capacity relocated is measured on the expected allocation, which is not subject to the capacity screen. Dispatch: 2024 merit order plus all EIA-860 proposed thermal units, with planned wind and solar netted from thermal demand. Location choice is evaluated with load beyond thermal capacity priced on a scarcity ramp rising linearly to $9,000/MWh over a band equal to 5% of thermal capacity; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC $185/t (Rennert et al. 2022); data center load $=$ nameplate $\times$ 0.70.

Table D.3: State data center moratoria ranked by Δ consumer surplus (2026–2030 build-out, 50% realization)

		Δ residential electricity costs ($M)		Total relocation effects	
	Capacity displaced (GW)	In-state	Out-of-state	Climate damages ($M)	Consumer surplus ($M)
Texas*	12.04	-586.1	850.2	+$1649.3	-$1913.4
Virginia*	10.35	-109.3	-117.2	+$734.8	-$508.3
California	1.56	-60.3	107.4	+$308.5	-$355.5
Louisiana	1.47	-17.6	85.6	+$159.3	-$227.2
New York*	0.93	-25.7	64.6	+$177.6	-$216.6
⋮					
Missouri	1.19	-5.2	17.0	-$168.3	+$156.5
Michigan*	1.15	-14.5	30.7	-$181.8	+$165.6
Oklahoma*	1.61	-45.7	50.2	-$197.2	+$192.7
Indiana	1.56	-16.7	42.6	-$255.7	+$229.8
Iowa	1.74	-5.3	32.9	-$280.8	+$253.2
Avg, proposed ban*	2.36	-54.3	68.6	+$114.5	-$128.7
Avg, no ban proposed	1.11	-11.7	10.2	-$0.7	+$2.2

Notes: Per-state moratorium on the announced 2026–2030 grid buildout at 50% realization (behind-the-meter excluded), using the Grid (2026–2030) preference estimates in Table 1: the state is dropped from every campus's choice set and its capacity redistributed at the price-feedback equilibrium. Capacity displaced is the baseline modeled capacity in the banned state. Residential electricity costs are positive for an increase, so an in-state saving is negative. Climate damages are the signed change in total-grid CO_2 valued at the SCC. Consumer surplus $=$ $-$in-state costs $-$ out-of-state costs $-$ climate damages. The five largest consumer-surplus decreases and increases are shown; $*$ marks the 16 states with a proposed or enacted moratorium as of September 2026 (NCSL tracker; Office of the Texas Governor 2026). Group rows are unweighted averages within each group (16 proposed-ban states, 32 other states in the contiguous U.S.). Moratorium allocations are expected allocations and are not subject to the generation capacity screen. Location choice is evaluated with load beyond thermal capacity priced on a scarcity ramp rising linearly to $9,000/MWh over a band equal to 5% of thermal capacity; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC $185/t (Rennert et al. 2022); data center load $=$ nameplate $\times$ 0.70; 2024 dispatch.

Table D.4: State data center moratoria ranked by Δ consumer surplus (2026–2030 build-out, 50% realization, augmented grid)

		Δ residential electricity costs ($M)		Total relocation effects	
	Capacity displaced (GW)	In-state	Out-of-state	Climate damages ($M)	Consumer surplus ($M)
Virginia*	9.95	-65.5	-4.7	+$1012.4	-$942.2
Texas*	15.82	-501.7	708.3	+$390.6	-$597.2
Ohio	3.41	-19.4	-8.0	+$304.5	-$277.1
Louisiana	1.31	-8.7	52.5	+$197.1	-$240.9
New York*	0.86	-21.8	39.4	+$152.2	-$169.8
⋮					
Michigan*	1.09	-9.9	18.6	-$185.6	+$176.9
Missouri	1.09	-3.5	9.6	-$183.1	+$177.0
Indiana	1.45	-11.3	24.3	-$258.8	+$245.8
Oklahoma*	1.52	-30.5	31.2	-$262.4	+$261.6
Iowa	1.57	-3.5	17.0	-$275.3	+$261.8
Avg, proposed ban*	2.51	-42.9	56.1	+$50.8	-$63.9
Avg, no ban proposed	1.03	-13.5	9.0	-$12.3	+$16.9

Notes: Per-state moratorium on the announced 2026–2030 grid buildout at 50% realization (behind-the-meter excluded), using the Grid (2026–2030) preference estimates in Table 1: the state is dropped from every campus's choice set and its capacity redistributed at the price-feedback equilibrium. Capacity displaced is the baseline modeled capacity in the banned state. Residential electricity costs are positive for an increase, so an in-state saving is negative. Climate damages are the signed change in total-grid CO_2 valued at the SCC. Consumer surplus $=$ $-$in-state costs $-$ out-of-state costs $-$ climate damages. The five largest consumer-surplus decreases and increases are shown; $*$ marks the 16 states with a proposed or enacted moratorium as of September 2026 (NCSL tracker; Office of the Texas Governor 2026). Group rows are unweighted averages within each group (16 proposed-ban states, 32 other states in the contiguous U.S.). Moratorium allocations are expected allocations and are not subject to the generation capacity screen. Dispatch: 2024 merit order plus all EIA-860 proposed thermal units, with planned wind and solar netted from thermal demand. Location choice is evaluated with load beyond thermal capacity priced on a scarcity ramp rising linearly to $9,000/MWh over a band equal to 5% of thermal capacity; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC $185/t (Rennert et al. 2022); data center load $=$ nameplate $\times$ 0.70.

Table D.5: Data center allocations and regional headroom (2026–2030 buildout, 50% realization, augmented grid)

		Unconstr.	Capacity-screened allocation	
Region	Headroom	Modeled	Random	Modeled
PJM	27.4	13.6	3.1	13.6
ERCOT	11.5	10.4	6.7	10.2
Southeast	27.8	9.6	6.5	9.6
MISO_North	16.6	5.4	8.5	5.5
SPP	17.2	3.6	9.6	3.6
NorthernGrid	4.8	3.1	4.7	3.1
MISO_South	12.0	1.9	3.4	1.9
WestConnect	8.9	1.8	4.9	1.8
CAISO	3.1	1.1	1.7	1.1
NYISO	3.5	0.6	0.7	0.6
ISONE	3.9	0.3	1.4	0.3
Total	136.8	51.2	51.2	51.2

Notes: Average data center load (GW) by dispatch region, announced 2026–2030 grid buildout at 50% realization. *Headroom*: flat data center load a region hosts while below thermal capacity for $\geq$99.9% of hours (thermal capacity minus 99.9th-percentile non-data-center thermal demand). *Unconstr. Modeled*: expected allocation at the price-feedback equilibrium, no screen. *Capacity-screened* columns place campuses whole under the generation capacity screen, averaged over 2,000 draws: Random in proportion to land area, Modeled from the equilibrium choice probabilities. Dispatch: 2024 merit order plus all EIA-860 proposed thermal units, with planned wind and solar netted from thermal demand. Data center load = nameplate $\times$ 0.70.

Table D.6: Utility-state characteristics by data center exposure

		Median sales (TWh)			Median price ($/MWh)			Data centers (MW)	
Group	*N*	Total	Res.	Non-res.	Res.	Non-res.	Median area	Stock 2024	Median addition
1. No data centers	1156	0.356	0.148	0.192	154.1	125.2	2393	0.0	–
2. Some DC, never treated	152	5.894	2.209	3.564	156.3	119.6	9020	13.6	4.0
3. Treated	50	12.169	3.901	8.023	153.1	104.6	20083	402.0	31.0

Notes: Utility-states in the local projections panel, grouped by data center exposure over 2000–2024. "No data centers" never host any data center capacity; "Some DC, never treated" host capacity but never make an annual addition of at least 100 MW, and together with the not-yet-treated form the control group; "Treated" make at least one annual addition of 100 MW or more. Sales, prices, and area are 2016 EIA-861 levels, so they describe the units rather than their treatment paths; prices are the average revenue per MWh in $2025. Area is the utility-state service-territory piece area. "Stock 2024" is data center capacity in the final panel year and "Median addition" is the typical size of a single positive annual addition, pooled across unit-years within each group; the no-data-center group has no additions by construction. All entries are medians across utility-states in the group. Campus data from Schlauch et al. (2026).

Table D.7: State data center subsidy programs

State	Year	Subsidy (nominal $M)	MW before subsidy year	Non-res. rate ($2025/MWh)	Elec. cost /MW/yr ($2025)	Subsidy /MW/yr ($2025)	Net elec. cost /MW/yr ($2025)
GA	2025	296.0	2,279	97.6	598,593	129,887	468,706
IA	2020	151.1	1,425	81.0	496,621	131,967	364,653
IL	2023	370.6	1,388	106.3	651,773	282,346	369,428
MN	2025	114.2	125	110.3	676,152	916,165	-240,013
NV	2024	140.0	1,040	97.5	597,718	138,204	459,514
OH	2025	127.4	2,389	91.2	559,324	53,320	506,005
TN	2025	103.6	796	100.3	614,997	130,210	484,787
TX	2025	1015.6	9,559	74.9	459,325	106,251	353,074
VA	2023	928.6	8,022	90.0	551,981	122,405	429,576
WA	2025	112.1	1,496	90.4	554,161	74,974	479,187

Notes: Subsidy totals are from Good Jobs First and are reported in nominal dollars for the indicated year. "MW before subsidy year" is the sum of nameplate MW capacity across all data center campuses in the state with opening year strictly before the subsidy year; campuses with unknown opening year are excluded. The non-residential retail rate is the state-level sales-weighted average non-residential electricity price in 2024, deflated to 2025 dollars using CPI-U. Electricity cost per MW per year assumes 70% utilization (i.e., $0.70 \times 8{,}760$ hours $\times$ rate). Subsidy per MW per year divides the CPI-deflated ($2025) total subsidy by MW before the subsidy year. Net electricity cost is the difference between annual electricity cost and the annual subsidy, both per MW.

Table D.8: LBNL wholesale prices on dispatch-model market-clearing prices

	Levels	Logs
Dispatch-model MCP	1.316^{*}	0.857^{***}
	(0.453)	(0.147)
95% CI	[0.31, 2.33]	[0.53, 1.18]
Region FE	Yes	Yes
Year FE	Yes	Yes
R^2	0.81	0.89
Regions	11	11
Observations	99	99

Notes: LBNL wholesale price regressed on dispatch-model MCP; 11-region region-year panel, 2016–2024 ($2025/MWh), with region and year fixed effects. Webb wild cluster bootstrap standard errors over region clusters. * $p < 0.05$, *** $p < 0.001$.

Table D.9: Focus firm summary statistics

	Pre-2015		2019-2024	
Firm	Campuses	MW	Campuses	MW
Amazon	7	1049	29	3850
Meta	3	625	10	2509
Microsoft	6	746	17	2382
Google	7	1817	12	1675
QTS	12	1209	5	810
Digital Realty	17	637	5	722
Focus firms total	52	6083	78	11948
All firms total	844	16501	281	28256

Notes: Number of distinct campuses and allocated capacity (MW) by firm and period. The six focus firms used to construct the firm Bartik instrument account for 37% of pre-2015 MW and 42% of 2019–2024 MW.

Table D.10: Firm Bartik IV using 2018–2025 retail prices

	OLS	IV
Res. price ($/MWh)	-3.76	-6.49
	[-7.55, 0.03]	[-27.58, -2.38]
Nonres. price ($/MWh)	-5.90	-7.47
	[-9.85, -1.95]	[-25.84, -4.13]
Nonres. sales (TWh)	+2.73	+3.34
	[2.26, 3.19]	[2.41, 9.18]
First-stage F (state-clust.)		7.0
Observations	369	369

Notes: OLS and firm Bartik instrumental-variables estimates of the effect of a +1 GW own data center capacity addition on the 2018–2025 change in residential and non-residential retail price ($2025/MWh) and non-residential sales (TWh). Endpoints are from EIA Form 861M (monthly): prices are revenue-weighted annual averages and sales are annualized totals, each taken over the calendar months present in both 2018 and 2025, restricting the sample to utility-states reporting in 861M in both years. Observations are utility-states in the contiguous U.S.; specifications include dispatch-region fixed effects and controls (2016 non-residential retail price, log 2016 non-residential sales, log land value), with standard errors clustered on state. OLS brackets are conventional 95% confidence intervals; IV brackets are Anderson and Rubin (1949) 95% confidence sets, robust to weak identification. The first-stage F is state-clustered. Source: EIA-861M; campus data from Schlauch et al. (2026).

Table D.11: Rotemberg-weight decomposition of the firm Bartik 2SLS estimate

Firm	Shock g (GW)	Rotemberg α	β res. price	β nonres. price	β nonres. sales
Amazon	3.85	0.576	-1.10	-1.25	2.90
Microsoft	2.38	0.168	-4.34	-1.89	2.46
QTS	0.81	0.114	0.09	-1.94	3.49
Digital Realty	0.72	0.096	0.92	-0.25	2.83
Meta	2.51	0.037	-13.80	-9.01	4.72
Google	1.68	0.010	-41.67	-37.40	9.14

Notes: Per-firm decomposition of the firm Bartik 2SLS estimate for the six focus firms. The shock g_f is the firm's national 2019–2024 capacity addition (GW); the Rotemberg weight α_f is the firm's contribution to the composite estimate (weights sum to one); β_f is the firm's just-identified IV estimate using its pre-2015 location shares as the instrument, per $+1$ GW, for each outcome (retail price, \$2025/MWh; non-residential sales, TWh). Unlike our main Bartik results, this analysis uses a pooled design, with no region FE, for computational ease. Source: EIA-861; campus data from Schlauch et al. (2026).

Table D.12: Persistence: firm-by-utility-state Poisson regression

Dependent Variable:	MW built 2019-2024					
	Amazon	Meta	Microsoft	Google	QTS	Digital Realty
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Constant	0.505***	0.615***	-0.010	0.213***	-0.877***	-2.48***
	(0.021)	(0.020)	(0.027)	(0.024)	(0.042)	(0.094)
Pre-period MW	0.009***	-0.093	0.027***	-0.150	0.009***	0.021***
	(4.18×10^{-5})	(0.673)	(0.0002)	(1.70)	(0.0001)	(0.0002)
Fit statistics						
Observations	1,359	1,359	1,359	1,359	1,359	1,359
Squared Correlation	0.526	1.34×10^{-5}	0.537	1.53×10^{-5}	0.269	0.982
Pseudo R^2	0.414	0.0004	0.424	0.0008	0.258	0.848
BIC	25,127.8	25,485.6	14,585.3	18,192.7	7,192.7	1,473.8

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Per-firm Poisson regression of firm-by-utility-state 2019–2024 MW on the firm's pre-period (pre-2015) MW in that utility-state. One column per focus firm.

Table D.13: Firm Bartik estimates excluding Dominion Virginia

	All utility-states		Excluding Dominion Virginia		
	OLS	IV	OLS	IV	IV, instrument rebuilt
Res. price ($/MWh)	+1.48	-2.81	+5.37	-8.80	-5.33
	[-3.64, 6.60]	[-15.01, 0.31]	[0.57, 10.16]	[-37.96, 6.93]	[-26.52, 10.56]
Nonres. price ($/MWh)	-3.14	-2.40	-4.41	-3.93	-1.13
	[-5.52, -0.75]	[-9.35, 0.64]	[-9.28, 0.45]	[-29.04, 9.73]	[-18.63, 13.85]
Nonres. sales (TWh)	+4.18	+3.13	+5.85	+4.56	+4.68
	[1.66, 6.71]	[2.66, 5.46]	[2.86, 8.83]	[2.16, 8.97]	[2.20, 7.58]
First-stage F (state-clust.)		7.7		20.6	22.3
Observations	1,253	1,253	1,252	1,252	1,252

Notes: OLS and firm Bartik instrumental-variables estimates of the effect of a +1 GW own data center capacity addition on the 2018–2024 change in residential retail price, non-residential retail price ($2025/MWh), and non-residential sales (TWh). The first two columns use all utility-states in the contiguous U.S.; the last three exclude Dominion Virginia (Virginia Electric & Power Co.). In the final column, the instrument is also rebuilt without Dominion Virginia's campuses: each firm's pre-2015 shares are computed over the remaining utility-states, and each firm's 2019–2024 shift excludes the capacity it built in Dominion Virginia. Specifications include dispatch-region fixed effects and controls (2016 non-residential retail price, log 2016 non-residential sales, log land value), with standard errors clustered on state. OLS brackets are conventional 95% confidence intervals; IV brackets are Anderson and Rubin (1949) 95% confidence sets, robust to weak identification. The first-stage F is state-clustered. Source: EIA-861; campus data from Schlauch et al. (2026).

Table D.14: Leave-one-firm-out firm Bartik 2SLS (residential price)

Dropped	Estimate	SE	Eff. F
None (all firms)	-2.81	2.25	7.7
Amazon	-4.82	4.71	2.0
Google	-2.22	1.73	15.4
Microsoft	-2.15	2.28	5.7
Meta	-2.41	1.26	260.2
Digital Realty	-3.06	2.40	7.1
QTS	-3.31	2.67	5.1

Notes: Firm Bartik 2SLS of the 2018–2024 change in residential retail price ($2025/MWh) on own data center capacity, per +1 GW, re-estimated dropping each firm's component from the composite instrument. Utility-state observations with dispatch-region fixed effects; standard errors clustered on state. "Eff. F" is the state-clustered first-stage F. Source: EIA-861; campus data from Schlauch et al. (2026).

Table D.15: Firm Bartik reduced form, pre- vs. post-period

	Res. price		Nonres. price		Nonres. sales	
	2018–2024	2012–2018	2018–2024	2012–2018	2018–2024	2012–2018
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Firm-Bartik instrument	-0.004*	-0.003	-0.003	-0.0006	0.004***	0.002***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.0002)
Average non-residential electricity price, 2016	-0.050	-0.109	-0.168***	-0.090	-0.0009	-0.0007
	(0.031)	(0.105)	(0.023)	(0.108)	(0.001)	(0.0006)
Log(non-residential electricity load, 2016)	1.82***	1.07*	0.876*	0.353	0.038	0.066
	(0.517)	(0.637)	(0.449)	(0.588)	(0.081)	(0.048)
Log(land value)	-1.54**	-0.081	-0.432	0.343	-0.015	-0.017
	(0.717)	(0.528)	(0.609)	(0.421)	(0.021)	(0.014)
Fixed-effects						
region	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	1,253	1,253	1,253	1,253	1,253	1,253
R^2	0.187	0.120	0.184	0.119	0.157	0.113
Within R^2	0.064	0.048	0.109	0.027	0.130	0.086

Clustered (STUSPS) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Reduced-form regressions of outcome changes on the firm Bartik with dispatch-region fixed effects and location-choice controls; utility-state observations, standard errors clustered on state. Within each outcome, the first column is the 2018–2024 change (actual) and the second the pre-boom 2012–2018 change (placebo). Outcomes: residential and non-residential retail price ($2025/MWh) and non-residential sales (TWh). Source: EIA-861; campus data from Schlauch et al. (2026).

Table D.16: Regional generation capacity, headroom, and the announced 2026–2030 build-out

	2024 grid		Planned 2026–2030		
Region	Thermal capacity (GW)	Headroom (GW avg)	Data centers (GW)	Fossil (GW)	Solar and wind (GW avg)
PJM	130.7	16.8	46.0	6.8	4.4
ERCOT	66.3	5.9	25.8	3.8	11.2
MISO_North	75.7	9.3	21.0	3.9	3.7
Southeast	126.1	18.4	17.9	7.4	3.1
NorthernGrid	25.4	3.3	10.3	0.9	3.7
SPP	47.9	7.0	9.2	9.8	1.6
MISO_South	38.7	6.4	7.1	5.1	0.8
WestConnect	31.6	4.6	7.0	2.5	5.2
CAISO	25.5	1.5	1.4	0.0	3.8
NYISO	21.4	2.6	1.1	0.0	1.6
ISONE	15.0	3.6	0.0	0.0	0.6
Total	604.3	79.3	146.8	40.3	39.7

Notes: One row per dispatch region, ordered by planned data center capacity. *Thermal capacity* is each region's total operating capacity on the 2024 dispatch supply curve, where a plant's capacity is the 99th percentile of its hourly gross load (EPA CEMS, 2024). *Headroom* is the flat data center load a region can host while remaining below thermal capacity in at least 99.9% of hours: thermal capacity minus the 99.9th percentile of 2024 hourly thermal demand, in average GW. *Data centers* is the announced 2026–2030 grid-connected pipeline (Aterio high-likelihood, net-new nameplate GW; behind-the-meter campuses excluded), with each campus assigned whole to the region holding the largest share of its utility-state's area. *Fossil* is nameplate GW of EIA-860 proposed generators (firm and speculative) with a planned in-service year through 2030; proposed battery storage is excluded. *Solar and wind* covers the same proposed generators but is reported in average GW, deflating each region's proposed nameplate by that region's 2024 annual mean capacity factor for the technology, computed from the same hourly profiles the dispatch model nets out; this assumes proposed units perform like the region's existing fleet of that technology. Nationally this converts 167.6 GW of proposed nameplate into 39.7 GW of average output. Data center capacity is nameplate; multiply by 0.70 to compare it against headroom in average load.

Table D.17: Regional thermal supply-curve capacity: 2024 baseline, 2026–2030 supply scenarios, and the announced buildout

BA region	2024 baseline (GW)	Optimistic: proposed plants, no retirements (Δ GW)	Mid-case: in-progress plants, no retirements (Δ GW)	Pessimistic: in-progress plants, planned retirements (Δ GW)	Planned grid buildout 2026–2030 (GW)
PJM	130.7	+6.6	+1.8	-7.0	45.7
Southeast	126.1	+6.4	+1.9	-3.1	18.1
MISO_North	75.7	+3.4	+0.6	-11.6	20.7
ERCOT	66.3	+3.4	+0.8	-1.1	22.3
SPP	47.9	+7.9	+1.2	-1.9	12.5
MISO_South	38.7	+4.8	+0.7	-2.7	7.0
WestConnect	31.6	+1.8	+0.2	-3.8	7.5
CAISO	25.5	+0.0	+0.0	-2.2	1.4
NorthernGrid	25.4	+0.9	+0.8	-1.8	10.0
NYISO	21.4	+0.0	+0.0	+0.0	1.1
ISONE	15.0	+0.0	+0.0	+0.0	0.0
Total	604.3	+35.3	+8.1	-35.1	146.3

Notes: Thermal (fossil) generating capacity by dispatch region. The 2024 baseline is each region's total capacity on the dispatch supply curve, where a plant's capacity is the 99th percentile of its hourly gross load (EPA CEMS, 2024). The three scenario columns report the net change in that capacity (Δ GW) under EIA-860 generator plans with a planned in-service year through 2030, where a proposed unit's capacity is its nameplate scaled by a per-technology capacity factor: *Optimistic* adds all proposed thermal generators (firm and speculative) and applies no retirements; *Mid-case* adds only in-progress units (EIA-860 status under construction or testing) and applies no retirements; *Pessimistic* adds the in-progress units and removes all scheduled retirements through 2030. The final column, the "planned grid buildout," is the announced 2026–2030 grid-connected data center pipeline (Aterio high-likelihood, net-new nameplate GW, area-weighted from utility-states to regions); it excludes announced behind-the-meter campuses and is a demand-side addition shown for comparison, not part of supply. Regions ordered by 2024 baseline capacity.

Table D.18: Regional renewable and battery-storage capacity: 2026–2030 supply scenarios and the announced buildout

BA region	Optimistic: proposed plants, no retirements (Δ GW)	Mid-case: in-progress plants, no retirements (Δ GW)	Storage (proposed) (GW)	Planned grid buildout 2026–2030 (GW)
ERCOT	+48.2	+14.1	32.3	22.3
PJM	+20.9	+6.3	1.3	45.7
WestConnect	+18.7	+4.9	10.5	7.5
MISO_North	+16.8	+8.1	1.7	20.7
Southeast	+14.6	+4.5	2.4	18.1
CAISO	+14.3	+6.7	13.4	1.4
NorthernGrid	+13.8	+1.9	9.7	10.0
NYISO	+7.8	+2.3	1.3	1.1
SPP	+5.7	+1.3	1.1	12.5
MISO_South	+4.7	+0.8	0.0	7.0
ISONE	+2.9	+2.0	0.6	0.0
Total	+168.4	+52.8	74.3	146.3

Notes: Renewable (solar, wind, hydro, geothermal) generating capacity and battery storage by dispatch region, in nameplate GW. Source: EIA-860 (2024 report year, via PUDL). The two scenario columns report the net change in renewable nameplate capacity (Δ GW) from EIA-860 generators with a planned in-service year through 2030: *Optimistic* adds all proposed renewables (firm and speculative); *Mid-case* adds only in-progress units (EIA-860 status under construction or testing). A with-retirements (pessimistic) column is omitted because scheduled renewable retirements through 2030 are negligible (≈0.2 GW nationally), leaving it indistinguishable from the mid-case. *Storage (proposed)* is all proposed battery storage (nameplate GW; pumped storage excluded). The "planned grid buildout" column is the announced 2026–2030 grid-connected data center pipeline (Aterio high-likelihood, net-new nameplate GW, area-weighted from utility-states to regions), a demand-side addition shown for comparison. Unlike the thermal table, capacities are nameplate and are not derated to operating capacity; nameplate overstates the firm-capacity contribution of variable renewables, whose annual capacity factors average roughly 25% (solar) and 35% (wind). Regions ordered by proposed renewable capacity (optimistic).

References

Alvarez, Fernando, David Argente, Joyce Chow, and Diana Van Patten. 2026. "Data Centers and Local Economies in the Age of AI: A Shift–Share Approach." National Bureau of Economic Research w35194, Cambridge, MA.

Anderson, T. W., and Herman Rubin. 1949. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *The Annals of Mathematical Statistics*, 20(1): 46–63.

Arkolakis, Costas, and Conor Walsh. 2023. "Clean Growth." https://www.nber.org/papers/w31615.

Aterio. 2026. "US Data Center Locations and Power Demand." https://www.aterio.io/datasets/lst_us_data_centers.

Balboni, Clare, and Joseph S. Shapiro. 2025. "Spatial environmental economics." In *Handbook of Regional and Urban Economics*. Vol. 6, , ed. Dave Donaldson and Stephen J. Redding, 585–652. Elsevier.

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute.

Benetton, Matteo, Giovanni Compiani, and Adair Morse. 2023. "When Cryptomining Comes to Town: High Electricity-use Spillovers to the Local Economy." National Bureau of Economic Research w31312, Cambridge, MA.

BloombergNEF. 2025. "AI and the Power Grid: Where the Rubber Meets the Road." https://about.bnef.com/insights/clean-energy/ai-and-the-power-grid-where-the-rubber-meets-the-road/.

BloombergNEF. 2026. "Six Things to Know About BNEF's New US Data Center Capacity Outlook." https://about.bnef.com/insights/data-centers/six-things-to-know-about-bnefs-new-us-data-center-capacity-outlook/.

Borenstein, Severin. 2025. "What Will Data Centers Do To Your Electric Bill?" *Energy Institute Blog*. https://energyathaas.wordpress.com/2025/09/29/what-will-data-centers-do-to-your-electric-bill/.

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

Borenstein, Severin, and James Bushnell. 2015. "The US Electricity Industry After 20 Years of Restructuring." *Annual Review of Economics*, 7(1): 437–463.

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.

Bray, Hiawatha. 2026. "Massachusetts electric bills: Should we buy power the N.H. way?" *The Boston Globe*.

Bruns, Adam. 2026. "DATA CENTERS: What's Ahead for The Global Data Center Pipeline?" *Site Selection Magazine*. https://siteselection.com/data-centers-whats-ahead-for-the-global-data-center-pipeline/.

Brülhart, Marius, Mario Jametti, and Kurt Schmidheiny. 2012. "Do agglomeration economies reduce the sensitivity of firm location to tax differentials?" *The Economic Journal*, 122(563): 1069–1093.

Castro-Vincenzi, Juanma. 2024. "Climate Hazards and Resilience in the Global Car Industry."

Cicala, Steve. 2022. "Imperfect Markets versus Imperfect Regulation in US Electricity Generation." *THE AMERICAN ECONOMIC REVIEW*, 112(2).

Darmouni, Olivier, Clemens Lehner, and Yuqi Zhang. 2026. "The Energy Cost of AI and Data Centers."

Devereux, Michael P., Rachel Griffith, and Helen Simpson. 2007. "Firm location decisions, regional grants and agglomeration externalities." *Journal of Public Economics*, 91(3): 413–435.

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor. 2025. "A Local Projections Approach to Difference-in-Differences." *Journal of Applied Econometrics*, 40(7): 741–758.

EPRI. 2024. "Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption."

Fajgelbaum, Pablo D, Eduardo Morales, Juan Carlos Suárez Serrato, and Owen Zidar. 2019. "State Taxes and Spatial Misallocation." *The Review of Economic Studies*, 86(1): 333–376.

Gagnon, Pieter, Brady Cowiestoll, and Marty Schwarz. 2023. "Cambium 2022 Scenario Descriptions and Documentation." National Renewable Energy Laboratory.

Goldmacher, Shane. 2026. "The Data Center Backlash Bursts Into the Midterms." *The New York Times*.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review*, 110(8): 2586–2624.

Gonzales, Luis E., Koichiro Ito, and Mar Reguant. 2023. "The Investment Effects of Market Integration: Evidence From Renewable Energy Expansion in Chile." *Econometrica*, 91(5): 1659–1693. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA20769.

Grant, Peter. 2024. "Exclusive | Arizona Deal Latest Sign of Booming Demand for Sites to Power AI." *Wall Street Journal*.

Guidi, Gianluca, Francesca Dominici, Nat Steinsultz, Gabriel Dance, Lucas Henneman, Henry Richardson, Edgar Castro, Falco J. Bargagli-Stoffi, and Scott Delaney. 2025. "The environmental burden of the United States' bitcoin mining boom." *Nature Communications*, 16(1): 2970.

Guimarães, Paulo, Octávio Figueirdo, and Douglas Woodward. 2003. "A Tractable Approach to the Firm Location Decision Problem." *The Review of Economics and Statistics*, 85(1): 201–204.

Hainmueller, Jens. 2012. "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies." *Political Analysis*, 20(1): 25–46.

Ham, Dasom, Owen Kay, and Catherine Hausman. 2026. "Transmission lowers US generation costs, but generator incentives are not aligned." *Proceedings of the National Academy of Sciences*, 123(9): e2524463123.

Hausman, Catherine. 2025. "Power Flows: Transmission Lines, Allocative Efficiency, and Corporate Profits." *American Economic Review*, 115(8): 2574–2615.

Hering, Garrett, and Susan Dlin. 2025. "Data center grid-power demand to rise 22% in 2025, nearly triple by 2030 | S&P Global." *S&P Global*.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2016. "Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors." *American Economic Review*, 106(12): 3700–3729.

Homeland Infrastructure Foundation-Level Data. 2023. "Electric Retail Service Territories." https://openenergyhub.ornl.gov/explore/dataset/electric-retail-service-territories/information/.

Houlihan Lokey. 2024. "Real Estate Market Update: Data Centers."

Howland, Ethan. 2026. "Data centers were 40% of PJM capacity costs in last auction: market monitor." *Utility Dive*. https://www.utilitydive.com/news/data-centers-pjm-capacity-auction/808951/.

Hsiao, Allan. 2025. "Sea Level Rise and Urban Adaptation in Jakarta." 30.

Kahn, Matthew E., and Erin T. Mansur. 2013. "Do local energy prices and regulation affect the geographic concentration of employment?" *Journal of Public Economics*, 101: 105–114.

Kay, Owen, Robert Reaser, and Reid Taylor. 2026. "Processing Power: The Effect of Data Centers on Wholesale Electricity Market." *Federal Reserve Bank of Dallas, Working Papers*, 2026(2606).

Kim, Dongjoo, Lin Dong, and Le Xie. 2026. "Flexibility-aware framework for efficient planner-initiated siting of data center." *Nature Communications*, 17(1): 6512.

Knittel, Christopher, Juan Ramon Senga, and Shen Wang. 2025. "Flexible Data Centers and the Grid: Lower Costs, Higher Emissions?" National Bureau of Economic Research w34065, Cambridge, MA.

Krishnamoorthy, Gayathri, Vincent Westfallen, and Killian McKenna. 2026. "The Role of the U.S. Electric Distribution System in Serving Data Center and Other Large Loads." *Technical Report*.

Mamkhezri, Jamal, Xiaochen Sun, and Yuting Yang. 2026. "Site and Scale: Data Centers, Grid Congestion, and Welfare." https://papers.ssrn.com/abstract=7261698.

M., Ana. 2026. "Data Center Tiers Explained: Tier I, II, III & IV (2026 Guide)." *Ingenious Build*. https://www.ingenious.build/blog-posts/data-center-tiers-explained.

Meeks, Robyn, Jacquelyn Pless, Zhiyuan Qi, and Zhenxuan Wang. 2026. "Who Pays for Growth? Evidence from Data Centers and the Grid."

Michael, Amelia, and Ben Cottier. 2026. "Total cost of ownership of a one-gigawatt AI data center." *Epoch AI*. https://epoch.ai/data-insights/ai-datacenter-cost-breakdown.

Millstein, Dev, Eric O'Shaughnessy, and Ryan Wiser. 2026. "Renewables and Wholesale Electricity Prices (ReWEP) tool." *Lawrence Berkeley National Laboratory*. https://emp.lbl.gov/renewables-and-wholesale-electricity-prices-rewep.

Moffatt, Ryan. 2024. "Community and Economic Development Annual Report and Development Retrospective."

Muller, Nicholas. 2026. "Measuring the Impact of Data Centers in the United States Economy: Monetary Damage from Air Pollution and Greenhouse Gas Emissions." National Bureau of Economic Research w35100, Cambridge, MA.

Musk, Elon. 2025. "The @xAI goal is 50 million in units of H100 equivalent-AI compute (but much better power-efficiency) online within 5 years." *Twitter*. https://x.com/elonmusk/status/1947704195844608094.

National Conference of State Legislatures. 2026. "Which States Are Banning Data Centers?" https://www.ncsl.org/fiscal/which-states-are-banning-data-centers.

Nolte, Christoph. 2020. "High-resolution land value maps reveal underestimation of conservation costs in the United States." *Proceedings of the National Academy of Sciences*, 117(47): 29577–29583.

Norris, Tyler. 2025. "The Puzzle of Low Data Center Utilization Rates." https://www.powerpolicy.net/p/the-puzzle-of-low-data-center-utilization.

O'Donnell, James, and Casey Crownhart. 2025. "We did the math on AI's energy footprint. Here's the story you haven't heard." *MIT Technology Review*. https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/.

Office of the Texas Governor. 2026. "Governor Abbott Directs TCEQ To Halt Data Center Permits." https://gov.texas.gov/news/post/governor-abbott-directs-tceq-to-halt-data-center-permits.

Pan Fang, Tommy, and Shane Greenstein. 2025. "Where the Cloud Rests: The Economic Geography of Data Centers." *Strategy Science*, 10(4): 404–420.

Pontecorvo, Emily. 2026. "Americans Now Blame Data Centers for Their Rising Power Bills -." *Heatmap News*.

Rana, Vishrut, Christian Kaps, and Serguei Netessine. 2026. "The Cost of Intelligence: Impacts of AI Data Centers on US Electricity Prices." https://papers.ssrn.com/abstract=7580459.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, David Smith, Kevin Cromar, Delavane Diaz, Frances C. Moore, Ulrich K. Müller, Richard J.

Plevin, Adrian E. Raftery, Hana Ševčíková, Hannah Sheets, James H. Stock, Tammy Tan, Mark Watson, Tony E. Wong, and David Anthoff. 2022. "Comprehensive evidence implies a higher social cost of CO2." *Nature*, 610(7933): 687–692.

RGGI. 2026. "Allowance Prices and Volumes." https://www.rggi.org/auctions/auction-results/prices-volumes.

Riot Platforms. 2023. "Riot Announces August 2023 Production and Operations Updates." https://www.riotplatforms.com/riot-announces-august-2023-production-and-operations-updates/.

Roberts, David. 2026. "For data centers, a little flexibility goes a long way." https://www.volts.wtf/p/for-data-centers-a-little-flexibility.

Saul, Josh, and Ari Natter. 2026. "Pennsylvania to Michigan, High Electric Bills Infuriate Voters Ahead of Midterms." *Bloomberg*.

Saul, Josh, Leonardo Nicoletti, Demetrios Pogkas, Dina Bass, and Naureen Malik. 2025. "AI Data Centers Are Sending Power Bills Soaring." *Bloomberg*.

Saussay, Aurélien, and Misato Sato. 2024. "The impact of energy prices on industrial investment location: Evidence from global firm level data." *Journal of Environmental Economics and Management*, 127: 102992.

Schlauch, Garrison, Coly Elhai, Simon Greenhill, Falco J Bargagli Stoffi, and Gianluca Guidi. 2026. "Mapping the Scale and Geography of U.S. Data Centers."

Seim, Katja. 2006. "An Empirical Model of Firm Entry with Endogenous Product-Type Choices." *The RAND Journal of Economics*, 37(3): 619–640.

Shehabi, Arman, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, and William Lintner. 2024. "United States Data Center Energy Usage Report." Lawrence Berkeley National Laboratory LBNL–1005775, 1372902.

Slattery, Cailin. 2025. "Bidding for Firms: Subsidy Competition in the United States." *Journal of Political Economy*, 133(8): 2563–2614.

Starner, Ron. 2023. "Data Centers: Grab Power While You Can." *Site Selection Magazine*. https://siteselection.com/grab-power-while-you-can/.

Suárez Serrato, Juan Carlos, and Owen Zidar. 2016. "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms." *American Economic Review*, 106(9): 2582–2624.

U.S. Energy Information Administration. 2023. "CE4.6.EL.ST Annual household site end-use electricity consumption in the United States by state—averages, 2020."

U.S. Energy Information Administration. 2024. "How much electricity does an American home use? - Frequently Asked Questions (FAQs)." https://www.eia.gov/tools/faqs/faq.php.

Watten, Asa, John Bistline, and Geoffrey Blanford. 2026. "Have Data Centers Raised Your Electric Bill? Causal Evidence from the United States." arXiv:2606.19777 [physics.soc-ph]. http://arxiv.org/abs/2606.19777.

Wiser, Ryan, Adria Brooks, Paul Donohoo-Vallett, Glenda Oskar, Jason Frost, Patrick Brown, and Trieu Mai. 2024. "Transmission Impact Assessment: Power Sector Infrastructure Deployment to Reduce Costs, Improve Reliability, and Lower Pollution."

Wiser, Ryan, Galen Barbose, Will Gorman, Eric O'Shaughnessy, Sydney Forrester, Paul Donohoo-Vallett, Peter Cappers, Jeffrey Deason, Ryan Hledik, and Long Lam. 2026. "Retail Electricity Price Trends and Drivers: Data Update- 2026 Edition." Lawrence Berkeley National Laboratory.