

Data Centers: Location Choice and Grid Externalities*

Coly Elhai, *Harvard University*

Yixin Zhou, *Harvard University*

September 25, 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 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, and therefore raise consumer electricity prices by less than they would if they were randomly located. 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, and Rebecca Diamond 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: 20% of all data center capacity is in Virginia, where data centers accounted for more than a quarter of Virginia’s total electricity consumption in 2023 ([EPRI 2024](#)). How does the geography of data center demand shape its effects on electricity prices and emissions?

In this paper, we use a new dataset on the data center buildout to estimate how additional demand affects prices and emissions, how these effects vary with data center geography, and the potential for policy to reduce data center harms. We find that data center location choice meaningfully affects total costs, and that data centers tend to prefer locations where they raise electricity prices by less. As a result, data centers have so far raised residential electricity costs by less than they would have if randomly allocated. Policies that change where data centers build can also change their costs. We use a model of data center location choice to show that, on average, proposed state-level data center moratoria harm consumers by \$20 million per year. However, moratoria in states with particularly emissions-intensive grids can actually improve consumer welfare by as much as \$180 million per year, mainly by reducing data centers’ total climate damages.

We deploy a dispatch model, a standard tool within the electricity literature, to calculate regional grid characteristics and their associations with data center location choice. We observe that the price and emissions impact of additional load vary substantially across regions. In some regions, the addition of an extra gigawatt of demand increases prices by \$1-\$3, and in others prices move by under \$0.30. There is also dispersion in marginal emissions intensity: the cleanest regions have marginal generation that is about half as emissions-intensive as the generation in the dirtiest regions. Correlating grid characteristics with data center growth over time (using data constructed in companion work, [Schlauch et al. 2026](#)), we find that data centers have a persistent preference for locations with shallow supply curves. This is, to our knowledge, a novel observation, and suggests that data centers are choosing to locate in places where their entry will have a smaller impact on prices. We find no consistent preference for locations with higher or lower emissions intensities.

Our dispatch model also allows us to estimate the total price and emissions effects of the data center buildout. We calculate that data centers have so far increased wholesale elec-

tricity prices by \$1–3 per MWh, relative to average wholesale prices between \$20 and \$80. Actual price effects are close to what could have been achieved under a price-minimizing allocation. We find that the climate costs of the data center buildout are large. The actual allocation led to annual emissions of about 200 megatons in 2025, about an eighth of the total from the U.S. power sector, and equivalent to \$37 billion in climate damages. Even the climate damage-minimizing placement guarantees climate damages over \$30 billion per year.

We compare estimated wholesale price effects against observed retail electricity price changes using shift-share and local projections difference-in-differences designs. Examining retail price trends lets us capture changes in the non-generation components of electricity prices, such as transmission and distribution costs. Our results suggest that data center-driven wholesale price increases have been outweighed by decreases in the other electricity cost components, meaning that the net effect of data centers on retail electricity prices has so far been negative. However, as data center load continues to grow, our dispatch model predicts that generation costs will increase nonlinearly.

To quantify how alternate policies could change the price and emissions effects of the data center buildout, we introduce a structural model of data center location choice. In this model, data center developers choose where to build based on private incentives, which are a function of local characteristics such as electricity prices, infrastructure, and land values. Their entry decisions impact electricity prices and aggregate emissions, which in turn affect consumer welfare. Because data center entry is a function of electricity costs, but also changes those costs, our solution concept is Bayesian Nash equilibrium. We embed 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 price elasticities between -0.2 and -1.1. We also find strong preferences for utilities with large non-residential electricity sales, utilities where firms have already built campuses, and locations with better-interconnected transmission networks.

Using our estimated parameters, we explore a set of location choice counterfactuals. 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. In implementing the latter, we find that the existing grid lacks the generation

capacity to support the full, planned data center buildout for 2026–2030. Our forward counterfactuals therefore assume that only 50% of planned data center capacity is actually realized by 2030, bringing us in line with data center capacity projections from other sources. In some versions of the forward counterfactuals, we also augment the grid with generation proposed to come online by 2030.

We find that by changing the geography of development, state-level data center bans can substantially shift the total price and emissions effects of data center entry. Such policies are popular, and have been proposed in 15 states as of July 2026 ([National Conference of State Legislatures 2026](#)). These bans can decrease electricity prices in the implementing state, but these savings may come at a cost because bans can force data centers to locate in places with less advantageous grid characteristics. Due to Virginia’s popularity as a data center destination, a ban in that state produces the most dramatic effects. A ban on new development in place 2021–2025 could have saved Virginia residents \$27 million in electricity costs per year, but also would have increased electricity bills outside the state by a total of \$38 million and increased total climate damages by \$506 million per year. However, other bans can actually increase welfare: a ban in Michigan would have saved its residents \$4.8 million per year and, while raising costs in other states by \$18 million, would have cut CO₂ emissions by about 1 megaton per year—\$194 million in avoided climate damages.

This project contributes to the well-established literature on the geography of economic activity. Although the most common frameworks for approaching spatial linkages between locations focus on trade linkages, other work has studied how commuting, migration, production, and information can tie the economic fates of different locations ([Allen and Arkolakis 2023](#)). Like [Arkolakis and Walsh \(2023\)](#), we instead study how the electric grid interacts with economic development across space. Ours is the first project we know of to embed a dispatch model into a model of firm location choice, allowing for feedback between the two. Like [Castro-Vincenzi \(2024\)](#), we are able to use our framework to speak to the interaction between firm location choice and the climate, though in our case we study the effect of firm location choice on climate change rather than the reverse.

We also contribute to the emerging literature on data centers and their impact on society. [Guidi et al. \(2025\)](#) and [Muller \(2026\)](#) both provide estimates of the emissions costs of data centers; we provide our own, using our dispatch model to ground our definition of marginal

electricity generation. Like [Benetton, Compiani and Morse \(2023\)](#), which focuses on cryptomines in upstate New York, we study the impact of new data center load on electricity prices, but with a broader geographic and sectoral scope because our work encompasses all of the data centers in the U.S. Several contemporaneous papers estimate price impacts from data center additions using a dispatch model ([Kay, Reaser and Taylor 2026](#); [Darmouni, Lehner and Zhang 2026](#)) and a range of identification strategies ([Meeks et al. 2026](#); [Watten, Bistline and Blanford 2026](#)). Our work pairs a dispatch model with causal identification and a model of location choice, allowing us to quantify both price impacts and the importance of location selection, compare dispatch model results with realized price data, and model policy counterfactuals.

In this project, we focus primarily on the impact of data centers on electricity generation costs, without modeling potential changes in transmission and distribution costs. We make this choice both because it is challenging to model how these costs evolve in response to load growth and because the timeline of these investments tends to be longer (on the order of decades) than the timeline for data center entry. Our finding that retail price effects of data center entry to date have been negative, along with work by other researchers ([Wiser et al. 2025](#); [Watten, Bistline and Blanford 2026](#)), suggests that load growth from data centers can decrease the transmission and distribution costs borne by other customers. However, these effects are highly dependent on local grid characteristics and the investment decisions made by utilities, making it difficult to extrapolate from these findings into the future.

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. In Section 5, we explore how our dispatch model price estimates align with those from several causal inference techniques. Next, in Section 6, we present 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 7. Section 8 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 2,743 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 duplicate matches based on utility size and type, but are still left with duplicate matches for about 20% of campuses.¹ For these duplicates, we allocate campus 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

¹For campuses matched with multiple utilities, we drop candidate matches 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 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 \$2025.²

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 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

²We 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 LMP data into gridcell-level average annual prices, average prices across the 100 highest-price hours in a year, and annual 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 [B.6](#)), where competitive wholesale markets do not exist.³

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 [B.2a](#)) and nominal terms (Figure [B.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

³We 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 toy 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 price changes 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



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 and MEE' mark the marginal emissions externality (the vertical gap between S^{soc} and S^{priv}) at D and $D + K$, respectively. 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-

capacity. This has been particularly salient in PJM, the RTO serving much of the Mid-Atlantic region, where capacity market costs have set records since 2024.⁴ However, these price impacts are the combined result of PJM’s particular capacity market mechanisms and aggressive projections of data center growth, 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 (T&D) network serving both data centers and other customers, may be spread out across multiple regions and customer classes. Since T&D 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* T&D 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](#)). Second, additional data centers on the grid might induce entry of new generation. Again, this entry might take years to be realized, due in part to long interconnection queues. The net effect on resident utility bills depends on local rate regulation and whether new generators come with upfront capital costs that outweigh efficiency benefits relative to the existing generator stock, given pre-existing usage patterns. Although modeling capacity market and infrastructure price pathways is out of scope for this project, we conduct several retrospective analyses of data center effects on retail electricity prices that capture the net effect across all short-term price pathways.

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.⁵ This growth has been geographically concentrated. Just five states

⁴The 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/kWh in 2025 relative to 2024 ([Wiser et al. 2026](#)).

⁵The average U.S. household consumed 10,791 kilowatthours (kWh) of electricity in 2022 ([U.S. Energy Information Administration 2024](#)), suggesting that 800,000 households together have a power draw of about

account for 58% of all data center capacity to date, and Texas and Virginia alone account for 41% (Figure 2a). Data center growth is large relative to baseline demand. These same 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



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. Both panels span 1990–2025.

1 GW. We assume data center power utilization at 70% of nameplate capacity.

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.⁶ Cryptomining data centers are even less concerned with latency. Since they do not need to deliver large amounts of

⁶On 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.⁷

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

⁷Riot 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.

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 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. Electricity is not only the limiting factor in where data centers can be built, but also these facilities’ largest ongoing expense. One study estimates electricity to account 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 1d](#) or [1c](#). 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 C.9 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. With these supply curves, we are able to calculate marginal slopes and marginal emissions at different levels of total demand.

4.1 Setup

We model dispatch at the plant-fuel group level.⁸ 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.⁹ 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:

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

where g indexes plants and t indexes time. Heat rate captures the efficiency of each plant (which we measure in units of MMBtu/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

⁸In practice, dispatch is actually done at the generator level. The typical plant has around 3 generators.

⁹In future, we plan to add dispatchable hydropower plants to this set.

for CO₂, CSAPR for SO₂ and NO_x).¹⁰ 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.¹¹

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 increasing $mc_{g,t}$ 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.¹²

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:

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

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:

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

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

¹⁰We 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.

¹¹We 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.

¹²This 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.

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 B.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. 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.


4.2 Supply curves

Regions vary significantly in the level, shape, and composition of their supply curves. Figure 4 overlays the supply curves for all 11 regions. CAISO, ISONE, 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 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. The ordering of regions by marginal supply curve slope is relatively stable over time, though fuel price changes have different impacts by region (Figure B.4).

Marginal emissions over the supply curve vary non-monotonically (Figure B.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 B.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 the mid-high cost range in PJM, and is not used at all in CAISO (Figure B.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 B.5). However, marginal emissions are consistently highest in coal-dependent regions like SPP and MISO-North and lowest in gas-dominant regions (ISONE, CAISO, NYISO).

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 B.13). Regressing measured on modeled prices yields a coefficient near one (Table C.10). 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 toy 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 gigawatt 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



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 tCO₂/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.¹³ 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:

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

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 cam-

¹³The 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).

puses 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.¹⁴

Trends in cumulative data center capacity (Figure B.14) reveal the same pattern of geographic concentration we noted at the state level (Figure 2) and highlight where this concentration 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)/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 \$2.99/MWh or as little as \$1.46/MWh (Figure 6b). Scaling by total residential electricity consumption in 2025 and assuming full cost pass-through, these figures represent either \$2 billion or \$4.2 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 \$0.34–0.42/MWh per GW of nameplate data center capacity.¹⁵ Price effects vary by region (Figure B.15), but are on the order of a few dollars per megawatt-hour even in regions like ERCOT and PJM that have substantial data center capacity.¹⁶ The national average of the data center price effect is smaller than the median under area-weighted randomization (\$1.78), consistent with data centers preferring shallow-slope locations.

¹⁴We 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.

¹⁵We 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.

¹⁶Although 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




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 campuses are 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 CO₂ emissions from data centers equaled 202.2 Mt, equivalent to \$37.4 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 248 Mt and climate damages of \$46 billion in 2025. 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 B.14), the two regions with the highest marginal emissions intensity in 2025 (Figure B.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 80% or increased climate damages by 23% 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 14% less CO₂, such an allocation would have come at the cost of 40% larger electricity price effects (Figure B.23).

Figure 7: Cumulative data center emissions effects



Notes: CO₂ 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 campuses are placed to minimize/maximize total CO₂ 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 Estimating net effects of data centers on prices

Our dispatch model describes how data center growth impacts 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. We use realized price data and two different identification strategies (shift-share and local projections difference-in-differences) to estimate the net effect of data centers on retail prices. This allows us to test for the relative importance of the wholesale price channel, which we do model, versus the price channels that we do not model.

In doing so, we must account for several features of the context. First, as demonstrated

in Section 4.3, data centers are choosing to locate in places with particular grid characteristics. Since grid characteristics also determine the size and direction of price effects, it is difficult to define an appropriate control group. Second, in defining controls we must also be careful to account for spatial spillovers. Our dispatch model implies that, absent transmission congestion, utilities whose neighbors receive data centers should also experience data center price effects.

We take two approaches to estimating the effects of data centers on observed prices at the utility-state level. First, we construct a firm Bartik instrument for data center expansion that exploits persistence in firms' location preferences (Section 5.1). We next make use of the timing of campus openings to estimate local projections difference-in-differences, with which we can trace out the dynamic effects of data center construction on prices (Section 5.2).

In light of the spatial spillovers concern, 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.

5.1 Bartik instrument

Several large data center companies show a preference for developing more data centers in utilities where they already have a presence (Figure B.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:

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

where s_{fi}^{pre} is firm f 's share of pre-2015 MW located in utility-state i and $\tilde{G}_{f,-i}$ 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 43% of MW built 2019–2024 (Table C.11).

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 (50% 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:

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

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 identifying assumption of this approach is that pre-2015 firm footprints are uncorrelated with price shocks in the 2019–2024 period, conditional on controls. We believe this assumption is reasonable. Wiser et al. (2026) highlights load growth, wildfires, storms, the expansion of Renewable Portfolio Standards, and growth in behind-the-meter solar as drivers of recent price growth. None of these factors is obviously correlated with early-period data center location choice except load growth. To the extent non-data center load growth is a confounder, we directly look for pre-period price trends in our placebo tests (Table C.16) and find none.

The results are presented in Table 1. IV confidence intervals are wide, but able to dis-

tinguish a negative effect of data center capacity on non-residential retail prices. The estimated effect on residential retail prices is negative but marginally insignificant. Both point estimates suggest that data center entry has historically reduced the non-generation costs of electricity. This does not contradict our dispatch model results, but rather implies that the cost channels that are not in the model attenuate the data center-induced generation cost increases predicted by the model. Indeed, these other cost effects appear to be an order of magnitude larger than the wholesale price effect, which we estimate at around \$0.34–0.42/MWh per GW of additional data center capacity.

The non-residential sales estimate serves as a robustness check: we expect new data center capacity to mechanically increase non-residential electricity sales in a given area. A 1 GW data center that is operating full-time at 70% capacity should consume 6.1 TWh per year, which is larger than the point estimate of 2.37 TWh but within the confidence interval we estimate. A smaller load effect could be explained by a campus slowly ramping up to full capacity usage over time, as is expected for multi-facility campuses (which usually open constituent facilities in phases).

As another robustness check, we run the same analysis on the subset of utility-states that report 2025 price data through EIA-861M (Table C.12). These utilities are much larger than the median. Our results are larger in magnitude, share the same signs as our main analysis, and are all significant at the 5% level.

Table 1: Firm Bartik estimates: data center capacity, retail prices, and load

	OLS	IV
Res. price (\$/MWh)	+1.83 [-4.34, 7.99]	-2.59 [-14.36, 0.43]
Nonres. price (\$/MWh)	-0.99 [-5.73, 3.74]	-3.49 [-12.72, -1.02]
Nonres. sales (TWh)	+1.83 [1.39, 2.28]	+2.37 [1.74, 6.29]
First-stage F (state-clust.)		7.3
Observations	1,247	1,247

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 continental utility-states; 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.

Following [Goldsmith-Pinkham, Sorkin and Swift \(2020\)](#), we decompose the Bartik 2SLS estimates into per-firm just-identified estimates and calculate Rotemberg weights to determine which firms are driving our results (Table [C.13](#)). We find that Amazon carries by far the largest weight, followed by Microsoft, QTS, then Digital Realty. Amazon is the firm with the largest post-period footprint (Table [C.11](#)), while all four are firms with strong location choice persistence (Table [C.14](#)). Reassuringly, β estimates from the per-firm just-identified regressions match the main Bartik results. Coefficients are uniformly positive for non-residential sales, and five of the six β values for residential price and non-residential price are negative. Furthermore, leave-one-firm-out firm Bartik 2SLS regressions yield consistently negative point estimates for the residential price β (Table [C.15](#)).

We run placebo specifications (Table [C.16](#)) to test for instrument validity. The Bartik instrument passes the price placebo tests: it is uncorrelated with residential and non-residential price changes in the pre-period. It fails the non-residential sales placebo, as expected since location persistence suggests that the same firms are also building new capacity during the placebo period.

5.2 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. To reduce selection issues, 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 C.8). We identify the causal 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 ≥ 100 MW addition. We include only events that occurred in or after 2001, leaving us with 57 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.¹⁷

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

$$y_{u,t_0+h} - y_{u,t_0-1} = \mu_e^h + \beta^h \text{dose}_e \cdot \mathbf{1}\{u = i\} + \varepsilon_{e,u,h}, \quad u \in \{i\} \cup C_e \quad (4)$$

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, we show results with and without entropy balancing ([Hainmueller 2012](#)). We balance unit-years on deregulated market status as well as non-residential price and $\log(1 + \text{non-residential sales})$ trends in t_{-3} through t_{-1} .


Across-region results are in Figure 8. Both unconditional and entropy-balanced estimates indicate negative effects of data center entry on residential and non-residential prices

¹⁷The median never-treated unit has just 10 MW of total data center capacity by period end, small enough that we would not expect significant price effects.

in at least the first three years of the data center’s life. In the entropy-balanced specification, these negative coefficients are significant in years zero through two for residential and non-residential prices. The estimated magnitudes of the negative price coefficients are large but noisy, ranging from roughly 9 to 20 \$/MWh in the first few years. For comparison, retail prices range from around \$120 to \$300 and our dispatch model predicts wholesale price increases of just \$0.34–0.42/MWh per GW, shown as the orange line in the price panels of Figure 8. Our dispatch model predicts that within-region estimates should be \$0.34–0.42/MWh per GW below the across-region values. Our estimation is not precise enough to detect this level of difference, but the direction and magnitudes of the within-region results (Figure B.19) are similar to those of the across-region specifications. Price coefficients rebound into the positive range around year four, though there are many fewer usable events with observations at this horizon, confounding interpretation.

The non-residential sales regressions again provide us with a test of our design. We estimate positive coefficients starting in year zero for both the unconditional and matched series. Coefficients are significant from year zero onward in the matched series and from year 1 onward in the unconditional series. Again, the magnitudes on the sales coefficients are of the expected sign but smaller than the 6.1 TWh we would expect at 70% capacity utilization (plotted in orange). One explanation for this lower sales coefficient is phased entry. The sales coefficients increase through at least the first few years after data center entry. Another explanation is that data center capacity utilization rates have changed over time, as lower-capacity uses give way to relatively high-utilization AI applications (Shehabi et al. 2024). Our sales effect estimates might be dampened relative to the 70% benchmark because many of our treatment events happen earlier in the period.

Figure 8: 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). An event is the first year a utility-state’s annual data center capacity addition reaches 100 MW; 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 ≥ 100 MW addition through t_0+5 . Series A is unweighted; series B reweights controls to match treated units on deregulated-market status and on the changes in non-residential price and $\log(1 + \text{non-residential sales})$ from $[t_0-3, t_0-1]$. The orange horizontal line is the dispatch-model benchmark: the model’s implied *wholesale* price effect of +1 GW of data center capacity (\$0.34–0.42/MWh) in the price panels and its implied electricity demand in the sales panel. Bands are 95% confidence intervals from state-clustered standard errors.

We can use the same design to analyze data center entry effects on wholesale prices. From LBNL, we have gridcell-year-level data on average wholesale prices, prices during the most expensive 100 hours, and negative price frequency (share of hours). We assign each data center campus to every gridcell within 50 km of the campus, then calculate gridcell-level cumulative data center capacity over our data period. As in the utility-state-level analysis, we define a treatment event as the first year with data center capacity additions of at least 100 MW, and the dose as the addition that triggers the event. We plot unconditional and entropy-balanced results, where the latter matches on pre-event trends in average wholesale prices and log population density. We cluster errors at the region level.

Results indicate serious selection issues, even balancing on price trends, making the analysis inconclusive (Figure B.20). Average price and top-100 price results show significant pre-trends in the across-region analysis. Selection issues may be more severe in these specifications because the geographic units are smaller and our controls less targeted. Within-

region results have less severe pre-trends, and indicate some increase in wholesale prices after data center entry (consistent with local congestion effects). However, within-region estimates are insignificant at nearly all time horizons, meaning they include the zero value predicted by our dispatch model.

5.3 Takeaways

Both the Bartik and LP-DiD approaches suggest negative retail price impacts from data centers. The magnitudes and significance of these estimated price impacts vary across specifications, but in all cases, point estimates are at least a few times larger than the wholesale price increases we estimate using our dispatch model. These negative effects are in line with other research finding that data centers have modestly reduced retail rates ([Watten, Bistline and Blanford 2026](#)), and with utility announcements of rate decreases made possible by the revenues from data centers ([Georgia Power 2026](#); [Morales 2026](#); [Penrod 2026](#)). Price decreases are likely the result of utilities spreading the fixed costs of existing infrastructure across a larger rate base.

Importantly, both of our designs assume that effects scale linearly with dose and are homogeneous across regions. This allows us to report effects per GW of data center addition, but also limits our ability to extrapolate from these estimates. In practice, we expect the price effects of data centers to be heterogeneous and nonlinear, depending on the shape of supply curves, baseline demand, available infrastructure capacity, and utility regulations (see Section 3.1). As a result, though our estimates capture the average of data centers' *past* effects on prices, they do not imply that data centers will decrease prices moving forward. As data center growth increases in speed and scale, campuses may no longer be able to find hosts able to serve them without building additional, costly infrastructure, which would raise other users' rates unless data center companies are contractually bound to pay for the upgrades. Furthermore, generation cost effects may increase as data center load pushes grids further up their supply curves.

In contrast, our dispatch model is able to capture these non-linearities, making it well-suited to estimating future price effects. In the remaining parts of this paper, we build a model of data center location choice and estimate counterfactuals using the resulting pa-

rameters, assuming that prices evolve as predicted by the dispatch model. Our results demonstrate how generation costs could change as data center load increases going forward, and how these costs would vary if policy altered where data centers are built.

6 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 location choice holding entry fixed, equivalent to assuming that firms have inelastic demand for data center compute. Statements from major tech firms with specific compute targets and timelines (e.g., [Musk 2025](#)) are in harmony with this. In a future extension, we plan to model how a range of entry elasticities would change our results.

Second, we model each campus as independently located, rather than one part of a company’s portfolio. Modeling within-firm diversification strategies, 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.

Finally, we model entry as a static decision. Entrants decide where to locate given contemporaneous information (including the choice strategies of same-period peers) but not how future entrants might change the costs they face. This is a less limiting 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). In a future extension, we will quantify the gap between the static version of the location choice problem and a full dynamic version.

6.1 Firm problem

Similar to other work in the firm location choice literature (Guimarães, Figueirido and Woodward 2003; Devereux, Griffith and Simpson 2007; Brühlhart, Jametti and Schmidheiny 2012), we use a conditional logit model. 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 $\tilde{p}_{i,t}(\mathbf{L}_t)$:

$$\tilde{p}_{i,t}(\mathbf{L}_t) = v \cdot \sum_{m=1}^{120} \delta^m p_{i,t}(\mathbf{L}_t) \quad (5)$$

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

¹⁸For 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).

¹⁹Estimates 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.

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

$$V_{ci} = \beta_0 K_c \tilde{p}_{i,t}(\mathbf{L}_t) + \beta' \mathbf{X}_{i,t} + \rho H_{ni,t} \quad (6)$$

$$U_{ci} = V_{ci} + \varepsilon_{ci} \quad (7)$$

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.

Because of our error distribution assumption, we can express the choice probability of campus c in location i and period t to be:

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

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 $\mathcal{C}_t = \{c \text{ s.t. } t(c) = t\}$ 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 $\mathbf{Pr}_t^*$:

$$Pr_{ci}^* = \frac{\exp(\beta_0 K_c \tilde{p}_{i,t}(\mathbf{Pr}_t^{*'} \mathbf{K}_t) + \beta' \mathbf{X}_{i,t} + \rho H_{ni,t})}{\sum_{i' \in \mathcal{I}} \exp(\beta_0 K_c \tilde{p}_{i',t}(\mathbf{Pr}_t^{*'} \mathbf{K}_t) + \beta' \mathbf{X}_{i',t} + \rho H_{ni',t})} \quad (9)$$

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

²⁰For 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.

have unique location preferences. Let $\tau(c) \in \mathcal{T} = \{\text{hyperscale, colocation, crypto}\}$ denote campus c 's exogenous type. We then write utility with type-specific preference parameters:

$$V_{ci} = \beta_0^{\tau(c)} K_c \tilde{p}_{i,t}(\mathbf{L}_t) + \beta^{\tau(c)'} \mathbf{X}_{i,t} + \rho^{\tau(c)} H_{ni,t} \quad (10)$$

Again, this yields 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.

6.2 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 $i = 1, \dots, \mathcal{I}$. 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:

$$CS_t = \sum_{i \in \mathcal{I}} \left(\underbrace{[\bar{p} - p_{i,t}(\mathbf{L}_t)] Q_{i,t}}_{\text{Electricity surplus}} \right) - \underbrace{\mathcal{D}(\mathbf{L}_t)}_{\text{Climate damages}} + \underbrace{G_t}_{\text{Government transfers}} \quad (11)$$

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 :

$$\frac{\Delta CS_t}{\Delta L_{i,t}} = \sum_{j \in \mathcal{I}} \left(-\frac{\Delta p_{j,t}(\mathbf{L}_t)}{\Delta L_{i,t}} Q_{j,t} \right) - \frac{\Delta \mathcal{D}(\mathbf{L}_t)}{\Delta L_{i,t}} + \frac{\Delta G_t}{\Delta L_{i,t}} \quad (12)$$

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

6.3 Model estimation

To estimate the model presented in Section 6, 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 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 util-

²¹Aterio 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 exclude expansions of already-opened campuses.

²²Since 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.

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

²⁴We measure the mean degree of both high-voltage (≥ 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.

ity 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 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 A for more details. We then add $\Delta MCP_{i,t}$ to the reference year non-residential retail rate in

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

²⁶Our 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. n.d.](#)).

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 this model can equivalently be estimated using PPML with campus fixed effects, as described in [Guimarães, Figueirido 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 under 1% in the realized period, and by around 4% in the future period.

6.4 Results

Our location choice parameter estimates are displayed in Table 2. 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 exclusively by omitted variable bias. Coefficient magnitudes by type line up with observed concentration in low-price areas (Figure B.21b), and imply price elasticities of -1.1 for hyperscale facilities and -0.2 for colocation and crypto at the median size for each type of campus (164 MW for hyperscale, 19 MW for colocation, 56 MW for crypto). These are in the same range as estimates from [Kahn and Mansur \(2013\)](#) on the electricity price elasticities of electricity-intensive manufacturing industries.

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

²⁷Due 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.

prefer to build in locations where they already have some presence. The pooled estimate suggests that firms are 5.4 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 grid campuses and insignificant for planned behind-the-meter campuses. This pattern is consistent with industry reporting saying that developers have moved toward secondary and tertiary markets as the established development hubs run out of spare capacity ([Starner 2023](#)). It is also consistent with our observed concentration of data centers in high-load utilities decreasing over time (Figure [B.22a](#)). The degree of the high-voltage grid, which captures interconnect- edness, is positively correlated with entry, with the largest coefficient on hyperscale cam- puses. Hyperscale facilities are the largest type on average, and therefore the most likely to want to connect directly to very high voltage lines. Mid-voltage degree is also positively correlated with entry.

Table 2: Campus location choice estimates: realized and announced buildouts

	2016–2025				2026–2030	
	Pooled	Hyperscale	Colocation	Crypto	Grid	BTM
β_0 : price $\times$ size ($K_c \bar{p}$, std.)	−0.975 (0.262)	−0.955 (0.398)	−1.373 (0.472)	−0.627 (0.242)	−0.838 (0.281)	−0.041 (0.091)
Log 2010 non-residential sales (std.)	+0.929 (0.156)	+1.136 (0.258)	+1.131 (0.261)	+0.601 (0.128)	+0.236 (0.152)	+0.159 (0.188)
Log utility-state area (std.)	+0.833 (0.239)	+0.504 (0.372)	+0.940 (0.295)	+1.137 (0.465)	+1.389 (0.186)	+1.545 (0.251)
Pre-2010 campus count (std.)	+0.128 (0.023)	+0.110 (0.057)	+0.143 (0.031)	+0.086 (0.066)	+0.151 (0.037)	−0.078 (0.065)
Log land value (std.)	+0.118 (0.196)	+0.048 (0.348)	+0.448 (0.148)	−0.985 (0.296)	+0.507 (0.169)	−0.634 (0.310)
Deregulated state	+0.783 (0.280)	+0.967 (0.629)	+0.768 (0.231)	+0.685 (0.477)	+0.522 (0.277)	+1.798 (0.459)
High-voltage (≥ 220 kV) mean degree (std.)	+0.825 (0.172)	+1.402 (0.250)	+0.651 (0.203)	+0.366 (0.344)	+0.691 (0.173)	+0.302 (0.337)
Mid-voltage (69–161 kV) mean degree (std.)	+0.586 (0.088)	+0.528 (0.174)	+0.700 (0.091)	+0.750 (0.216)	+0.201 (0.268)	+0.868 (0.175)
ρ : firm incumbency H	+1.689 (0.449)	+1.962 (0.940)	+1.042 (0.421)	+2.609 (0.485)	+1.392 (0.261)	+2.369 (0.614)
Campuses (choosers)	413	122	201	90	462	65
of which $H = 1$ at chosen location	95	49	38	8	102	8

Notes: Campus-level conditional logit over the continental U.S. utility-state choice set; utility-states missing reference-year price or sales 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 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.

7 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 scenar-

ios use 2019 region-level supply curves and baseline demand, then add 2021–2025 data center entrants. The forward-looking scenarios use 2024 supply curves and demand and add 2026–2030 data center entrants.

We model four different types of data center allocation. First, we model data center placement according to the location preferences we estimate in Section 6.4. 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/tCO₂ (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 6.2, 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 CO₂. 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/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.

7.1 Backward-looking reallocation

Since data centers are more sensitive to local characteristics like total sales and firm incumbency than they are to price, price-based policies do not induce enough data center relocation to shift outcomes much. A carbon tax reduces total climate damages by \$70 billion per year (Table 3), 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. Allowing reallocation in response to the policy slightly increases data centers' impact on residential electricity costs while decreasing their climate damages, implying that data centers are relocating toward places with steeper supply curves but cleaner marginal emissions. The Virginia electricity surcharge has negligible effects on either prices or emissions, and causes just 620 MW (about 10% of Virginia's baseline total) to relocate. This \$11 surcharge, which amounts to about an 11% price increase per MWh for data centers, thus does not directly lead to significant electricity price decreases for Virginia ratepayers, but does yield \$370 million in revenues at the cost of a 10% decrease in in-state data center capacity.

Table 3: Counterfactual data center allocations (2021–2025 buildout)

Outcome	Baseline	Carbon tax		Virginia surcharge	
		Locations fixed	Firms relocate	Locations fixed	Firms relocate
Price effect (\$/MWh)	0.72	5.64	5.69	0.72	0.72
Incremental CO ₂ (Mt)	115.54	115.97	107.91	115.54	115.72
Total climate damages (\$B/yr)	\$303.0	\$232.4	\$230.9	\$303.0	\$303.0
Δ residential electricity costs (\$B/yr)	\$1.01	\$7.86	\$7.93	\$1.01	\$1.01
Δ climate damages (\$B/yr)	\$21.37	\$21.45	\$19.96	\$21.37	\$21.41
Δ government transfers (\$B/yr)	\$0.00	\$21.45	\$19.96	\$0.41	\$0.37
Δ consumer surplus (\$B/yr)	-\$22.38	-\$7.86	-\$7.93	-\$21.96	-\$22.04
Capacity in Virginia (GW)	6.12	6.12	5.71	6.12	5.53
Capacity relocated vs. baseline (GW)	0.00	0.00	4.05	0.00	0.62

Notes: Data center campuses built 2021–2025 reallocated across utility-states using the type-specific preference estimates in Table 2. *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/tCO₂ 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 CO₂ 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 at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC \$185/t (2025\$); data center load = nameplate $\times$ 0.70; 2019 dispatch.

In our moratorium counterfactual results (Table 4), we find that states that ban data centers can either impose significant costs on their neighbors or increase total social welfare, 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 decreased total (in-state plus out-of-state) residential electricity costs by \$8 million, but increased total climate damages by \$424 million per year. These large climate damages make California’s ban the most costly per MW of displaced capacity. Virginia, on the other hand, is part of a region (PJM) with a relatively flat supply curve. While a ban there would have decreased annual in-state electricity costs by \$27 million, it would have *increased* total electricity costs by raising out-of-state costs by \$38 million. This is despite the ban’s creating cost-savings in the other states in PJM (Figure 9). A Virginia ban would also have increased

total climate damages 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 been a net boon for consumers 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 \$194 million while slightly increasing total electricity costs.


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

	Capacity displaced (GW)	Δ residential electricity costs (\$M)		Total relocation effects	
		In-state	Out-of-state	Climate damages (\$M)	Consumer surplus (\$M)
Virginia*	6.17	-27.1	37.7	+\$506.0	-\$516.6
California	1.65	-63.7	56.0	+\$423.9	-\$416.2
Texas	3.38	-75.7	105.8	+\$126.1	-\$156.2
New York*	0.44	-22.4	16.2	+\$108.9	-\$102.6
Ohio	1.40	-5.0	5.6	+\$89.6	-\$90.1
:					
Missouri	0.42	-1.0	5.8	-\$70.5	+\$65.7
Arizona	0.47	-12.1	9.8	-\$80.9	+\$83.1
Iowa	0.40	-0.6	6.3	-\$89.9	+\$84.1
Indiana	0.49	-2.3	9.5	-\$109.5	+\$102.3
Michigan*	0.89	-4.8	18.0	-\$194.2	+\$181.1
Avg, proposed ban*	0.81	-5.4	6.6	+\$18.8	-\$20.0
Avg, no ban proposed	0.55	-6.6	7.7	-\$0.9	-\$0.1

Notes: Per-state moratorium on the modeled 2021–2025 buildout, using the type-specific preference estimates in Table 2: 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 CO₂ 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 states with a proposed moratorium ([NCSL tracker](#), July 2026); group rows are unweighted averages over all continental states. Location choice is evaluated with load beyond thermal capacity priced at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC \$185/t (2025\$); data center load = nameplate $\times$ 0.70; 2019 dispatch.

Figure 9: Change in residential electricity costs by state under a California or Virginia data center moratorium (2021–2025 buildout)



Notes: Change in annual residential electricity costs (\$M per year) in each continental state when the out-lined 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; the Texas total includes households served by competitive retailers, which have no franchised territory. Warm shading is a cost increase, cool a saving; the two panels share one color scale. Climate damages are not included, so the social cost of carbon does not enter. Location choice is evaluated with load beyond thermal capacity priced at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. Data center load = nameplate $\times$ 0.70; 2019 dispatch.

7.2 Forward-looking reallocation

To project forward how future data center growth will impact prices and emissions, we model the addition of data centers planned to enter 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 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, ERCOT’s planned data center capacity is four times larger than the headroom available on its grid in 2024 (Table C.17).

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 on-

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

line 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 heat rates 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. We find that price effects scale superlinearly with buildout. The 2026–2030 buildout increases average prices by \$2.81/MWh, 3.9 times larger than the price effect of the 2021–2025 data center additions (Table C.3). Notably, on the augmented grid, data center location preferences no longer protect against large price increases. Modeled preferences yield price effects that are larger than those induced by random allocations. Although emissions scale linearly with data center capacity, the magnitude of these emissions is still huge. On the 2024 grid, the climate damage-minimizing allocation leads to more than 250 megatons of CO₂ emissions per year, equivalent to nearly \$50 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, and also changes which locations are most costly. Under the augmented version of the 2024 grid, data center price effects are 16% smaller than price effects estimated using the baseline 2024 grid (Figure B.24). Incremental CO₂ emissions are 10% smaller (Figure B.25). When new generation is added (Figure B.26), marginal emissions also fall: ERCOT and SPP respectively experience 21% and 15% drops in their marginal emissions rates. The average proposed moratorium actually becomes welfare-positive under the augmented 2024 grid, largely because a ban in Virginia relocates data centers to places like ERCOT and SPP, where marginal emissions fall more than they do in PJM (Tables C.5 and C.6).

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.

²⁹Within 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.

8 Conclusion

In this paper, we develop a theoretical framework for understanding how location choice shapes the effects that data centers have on electricity prices and emissions. We build a dispatch model to show that data centers have only modestly increased electricity generation costs, and that price effects are close to what could have been achieved under price-minimizing placement. Empirical analysis on the effect of data center growth on retail prices shows that these cost increases have so far been outweighed by the consumer cost savings data centers have created in the non-generation components of electric bills.

However, this will change in the future. 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.

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. By taking unilateral action to block local data center development, state policymakers can reduce data center-induced costs for their constituents, but could also push data centers into more socially costly locations and thus substantially increase the costs borne by people outside their states. Investing in new generation, or taking careful steps to direct data centers to places with flatter supply curves and cleaner grids, could reduce the costs of data center growth going forward.

9 Appendix

A Mapping load changes to local retail price changes

A.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 B.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.

A.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 0.97. However, utility-state-specific regressions reveal substantial heterogeneity. Though the median slope is 1.02, the interquartile range spans 0.67 to 1.4 (Figure A.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 A.1: Utility-state-specific fit of local MCPs on regional MCPs



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 (0.969).


B Figures

Figure B.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 B.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 B.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 B.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 B.5: Marginal emissions intensity by region and year




Notes: For each of the 11 dispatch regions and each dispatch year (2014–2025), the simple annual mean incremental CO₂ emissions per MWh of a +1 GW load band added to that year’s annual merit order (metric tCO₂/MWh; emissions are physical, not demand-weighted). February 2021 ERCOT and SPP (Winter Storm Uri) hours are excluded. EPA CEMS hourly data.

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




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




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 B.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 B.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, ISONE, 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 US.

Figure B.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 CO₂ emissions intensity of the region's coal and natural-gas fleets (metric tCO₂/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 tCO₂/MWh). Per-plant CO₂ 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 B.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 B.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 B.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 B.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 $\times$ 0.70. February 2021 is excluded for ERCOT and SPP (Winter Storm Uri).

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




Notes: Monthly national marginal emissions intensity of data center load (metric tCO₂ per MWh): the incremental grid CO₂ attributable to data center load divided by that load. Incremental CO₂ 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 B.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 tCO₂ per MWh), evaluated year by year on each year's own dispatch grid: the incremental grid CO₂ attributable to data center load divided by that load. Incremental CO₂ 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 B.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). Continental U.S.; campuses with a reported opening year and capacity. Source: campus data from [Schlauch et al. \(2026\)](#).

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




Notes: Within-region counterpart to Figure 8, identical in every respect except the source of the clean controls. 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 (campus data from [Schlauch et al. \(2026\)](#)), 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 drawn from *inside* its own dispatch region — utility-states with no ≥ 100 MW addition through t_0+5 , so incidental sub-threshold data center capacity is allowed — so region-wide price movements common to the host and its neighbors difference out and the coefficient is the within-region differential. The three plots are the residential retail price, the non-residential retail price (\$2025/MWh), and non-residential sales (TWh). Horizons -3 to 5 , with $h = -1$ the reference. Series A is unweighted; series B reweights controls to match treated units on deregulated-market status and on the changes in non-residential price and $\log(1 + \text{non-residential sales})$ from three years before to the year before treatment. Bands are 95% confidence intervals from state-clustered standard errors.

Figure B.20: Local projections DiD: wholesale price effects of data center entry




(a) Across-region controls



(b) Within-region controls


Notes: Event-stacked local projections difference-in-differences (Dube et al. 2025) on an LBNL gridcell annual panel, with event fixed effects. Within each panel, the three plots are the average wholesale price (\$2025/MWh; 2012–2024); the mean price of the 100 highest-priced hours in the year (\$2025/MWh; 2016–2024), winsorized at the 95th percentile of the cell-year distribution ($\approx \$1,189/\text{MWh}$) to limit the influence of extreme scarcity events (e.g. Winter Storm Uri, February 2021); and the share of hours with negative wholesale prices (percent; 2012–2024). An event is the first year a gridcell’s annual data center capacity addition reaches 100 MW (campus data from Schlauch et al. (2026)); the dose is that crossing-year addition, rescaled so estimates are per +1 GW. Each treated gridcell is stacked with its own clean controls: gridcells with non-zero data center capacity by 2024 but no ≥ 100 MW addition through t_0+5 , drawn from *outside* the treated cell’s dispatch region in panel (a) and from *inside* it in panel (b). Horizons -3 to 5 , with $h = -1$ the reference. Series A is unweighted; series B reweights controls to match treated cells on the change in average wholesale price from three years before to the year before treatment and on log population density. Both series use events with wholesale price data in both of those years (167 of 317 events; 154 for the top-100-hour price). In panel (a), the orange horizontal line on the average-wholesale-price plot is the dispatch-model benchmark: the model’s implied wholesale price effect of +1 GW of data center capacity ($\$0.34\text{--}0.42/\text{MWh}$, which we plot wider than its true range so that it is visible at this scale), computed as the demand-weighted hourly price wedge from removing data center load (utilization rate of 0.70). Bands are 95% confidence intervals from region-clustered standard errors.

Figure B.21: Data center capacity concentration by non-residential price



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 continental utility-states 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 B.22: Data center capacity concentration by non-residential sales



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 buildout (dashed, ranked by 2024 sales). Realized campuses are from [Schlauch et al. \(2026\)](#). Sample is the 1,274 continental utility-states 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 B.23: Cross-objective effects of the optimizing allocations




(a) Price effect of CO₂-optimizing allocations

(b) CO₂ 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 CO₂ (million metric tons 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 B.24: 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 forward-period pooled preference estimates in Table 2. *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. *Price-minimizing* and *Climate-minimizing*: single greedy allocations that minimize residential electricity costs or climate damages under the same screen. 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 at \$9,000/MWh; household costs are evaluated with that load unserved, so prices are bounded by the most expensive generator. Data center load = nameplate × 0.70.

Figure B.25: Incremental emissions by data center allocation, 2024 grid and augmented grid (2026–2030 buildout, 50% realization)



Notes: Incremental CO₂ (million metric tons 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 forward-period pooled preference estimates in Table 2. *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. *Price-minimizing* and *Climate-minimizing*: single greedy allocations that minimize residential electricity costs or climate damages under the same screen. 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 at \$9,000/MWh; emissions are evaluated with that load unserved. Data center load = nameplate × 0.70.

Figure B.26: 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 CO₂ from the same 1 GW addition, averaged across hours (metric tCO₂/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 × 0.70.

C Tables

Table C.1: Social costs of imposed data center allocations (2021–2025 buildout)

Outcome	Random	Baseline	Social cost min.	
			Min. price effect	Min. climate damages
Price effect (\$/MWh)	0.75	0.72	0.58	1.16
Incremental CO ₂ (Mt)	123.98	115.54	134.82	96.69
Total climate damages (\$B/yr)	\$304.6	\$303.0	\$306.6	\$299.5
Δ residential electricity costs (\$B/yr)	\$1.04	\$1.01	\$0.81	\$1.62
Δ climate damages (\$B/yr)	\$22.94	\$21.37	\$24.94	\$17.89
Δ consumer surplus (\$B/yr)	-\$23.98	-\$22.38	-\$25.76	-\$19.50

Notes: Data center campuses built 2021–2025 reallocated across utility-states. *Random* and *Baseline* report the median of each row over 2,000 campus-level draws subject to a generation capacity screen: *Random* draws utility-states in proportion to land area, *Baseline* from the equilibrium choice probabilities under the type-specific preference estimates in Table 2. *Min. price effect* and *Min. climate damages* are single greedy cost-minimizing allocations under the same screen. Because rows are medians taken separately, the consumer surplus row need not equal the sum of its components. Price effect is the residential-load-weighted retail price change from data center load. Incremental CO₂ 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 = $-\text{costs} - \text{climate damages}$. Location choice is evaluated with load beyond thermal capacity priced at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC \$185/t (2025\$); data center load = nameplate $\times$ 0.70; 2019 dispatch.

Table C.2: Social costs of imposed data center allocations (2026–2030 buildout, 50% realization)

Outcome	Random	Baseline	Social cost min.	
			Min. price effect	Min. climate damages
Price effect (\$/MWh)	2.80	2.81	2.32	3.68
Incremental CO ₂ (Mt)	281.69	278.35	283.18	256.08
Total climate damages (\$B/yr)	\$276.1	\$275.5	\$276.4	\$271.4
Δ residential electricity costs (\$B/yr)	\$3.83	\$3.85	\$3.18	\$5.03
Δ climate damages (\$B/yr)	\$52.11	\$51.49	\$52.39	\$47.38
Δ consumer surplus (\$B/yr)	-\$55.96	-\$55.35	-\$55.57	-\$52.41

Notes: Announced 2026–2030 grid buildout (Aterio high-likelihood pipeline; behind-the-meter excluded) at 50% realization, reallocated across utility-states. *Random* and *Baseline* report the median of each row over 2,000 campus-level draws subject to a generation capacity screen: *Random* draws utility-states in proportion to land area, *Baseline* from the equilibrium choice probabilities under the forward-period pooled preference estimates in Table 2. *Min. price effect* and *Min. climate damages* are single greedy cost-minimizing allocations under the same screen. Because rows are medians taken separately, the consumer surplus row need not equal the sum of its components. Price effect is the residential-load-weighted retail price change from data center load. Incremental CO₂ 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. Location choice is evaluated with load beyond thermal capacity priced at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC \$185/t (2025\$); data center load = nameplate $\times$ 0.70; 2024 dispatch.

Table C.3: Policy counterfactuals with data center locations fixed and with firms relocating (2026–2030 buildout, 50% realization)

Outcome	Baseline	Carbon tax		Virginia surcharge	
		Locations fixed	Firms relocate	Locations fixed	Firms relocate
Price effect (\$/MWh)	2.81	17.28	17.37	2.81	2.80
Incremental CO ₂ (Mt)	278.35	259.15	254.18	278.35	278.35
Total climate damages (\$B/yr)	\$275.5	\$239.8	\$238.8	\$275.5	\$275.5
Δ residential electricity costs (\$B/yr)	\$3.85	\$23.62	\$23.74	\$3.85	\$3.83
Δ climate damages (\$B/yr)	\$51.49	\$47.94	\$47.02	\$51.49	\$51.49
Δ government transfers (\$B/yr)	\$0.00	\$47.94	\$47.02	\$0.64	\$0.61
Δ consumer surplus (\$B/yr)	-\$55.35	-\$23.62	-\$23.74	-\$54.71	-\$54.73
Capacity in Virginia (GW)	9.43	9.43	9.36	9.43	9.02
Capacity relocated vs. baseline (GW)	0.00	0.00	3.15	0.00	0.40

Notes: Announced 2026–2030 grid buildout (Aterio high-likelihood pipeline; behind-the-meter excluded) at 50% realization, reallocated across utility-states using the forward-period pooled preference estimates in Table 2. *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/tCO₂ 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 CO₂ 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 at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC \$185/t (2025\$); data center load = nameplate $\times$ 0.70; 2024 dispatch.

Table C.4: Policy counterfactuals with data center locations fixed and with firms relocating (2026–2030 buildout, 50% realization, augmented grid)

Outcome	Baseline	Carbon tax		Virginia surcharge	
		Locations fixed	Firms relocate	Locations fixed	Firms relocate
Price effect (\$/MWh)	2.35	13.83	13.64	2.35	2.35
Incremental CO ₂ (Mt)	250.95	221.84	217.75	250.95	250.82
Total climate damages (\$B/yr)	\$228.4	\$197.7	\$197.0	\$228.4	\$228.4
Δ residential electricity costs (\$B/yr)	\$3.21	\$18.90	\$18.64	\$3.21	\$3.21
Δ climate damages (\$B/yr)	\$46.43	\$41.04	\$40.28	\$46.43	\$46.40
Δ government transfers (\$B/yr)	\$0.00	\$41.04	\$40.28	\$0.60	\$0.57
Δ consumer surplus (\$B/yr)	-\$49.63	-\$18.90	-\$18.64	-\$49.04	-\$49.03
Capacity in Virginia (GW)	8.93	8.93	8.86	8.93	8.52
Capacity relocated vs. baseline (GW)	0.00	0.00	2.78	0.00	0.40

Notes: Announced 2026–2030 grid buildout (Aterio high-likelihood pipeline; behind-the-meter excluded) at 50% realization, reallocated across utility-states using the forward-period pooled preference estimates in Table 2. *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/tCO₂ 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 CO₂ 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 at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC \$185/t (2025\$); data center load = nameplate $\times$ 0.70.

Table C.5: State data center moratoria ranked by Δ consumer surplus (2026–2030 buildout, 50% realization)

	Capacity displaced (GW)	Δ residential electricity costs (\$M)		Total relocation effects	
		In-state	Out-of-state	Climate damages (\$M)	Consumer surplus (\$M)
California	3.20	-104.8	209.4	+\$885.5	-\$990.2
Texas	9.96	-207.7	778.0	+\$93.1	-\$663.4
New York*	1.41	-39.9	82.9	+\$388.5	-\$431.6
Georgia*	4.07	-43.5	28.4	+\$435.4	-\$420.3
Florida	2.42	-61.2	43.2	+\$242.2	-\$224.3
⋮					
Indiana	1.34	-11.4	32.5	-\$187.3	+\$166.1
Washington	1.54	-65.9	-44.5	-\$61.2	+\$171.7
Oregon	1.58	-32.1	-79.6	-\$62.2	+\$173.8
Michigan*	1.79	-17.7	45.0	-\$243.7	+\$216.4
Illinois	3.74	-14.6	-7.3	-\$224.1	+\$246.0
Avg, proposed ban*	1.59	-15.8	20.2	+\$16.1	-\$20.5
Avg, no ban proposed	1.49	-23.0	32.0	+\$17.5	-\$26.6

Notes: Per-state moratorium on the announced 2026–2030 grid buildout at 50% realization (behind-the-meter excluded), using the forward-period pooled preference estimates in Table 2: 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 CO₂ valued at the SCC. Consumer surplus = $-\text{in-state costs} - \text{out-of-state costs} - \text{climate damages}$. The five largest consumer-surplus decreases and increases are shown; * marks states with a proposed moratorium (NCSL tracker, July 2026); group rows are unweighted averages over all continental states. Location choice is evaluated with load beyond thermal capacity priced at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC \$185/t (2025\$); data center load = nameplate $\times$ 0.70; 2024 dispatch.

Table C.6: State data center moratoria ranked by Δ consumer surplus (2026–2030 buildout, 50% realization, augmented grid)

	Capacity displaced (GW)	Δ residential electricity costs (\$M)		Total relocation effects	
		In-state	Out-of-state	Climate damages (\$M)	Consumer surplus (\$M)
Texas	9.99	-143.5	482.0	+\$1124.1	-\$1462.6
California	3.49	-215.2	144.3	+\$781.7	-\$710.8
New York*	1.39	-36.7	63.5	+\$309.9	-\$336.7
Georgia*	4.03	-36.7	-20.0	+\$336.0	-\$279.2
Louisiana	0.82	-5.7	32.9	+\$132.5	-\$159.6
⋮					
Iowa	1.29	-2.9	14.1	-\$254.0	+\$242.8
Indiana	1.33	-9.8	21.2	-\$261.6	+\$250.2
Michigan*	1.77	-15.3	30.1	-\$344.3	+\$329.5
Illinois	3.70	-9.7	13.6	-\$380.7	+\$376.7
Virginia*	8.93	-62.4	39.6	-\$433.9	+\$456.7
Avg, proposed ban*	1.57	-12.5	10.9	-\$35.4	+\$37.0
Avg, no ban proposed	1.49	-21.7	20.7	+\$41.9	-\$40.8

Notes: Per-state moratorium on the announced 2026–2030 grid buildout at 50% realization (behind-the-meter excluded), using the forward-period pooled preference estimates in Table 2: 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 CO₂ valued at the SCC. Consumer surplus = $-\text{in-state costs} - \text{out-of-state costs} - \text{climate damages}$. The five largest consumer-surplus decreases and increases are shown; * marks states with a proposed moratorium (NCSL tracker, July 2026); group rows are unweighted averages over all continental states. 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 at \$9,000/MWh; household costs and emissions are evaluated with that load unserved, so prices are bounded by the most expensive generator. SCC \$185/t (2025\$); data center load = nameplate $\times$ 0.70.

Table C.7: Data center allocations and regional headroom (2026–2030 buildout, 50% realization, augmented grid)

Region	Headroom	Unconstr. Modeled	Capacity-screened allocation	
			Random	Modeled
PJM	27.4	12.7	2.9	12.8
Southeast	27.8	8.3	6.0	8.3
MISO North	16.6	6.3	8.2	6.3
ERCOT	11.5	6.0	5.2	6.0
SPP	17.2	4.2	10.0	4.2
NorthernGrid	4.8	3.9	4.8	3.8
WestConnect	8.9	3.0	6.9	3.0
CAISO	3.1	2.3	2.0	2.3
MISO South	12.0	1.7	3.1	1.7
ISONE	3.9	1.6	1.2	1.6
NYISO	3.5	1.0	0.6	1.0
<i>Total</i>	136.8	51.0	51.0	51.0

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 C.8: Utility-state characteristics by data center exposure

Group	N	Median sales (TWh)			Median price (\$/MWh)		Median area	Data centers (MW)	
		Total	Res.	Non-res.	Res.	Non-res.		Stock 2024	Median addition
1. No data centers	1131	0.351	0.146	0.187	153.6	125.3	2392	0.0	–
2. Some DC, never treated	166	3.419	1.586	1.358	156.9	125.0	7759	10.0	3.5
3. Treated	59	4.027	1.807	1.743	149.5	113.6	22611	393.0	28.0

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 C.9: 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., $0.70 \times 8,760 \text{ hours} \times \text{rate}$). 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 C.10: LBNL wholesale prices on dispatch-model market-clearing prices

	Levels	Logs
Dispatch-model MCP	1.314* (0.452)	0.846*** (0.149)
95% CI	[0.31, 2.32]	[0.51, 1.18]
Region FE	Yes	Yes
Year FE	Yes	Yes
R^2	0.81	0.88
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 C.11: Focus firm summary statistics

Firm	Pre-2015		2019-2024	
	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	279	28016

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 43% of 2019–2024 MW.

Table C.12: Firm Bartik IV using 2018–2025 retail prices

	OLS	IV
Res. price (\$/MWh)	-2.70 [-7.39, 2.00]	-6.48 [-26.85, -2.37]
Nonres. price (\$/MWh)	-5.27 [-9.27, -1.26]	-7.39 [-24.71, -3.88]
Nonres. sales (TWh)	+2.65 [2.17, 3.14]	+3.36 [2.42, 9.60]
First-stage F (state-clust.)		6.8
Observations	365	365

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 continental utility-states; 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 C.13: 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	0.583	-0.72	-1.20	2.02
Microsoft	2.38	0.172	-3.63	-3.26	2.15
QTS	0.81	0.101	-0.19	-1.46	1.85
Digital Realty	0.72	0.096	1.21	0.04	1.79
Meta	2.51	0.038	-10.02	-10.82	5.48
Google	1.67	0.011	-36.72	-37.50	10.01

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 C.14: Persistence: firm-by-utility-state Poisson regression

Dependent Variable:	MW built 2019-2024					
Model:	Amazon (1)	Meta (2)	Microsoft (3)	Google (4)	QTS (5)	Digital Realty (6)
<i>Variables</i>						
Constant	0.508*** (0.021)	0.618*** (0.020)	-0.007 (0.027)	0.216*** (0.024)	-0.853*** (0.042)	-2.48*** (0.094)
Pre-period MW	0.009*** (4.18×10^{-5})	-0.093 (0.673)	0.027*** (0.0002)	-0.150 (1.71)	0.009*** (0.0001)	0.021*** (0.0002)
<i>Fit statistics</i>						
Observations	1,355	1,355	1,355	1,355	1,355	1,355
Squared Correlation	0.526	1.35×10^{-5}	0.537	1.56×10^{-5}	0.271	0.982
Pseudo R^2	0.414	0.0004	0.424	0.0008	0.245	0.849
BIC	25,114.5	25,470.8	14,577.3	17,975.9	7,243.5	1,465.0

IID standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Per-firm Poisson regression of firm-by-utility-state 2020–2025 MW on the firm's pre-period (pre-2015) MW in that utility-state. One column per focus firm.

Table C.15: Leave-one-firm-out firm Bartik 2SLS (residential price)

Dropped	Estimate	SE	Eff. F
None (all firms)	-2.59	2.10	7.3
Amazon	-4.50	4.53	1.9
Google	-2.06	1.64	14.0
Microsoft	-1.92	2.17	5.4
Meta	-2.37	1.23	158.5
Digital Realty	-2.82	2.24	6.7
QTS	-2.90	2.38	5.2

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 C.16: Firm Bartik reduced form, pre- vs. post-period

	Res. price		Nonres. price		Nonres. sales	
	2018–2024 (1)	2012–2018 (2)	2018–2024 (3)	2012–2018 (4)	2018–2024 (5)	2012–2018 (6)
<i>Model:</i>						
<i>Variables</i>						
Firm-Bartik instrument	-0.003 (0.002)	-0.003 (0.002)	-0.004*** (0.002)	-0.001 (0.002)	0.003*** (0.0004)	0.002*** (0.0002)
Average non-residential electricity price, 2016	-0.036 (0.034)	-0.104 (0.101)	-0.134*** (0.028)	-0.077 (0.105)	-0.002 (0.002)	-0.003* (0.002)
Log(non-residential electricity load, 2016)	1.59*** (0.539)	0.883 (0.643)	1.23*** (0.452)	0.560 (0.653)	-0.033 (0.061)	-0.020 (0.053)
Log(land value)	-1.46* (0.742)	-0.174 (0.549)	-0.379 (0.661)	0.241 (0.442)	0.043 (0.046)	0.025 (0.031)
<i>Fixed-effects</i> region	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	1,247	1,247	1,247	1,247	1,247	1,247
R ²	0.167	0.112	0.160	0.105	0.206	0.075
Within R ²	0.044	0.040	0.077	0.021	0.159	0.050

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 C.17: Regional generation capacity, headroom, and the announced 2026–2030 build-out

Region	2024 grid		Planned 2026–2030		
	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	24.0	3.8	11.2
MISO_North	75.7	9.3	20.0	3.9	3.7
Southeast	126.1	18.4	17.9	7.4	3.1
SPP	47.9	7.0	11.4	9.8	1.6
NorthernGrid	25.4	3.3	10.3	0.9	3.7
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
<i>Total</i>	604.3	79.3	146.3	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 C.18: 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 C.19: 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

- Allen, Treb, and Costas Arkolakis.** 2023. "Economic Activity across Space: A Supply and Demand Approach." *Journal of Economic Perspectives*, 37(2): 3–28.
- 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.
- 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 Bushnell.** 2015. "The US Electricity Industry After 20 Years of Restructuring." *Annual Review of Economics*, 7(1): 437–463.

- 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/>.
- Brühlhart, 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."
- Gagnon, Pieter, Brady Cowiestoll, and Marty Schwarz.** 2023. "Cambium 2022 Scenario Descriptions and Documentation." National Renewable Energy Laboratory.
- Georgia Power.** 2026. "Georgia Power's contract with OpenAI approved; Latest approval part of portfolio of large-load contracts delivering ap-

proximately \$950 million in annual savings beginning in 2029.” *Georgia Power*. <https://www.georgiapower.com/news-hub/press-releases/contract-openai-approved-part-portfolio-delivering-950-million-annual-savings.html>.

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.

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*.

- 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/>.
- 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.
- 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*.
- 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.”

Morales, Jowi. 2026. “Indiana power provider proposes \$59 million rate cut for state, says data centers and other large customers are driving increased revenue — move could potentially save residential users \$100 annually.” *Tom’s Hardware*, Section: Data Centers. <https://www.tomshardware.com/tech-industry/data-centers/indiana-power-provider-proposes-usd59-million-rate-cut-for-state-says-data-centers-and>

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/>.

Penrod, Emma. 2026. “Data center growth has helped PG&E cut rates 11% since 2024, CEO says.” *Utility Dive*. <https://www.utilitydive.com/news/data-center-growth-has-helped-pge-cut-rates-11-since-2024-ceo-says/812230/>.

- Pontecorvo, Emily.** 2026. "Americans Now Blame Data Centers for Their Rising Power Bills -." *Heatmap News*.
- 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*.
- Schlauch, Garrison, Coly Elhai, Simon Greenhill, Falco J Bargagli Stoffi, and Gianluca Guidi.** 2026. "Building a Comprehensive Database 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/>.
- 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.** n.d.. "Transmission Impact Assessment: Power Sector Infrastructure Deployment to Reduce Costs, Improve Reliability, and Lower Pollution."
- Wiser, Ryan, Eric O'Shaughnessy, Galen Barbose, Peter Cappers, and Will Gorman.** 2025. "Factors influencing recent trends in retail electricity prices in the United States." *The Electricity Journal*, 38(4): 107516.
- 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.