MIT CEEPR
Center for Energy and Environmental Policy Research

Working Paper Series

Who Pays for Growth? Evidence from Data Centers and the Grid

Robyn Meeks, Jacquelyn Pless, Zhiyuan Qi, and Zhenxuan Wang

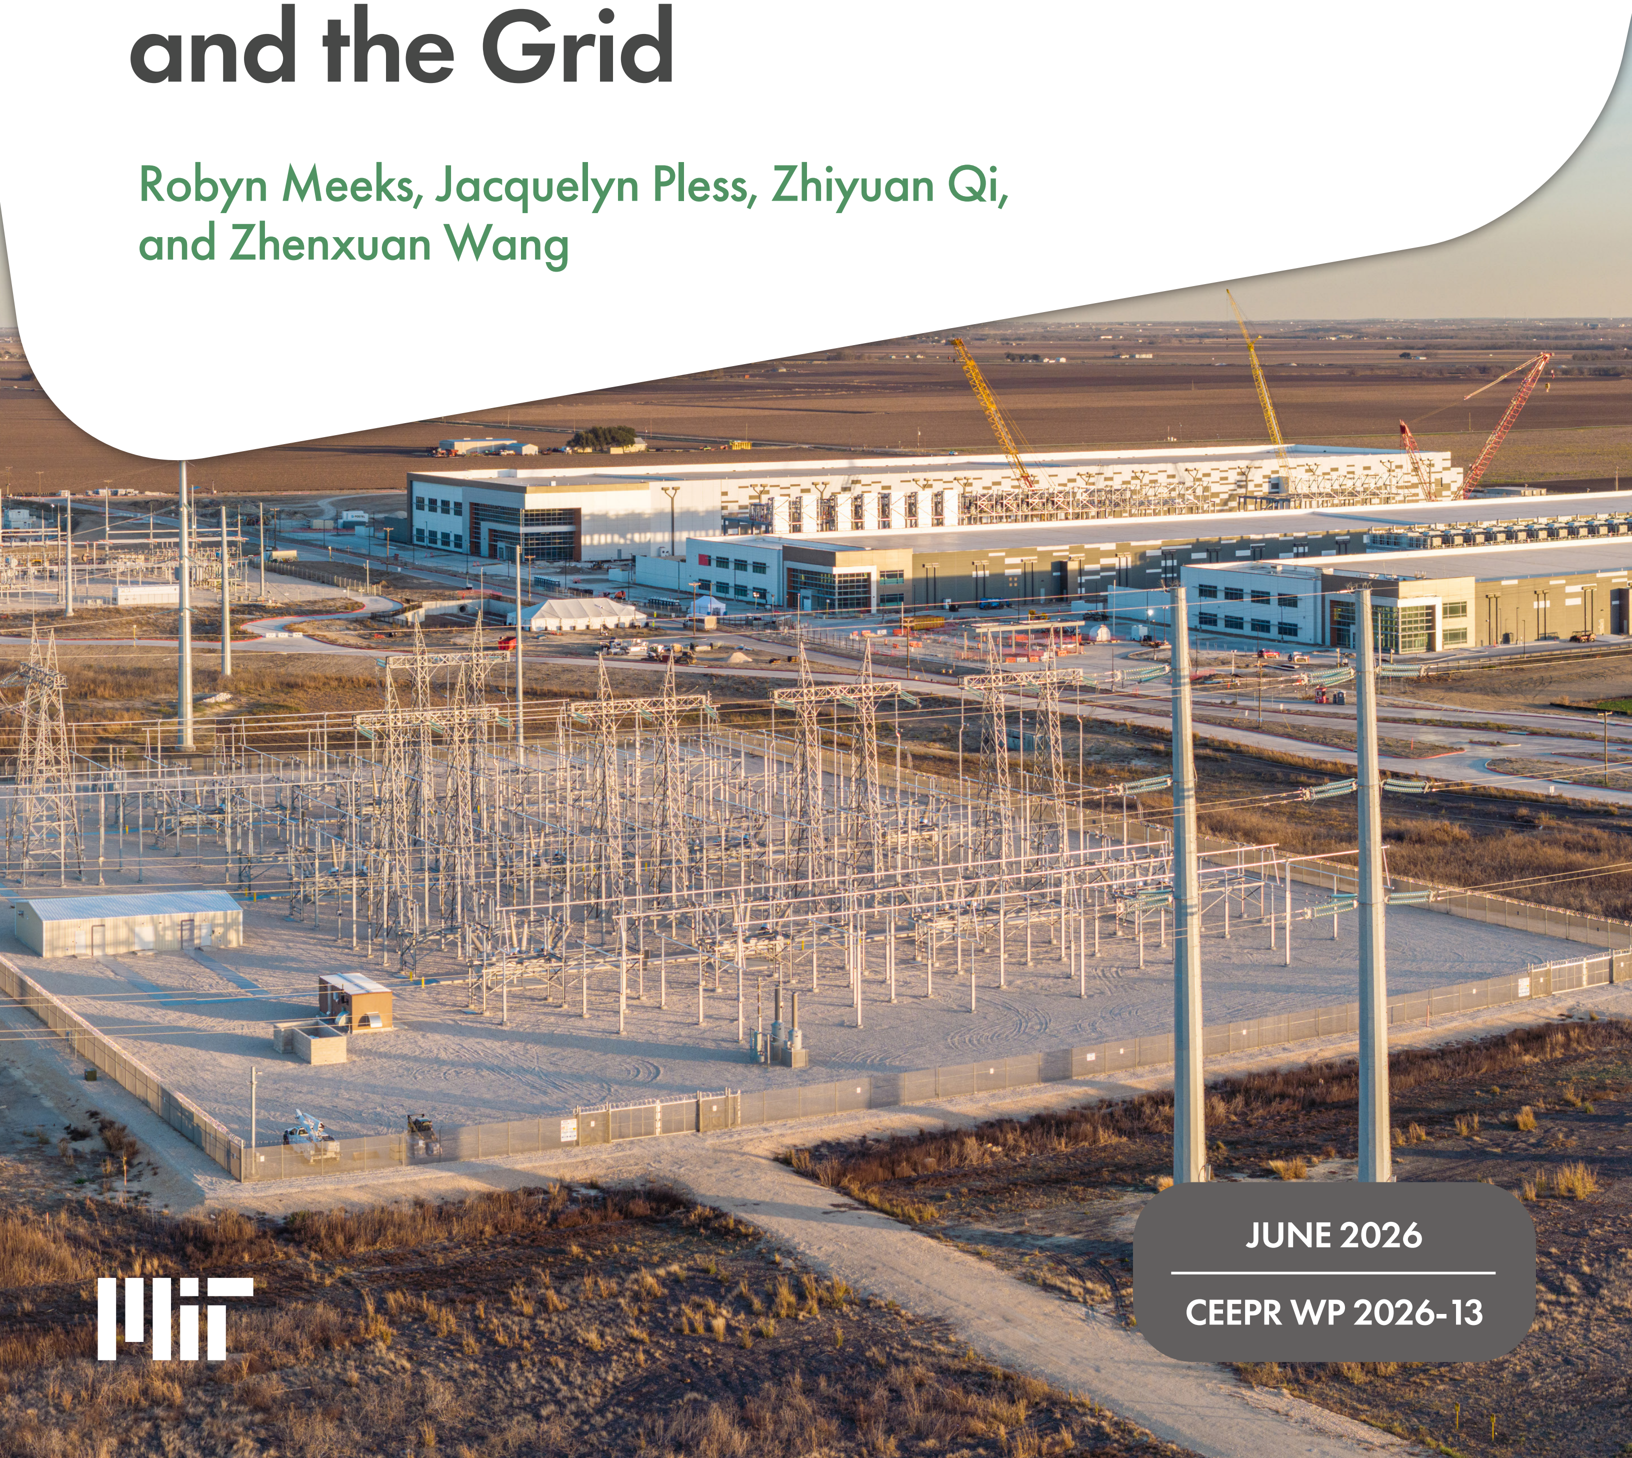

JUNE 2026

CEEPR WP 2026-13

Working Paper Series.

Since 1977, the Center for Energy and Environmental Policy Research (CEEPR) has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. Drawing on the unparalleled resources available at MIT, affiliated faculty and research staff as well as international research associates contribute to the empirical study of a wide range of policy issues related to energy supply, energy demand, and the environment.

An important dissemination channel for these research efforts is the MIT CEEPR Working Paper series. CEEPR releases Working Papers written by researchers from MIT and other academic institutions in order to enable timely consideration and reaction to energy and environmental policy research, but does not conduct a selection process or peer review prior to posting. CEEPR's posting of a Working Paper, therefore, does not constitute an endorsement of the accuracy or merit of the Working Paper. If you have questions about a particular Working Paper, please contact the authors or their home institutions.

Who Pays for Growth?
Evidence from Data Centers and the Grid*

Robyn Meeks[†] Jacquelyn Pless*[‡] Zhiyuan Qi[§] Zhenxuan Wang[¶]

June 19, 2026

Abstract

Artificial intelligence is driving rapid growth in electricity demand through the expansion of data centers. Meeting this demand requires complementary infrastructure investment, but who pays for that investment depends on the institutions that finance, regulate, and allocate its costs. We study this question by linking facility-level data on data center entry to utility-level data on retail electricity prices, infrastructure investment, operational outcomes, and regulatory rate case proceedings across the United States. Data center entry between 2010 and 2024 increases average retail electricity prices by 2.7%, with effects of 2.1% for residential customers, 2.8% for commercial customers, and 4.2% for industrial customers. These effects are concentrated among investor-owned utilities, where average retail prices rise by 5.6%, while effects are much smaller among publicly-owned utilities and absent among cooperatives. Price increases are also larger in states with deregulated electricity generation. To explain these patterns, we examine utility investment, operations, and rate case outcomes. Data center entry induces real infrastructure and operational responses, including construction work in progress, generation investment, transmission spending, and larger authorized rate bases. But pass-through depends on regulatory environment: vertically integrated utilities in regulated states appear to absorb load growth partly through improved utilization, while utilities in deregulated states exhibit more active cost recovery and greater exposure to system constraints. The results show that the incidence of infrastructure-intensive technological change depends not only on the technology itself, but also on the institutions governing complementary infrastructure investment and cost recovery.

Keywords: incidence; infrastructure investment; electricity prices; technological change; rate-of-return regulation
JEL Codes: Q41, L94, H23, L51, O33

*Zhenxuan Wang gratefully acknowledges the funding support from the Institute for Humane Studies (grant no. IHS020420) for data purchase.

†Sanford School of Public Policy, Duke University. Email: robyn.meeks@duke.edu.

‡*Corresponding author. MIT Sloan School of Management. Email: jpless@mit.edu.

§Department of Economics, North Carolina State University. Email: zqi8@ncsu.edu.

¶Department of Agricultural and Resource Economics, North Carolina State University. Email: zhenxuan.wang@ncsu.edu.

1 Introduction

Artificial intelligence (AI) is unleashing one of the largest infrastructure booms in modern history. Behind the rapid advances in frontier AI models, cloud computing, and digital services lies an enormous and fast-growing physical footprint: hyperscale data centers consuming vast quantities of electricity. Global investment in data centers that rely crucially on a stable source of electricity supply nearly doubled between 2022 and 2024 (IEA, 2025), reaching roughly half a trillion dollars, and projections indicate that electricity use from data centers in the U.S. alone could more than double by 2030 (Shehabi et al., 2024). The scale and speed of this expansion have transformed data centers from a niche component of the digital economy into one of the fastest growing sectors of the U.S. economy and the most consequential source of electricity demand growth in the twenty first century.

This surge arrives after decades of relatively flat U.S. electricity demand and represents a fundamentally different form of load growth from what electricity systems have historically been built to accommodate. Unlike traditional demand growth—which tends to be gradual, diffuse, and tied to population or economic expansion—data centers arrive in large, concentrated increments and operate continuously at exceptionally high utilization rates. Meeting this demand often requires substantial investment in electricity generation, transmission, and distribution infrastructure, placing pressure on grids that are already facing reliability concerns and long interconnection queues. As a result, concerns over infrastructure investment, electricity prices, and the allocation of costs across ratepayers have become central to debates over data center expansion. While data centers may bring capital investment, jobs, tax revenue, and other local economic benefits, they also raise a fundamental question around who will pay for the infrastructure needed to support them. Public concern increasingly centers on whether incumbent ratepayers—households, businesses, and large industrial customers—will ultimately bear the costs of serving large new loads as utilities and regulators face mounting pressure to expand capacity quickly.

In this paper, we ask: who pays for growth? We study how data center entry affects U.S. retail electricity rates, with particular attention to how these effects differ across residential, commercial, and industrial customers. The effect on retail electricity prices is theoretically ambiguous. Large new loads may increase generation, transmission, distribution, and reliability costs. But they may also lower average costs by expanding the customer base over which fixed investments are recovered, especially if they improve system utilization. Moreover, the incidence of these costs need not be uniform. Average retail rates can mask substantial differences in how infrastructure costs are allocated across customer classes, utilities, and regulatory environments. We therefore study not only whether data centers raise electricity prices, but also how electricity sector institutions shape who ultimately bears the costs of demand growth.

To do so, we construct a novel panel linking facility-level data on data center entry and expansion to utility-level retail electricity prices across the U.S. Our data combine detailed information on the location, timing, and scale of data centers with utility-level retail rates by customer class, utility ownership, and regulatory structure. We also assemble utility-level data on infrastructure investments, operational outcomes, rate cases, customer advances for construction, large-load tariffs, and pre-existing system constraints. This allows us to trace the effects of data center entry from retail prices to the underlying mechanisms through which utilities accommodate large-load growth and recover associated costs.

Our empirical analysis uses the Callaway and Sant'Anna (2021) estimator and we begin by examining event-study type models that compare changes in outcomes for utility-state pairs whose service territories receive a large data center to those that never do. These specifications allow us to assess pre-trends, characterize the timing of pass-through, and examine how price effects evolve as utilities invest in infrastructure and recover associated costs. A key empirical challenge is that data centers do not locate randomly. Firms select locations based on factors that may also affect electricity prices, including electricity access, land availability, permitting, tax incentives, grid capacity, and digital connectivity. We address this concern in several ways. Our baseline design absorbs time-invariant differences across utility-state pairs and common annual shocks. We also show robustness to allowing utilities and regions with different baseline characteristics to follow different time paths, to dropping potentially contaminated control units within multi-state utilities, and to a matched difference-in-differences design that balances never-treated utilities to treated utilities on pre-period electricity prices, customer counts, manufacturing intensity, and income. Across these checks, the estimated price effects remain similar and, in the matched design, generally become larger.

We find that data center entry increases retail electricity rates, but the burden is not distributed uniformly. Average retail prices rise by 2.7% following data center entry. Residential rates increase by 2.1%, commercial rates by 2.8%, and industrial rates by 4.2%. The positive residential effect indicates that at least some of the costs of data center growth are passed through to incumbent households. The larger effects for commercial and industrial customers suggest that non-residential customers bear a larger share of the price increases, though these estimates should be interpreted with care: depending on utility tariff classifications, data centers themselves may be recorded as commercial, industrial, or special large-load customers. As a result, the commercial and industrial effects may reflect both pass-through to incumbent firms and the prices paid by the new large-load customers generating the demand shock.

Importantly, the incidence of data center-driven load growth differs across utility ownership types and regulatory environment, indicating that the consequences of data center expansion depend not only on the scale of new electricity demand, but also on the institutions through which infrastructure costs are financed, recovered, and allocated. Retail price increases are concentrated among investor-owned utilities (IOUs), which serve more than 70% of U.S. electricity customers. Among IOUs, data center entry increases average retail prices by 5.6%, with residential prices rising by 5.2%, commercial prices by 4.6%, and industrial prices by 8%. Effects are smaller among public utilities and absent among cooperatives. We also find larger price effects for utilities operating in states with deregulated electricity markets. Average retail prices rise by 3.8% in deregulated states, compared to a smaller and statistically insignificant increase in regulated states. A two-way split shows that ownership is the first-order source of heterogeneity: price effects are concentrated among IOUs in both regulated and deregulated states, but they are generally larger among IOUs operating in deregulated markets.

In residential electricity bill terms, these effects are meaningful but relatively modest, on average. When evaluated at pre-treatment prices and sales, the residential effect corresponds to roughly $3 per month for the average household, rising to about $7 per month among IOU customers. That said, we remain cautious in our interpretation with respect to electricity bill implications, as our estimates reflect realized average pass-through over our 2010–2024 study period rather than an upper bound on the potential ratepayer consequences of the most recent and upcoming waves of hyperscale data center growth.[1]

We next provide evidence on the mechanisms linking data center growth to higher retail rates. We start by showing that data center entry induces real infrastructure and operational responses among IOUs, but the nature of the response differs sharply by regulatory environment. In deregulated states, electric construction work in progress (CWIP) increases immediately—in fact, it increases even the year before data center activation—and transmission operations and maintenance (O&M) starts to increase a few years prior to data center entry as well, suggesting that these utilities begin investing in grid upgrades and electricity-related capital projects in anticipation of new load before those assets are placed in service. In regulated states, utilities also appear to accommodate

[1] Due to our research design, we also omit utility-state pairs that receive their first data center before 2010, and some of these have very large data centers or many of them (e.g., some utility-states in Virginia).

data center growth through transmission O&M as well as generation investment, and load factors improve a few years later. Increases in electric CWIP and completed electric plants in service also appear a few years later in regulated states. This combination is consistent with absorption through planning, utilization gains, and fixed-cost spreading, helping to at least partly explain why regulated states exhibit muted contemporaneous price effects despite evidence of infrastructure response.

Regulatory cost recovery further explains how infrastructure costs pass through to retail prices. We find that data center entry is associated with larger authorized rate bases in both regulated and deregulated states, indicating that additional capital is recognized by regulators as prudent, used and useful, and eligible for cost recovery. However, rate base expansion does not translate mechanically into measured annual revenue increases. Requested revenue requirements rise in deregulated states, but authorized revenue requirements do not increase proportionally, and we find no evidence that regulators authorize higher returns on equity or higher returns on the rate base. Instead, the cost-recovery mechanism operates through a larger approved capital base, more active regulatory proceedings, and faster cost-recovery channels. In deregulated states, data center entry increases the probability of filing any rate case, general rate cases, rider proceedings, and the number of cases, while case duration declines. These patterns suggest more frequent and faster regulatory engagement. We also find increases in customer advances for construction in deregulated states, suggesting that utilities attempt to assign some infrastructure costs directly to large-load customers, but this does not fully offset cost pass-through.

Finally, system constraints and wholesale market exposure also appear to partly explain why pass-through is stronger in deregulated markets. Price effects are larger in low-reserve-margin regions, particularly in deregulated states, consistent with a capacity adequacy mechanism: when systems have limited spare capacity, large new loads are more likely to require expensive generation procurement, capacity payments, reliability expenditures, or accelerated infrastructure investment. We also find larger effects in fossil-heavy deregulated regions specifically, consistent with exposure to marginal generation and fuel costs. These results suggest that deregulated utilities face more immediate cost pressure when new load enters constrained systems, while regulated utilities may smooth those costs through integrated resource planning, owned generation, long-term procurement, and gradual rate recovery.

Taken together, the mechanism results show that data centers generate infrastructure and capital accounting responses in both regulated and deregulated settings, but retail price increases are concentrated where cost pressures are not offset by utilization gains and are transmitted more quickly through market and regulatory cost-recovery channels.

Related Literature and Contributions. This paper contributes to a nascent but rapidly growing literature examining the economic consequences of data center expansion. Contemporaneous work studies how data centers affect local economic activity, electricity markets, and environmental outcomes. For example, Alvarez et al. (2026) and Yue and Zeng (2026) examine the effects of data center growth on local employment, wages, housing markets, and electricity prices, emphasizing agglomeration benefits and local economic development. Other work studies the environmental consequences of data centers (Knittel, Senga and Wang, 2025; Muller, 2026; Hernandez-Cortes, Meng and Weber, 2026). Our analysis complements these studies by shifting attention from the aggregate economic and environmental consequences of data center growth to the distributional consequences of supporting that growth. In particular, we examine how the costs associated with expanding electricity infrastructure are allocated across households and firms and the underlying institutional mechanisms shaping how costs are distributed.

With respect to the electricity sector, our paper is closely related to Kay, Reaser and Taylor (2026), who show that data center growth can substantially increase wholesale electricity prices and generation costs, particularly in regions facing supply constraints. However, higher wholesale market effects do not mechanically determine retail prices. Our analysis asks the incidence question: how do those wholesale and infrastructure costs pass through to retail prices, and who ultimately pays? Translating higher generation, capacity, transmission, and distribution costs into retail electricity prices requires a set of institutional processes, including utility investment decisions, rate cases, tariff design, customer class cost allocation, and ownership-specific approaches to cost recovery. We therefore focus on retail incidence rather than wholesale market outcomes. Our results show that the same underlying demand shock can have very different consequences depending on the institutions governing infrastructure finance and cost recovery.

More broadly, this paper contributes to the public economics literature on the costs and benefits of infrastructure investments. Existing work documents the economic gains from transportation, rail, communications, and other network infrastructure (e.g., Donaldson (2018) and Allen and Arkolakis (2022), among others). That literature has largely emphasized how infrastructure expands market access, raises productivity, and reshapes economic activity. Less is known about who ultimately pays when technological change requires new infrastructure. The setting of data center growth is useful because the benefits of AI and digital infrastructure are broadly dispersed, while the costs of complementary electricity infrastructure are mediated through local utilities, regulators, and ratepayers. Our findings therefore highlight a general feature of infrastructure-intensive technological change: its distributional consequences depend not only on the technology itself, but also on the institutions that finance complementary infrastructure and allocate its costs.

This paper also relates to several strands of literature in energy economics on retail pricing, cost pass-through, and electricity market design. Existing work examines how wholesale electricity costs, fuel costs, and emissions costs are transmitted through electricity markets (e.g., Fabra and Reguant (2014) and Borenstein and Bushnell (2015)). Others show how restructuring and wholesale market mechanisms affect production costs, procurement, and dispatch efficiency (Cicala, 2015, 2022). Recent studies on transmission constraints emphasize that incomplete grid integration can raise generation costs, limit gains from trade, and create political economy barriers to efficient grid expansion (Davis, Hausman and Rose, 2023; Hausman, 2025; Ham, Kay and Hausman, 2026). We build on these literatures by studying the retail incidence of a new source of load growth, showing that system constraints matter not only for wholesale market efficiency, but also for the pass-through of infrastructure-intensive demand shocks to households and firms.

Finally, our paper contributes to the literature on utility regulation. Classic work on cost-of-service and rate-of-return regulation emphasizes how regulation shapes utility investment behavior, cost minimization incentives, and capital expansion (Averch and Johnson, 1962; Joskow, 1974; Joskow and Rose, 1989). Recent work revisits these issues in modern electricity markets, showing that regulated returns, rate case incentives, and restructuring can shape capital investment and consumer costs (Dunkle Werner and Jarvis, 2025; Cicala, 2025; Gowrisankaran, Langer and Reguant, 2026). We contribute to this literature by studying how utility ownership, regulatory environment, rate-base treatment, rider proceedings, and cost-causation mechanisms shape the incidence of large-load growth. In contrast to work focused primarily on generation efficiency or utility investment incentives, we examine how a demand shock generated by technological change is translated into retail rates through regulatory institutions.

The rapid expansion of AI and digital infrastructure is likely to reshape electricity systems for decades to come. As policymakers, utilities, and regulators confront the challenge of financing large-scale infrastructure expansion, understanding who ultimately bears the costs of growth will become increasingly important. Our findings suggest that these distributional consequences are not determined by technology alone—they also depend on the institutional structures through which infrastructure investments are financed, regulated, and recovered. The economic consequences of infrastructure-intensive technological change depend fundamentally on the institutions that support it.

Outline. Section 2 provides background on data center growth, electricity markets, and utility regulation. Section 3 provides a conceptual framework. Sections 4 and 5 describe our data and empirical strategies. Section 6 presents the main price results and institutional heterogeneity. Section 7 examines mechanisms. Section 8 presents distributional implications across places. We conclude in Section 9.

2 Background

In this section we describe data centers' electricity demand and the investments needed to meet that demand, how utilities recover costs of large capital investments, and why incidence is ambiguous.

2.1 Data Centers, Electricity Demand, and Infrastructure Investments

Growth in energy demand from data centers has been unusually large compared to other sources of demand and non-linear over time. Data centers are energy intensive due to their continuous operation, computing equipment, and cooling needs (Satchwell et al., 2025). U.S. data center electricity use roughly doubled between 2000 and 2005, then roughly held constant through the early-to-mid 2010s even though the industry was expanding substantially (Shehabi et al., 2024). This leveling-off has been attributed to efficiency gains: improved cooling and power management, higher server utilization, computational efficiency improvements, reduced server idle power, and a structural shift of workloads (Masanet et al., 2020; Shehabi et al., 2024). That period of flat energy use ended approximately a decade ago, with U.S. data center energy use more than doubling between 2017 and 2023, reaching about 176 TWh, or 4.4% of total U.S. electricity consumption. Scenario projections have that further increasing to 6.7% to 12.0% of U.S. electricity by 2028 (Shehabi et al., 2024).

This growth in demand requires investments in new or expanded infrastructure: generation capacity is needed to meet the increasing demand, as well as transmission and distribution infrastructure to connect that demand growth to generation. That the majority of data centers are concentrated within some states intensifies the impacts on the local infrastructure and investment needs (Satchwell et al., 2025).

How much capacity utilities actually need to build is difficult to forecast. This is because realized load from these customers can depart sharply from what was projected: a data center's draw varies across the AI model lifecycle, running heavier during training phases and more erratically afterward, producing big shifts that are difficult to plan for (Satchwell et al., 2025). There are also challenges once a utility has committed to investing in long-lived assets sized to an expected load. If the demand does not materialize as predicted, a utility can be left carrying the costs of investments that it cannot fully recover from that particular customer and those costs could fall on the rest of the ratepayer base (Satchwell et al., 2025). The question of who bears the cost and risk of such investments is of considerable concern (Shehabi et al., 2024). To date no common approach has emerged for handling this forecast risk, and individual commissions and utilities are working out their own approaches (Satchwell et al., 2025).

Further, utilities' ability and incentives to make such capital investments to meet such demand growth may vary by utility ownership, and how they can recover the costs of any investments on the regulation guiding that. We explore the roles of utility ownership and regulatory status in the sections that follow.

2.2 Utility Ownership and Regulatory Oversight

Identical load shocks may generate very different retail incidence depending on the institutions responsible for the generation, transmission and distribution of electricity, as well as the regulatory entities overseeing these institutions.

There are three types of utilities in the U.S., which differ fundamentally in their ownership structure, and those ownership differences carry through to their finances. Investor-owned utilities (IOUs) are for-profit companies that are owned by shareholders. IOUs can be large entities covering multiple states, serving approximately three quarters of the country's population. In contrast, both public power utilities (also known as municipals or munis) and cooperatives are not-for-profit and consumer-owned: munis are owned by cities (or counties) and governed by the city council, while co-ops are owned by their members and run by a customer-elected board. These ownership structures in turn shape the utilities' incentives, with IOUs operating to earn a profit for its shareholders while co-ops operate at cost, returning any revenue above costs to its members.

The types of utilities also differ in their regulation and oversight, which determines the process through which electricity rates are set. IOUs face the most layered oversight, with their interstate generation, transmission, and power sales—if they exist —regulated by the Federal Energy Regulatory Commission (FERC) and their distribution and retail sales regulated by state public service commissions, often referred to as Public Utilities Commissions. In contrast, the consumer-owned utilities (COUs) are generally not regulated by state commissions.[2] Instead rates are typically set by their governing entities: the city council for munis and the member-elected board for co-ops. Like IOUs, however, COUs can be subject to a state's public utility laws addressing issues such as the timing of customer notice prior to rate adjustments.

2.3 Cost Recovery, Rate Cases, and Deregulation

The state regulatory commissions carry out multiple functions that determine electric retail incidence and utility cost recovery, including both setting rates and overseeing utilities' major financial decisions.

2.3.1 Rate Cases

Much of this work occurs through rate cases, which are the formal regulatory proceedings that serve as the principal vehicle for any significant change to prices. The process usually begins when a utility applies for a rate increase, arguing that its existing rates no longer give it a fair opportunity to earn its allowed return. Utilities often file every few years while others go longer between cases; some states, however, mandate a regular schedule for rate cases. Once filed, rate cases move through a fairly standard sequence of events, which takes time and eventually culminates in the commission specifying new rates.[3]

Rate cases function to determine utilities' revenue requirements (by setting the utility's rate base and allowed rate of return), apportion that requirement across the different classes of customers, and determine tariff structures that will generate the utilities' allowable revenue. We detail these components here.

1. Determining the revenue requirement. This step is the core of a general rate case and is built around a "test year"— either a recently completed historical year or a forecasted future year — so that costs, investments, and sales all reflect the same period. The commission applies a basic formula to calculate the revenue requirement as rate base × rate of return + operating expenses. The rate base is the utility's net long-lived investment

[2]Exceptions do exist, as noted in Lazar (2016).

[3]These steps include: the utility submits its proposed tariffs along with detailed supporting evidence; interested entities can join as parties, including the utility, the commission's staff, the consumer advocate, and any intervening entities that the commission admits, such as industrial, low-income, or environmental groups; the commission holds evidentiary hearings and separate sessions for public comment; lastly, the commission deliberates and issues a final order specifying the new rates and when they take effect.

used to serve its customers (e.g., the original cost of plant in service less accumulated depreciation, adjusted for working capital, deferred taxes, and other approved items). The rate base includes items such as power plants, poles, wires, transformers, as well as office buildings and related office goods such as vehicles and computers. To be included in the rate base an investment must be both "used and useful" and prudently incurred. The rate of return is a weighted cost of capital across the utility's funding sources—common equity, preferred equity, and debt—with the cost of common equity being the most commonly contested element (Lazar, 2016). Operating expenses (i.e., labor, fuel, purchased power, taxes, and depreciation) are recovered on a dollar-for-dollar basis.

2. Allocating the revenue requirement across customer classes. Once the total revenue requirement is set, the commission decides how much each class of customers should contribute, based on the usage characteristics of each class.[4] Costs are assigned according to cost causation, with some costs allocated by the number of customers, some by peak demand, and some by total energy consumption.[5] Allocation of costs can be driven by many factors, including individual judgment, politics, and economic development goals, amongst others.

3. Designing the rate within each customer class. The final step translates each class's allocated share into the actual prices customers pay, where essentially every rate element is an allocated portion of the revenue requirement divided by a billing determinant such as kWh, kW, or the number of customers. Rate design is normally handled within the general rate case. Addressing a single rate in isolation has been criticized with the concern that changing one customer's rate can shift costs onto others (Lazar, 2016).

2.3.2 Cost Recovery, Regulatory Lag, and Riders

The actual cost recovery for capital investments generally happens once the asset is "used and useful" and has been found to be prudently incurred, at which point the investment enters the rate base and the utility earns its allowed rate of return on the investment.

Given rates are reset only periodically through rate cases, any investment a utility makes today does not begin generating revenue from customers until a later proceeding

[4]Notably, there is no national uniformity in how classes are defined — most utilities separate residential customers, and many add commercial, industrial, and other classes.

[5]There are two broad methodologies used to assign costs, both of which tend to assign higher per-kWh costs to residential and small business customers (who need more distribution investment and have "peakier" usage) than to industrial customers. These methodologies are: embedded cost studies, used by about 30 states, which divide the same historic (or future test-year) accounting costs that make up the revenue requirement among the classes so that the class totals sum to that requirement; and marginal cost studies, used by about 20 states, which apportion costs in proportion to the incremental cost of adding service.

folds it in. The gap between when a utility's costs actually change and the time those changes are reflected in the rates customers pay is referred to as "regulatory lag." Regulatory lag becomes important during infrastructure construction: an investment, particularly multi-year projects, could be on the utility's books for years before any rate case can incorporate the investment into rates. Mechanisms have been created to mitigate this lag.[6]

There is a faster, narrower alternative to waiting for a full rate case, and these mechanisms exist largely as a response to regulatory lag. For certain categories of spending—such as fuel and purchased power costs (through fuel adjustment clauses), and capital investments in infrastructure, smart-grids, or storm hardening (through infrastructure "trackers" or riders)—some utilities can recover costs between rate cases through single-issue filings that bypass the lag for those specific line items. Consumer advocates have been wary of these faster mechanisms, as they tend to address rising costs in isolation, whereas declining costs are only revisited in a full rate case (Lazar, 2016).

2.3.3 Regulated versus Deregulated States

There are some differences in the way rate cases play out depending on whether the IOU is operating in a regulated or deregulated state. Regulated states have vertically integrated utilities that are responsible for generation, transmission and distribution of electricity. Because they are responsible for—and will make capital investments related to—all three, their rate base (and therefore their rate cases) can include investments in generation, transmission and distribution-related investments.

In contrast, IOUs operating in deregulated states are not vertically integrated, but instead are responsible for separate components, such as electricity distribution. This means distribution-only entities will make investments in the distribution-related infrastructure (e.g., substations, wires, smart meters, and operation and maintenance of that infrastructure) and these investments will comprise their rate base; however, generation-related investments would not show up in rate cases for such a utility. Distribution-only IOUs, however, buy their power from other entities and therefore are exposed to price changes in the energy and capacity markets so the rate design must allow them to recover any associated fluctuations in the market prices.

[6] One mechanism is an allowance for funds used during construction (AFUDC), which accrues the cost of equity and debt tied up in the project during the build and adds it to the rate base when the asset enters service. This allows the utility to eventually earn on those carrying costs, but only later. Another is construction work in progress (CWIP), which places unfinished investment into rate base during construction so the utility earns a current return before the plant is finished. This effectively eliminates more of the lag.

3 Conceptual Framework

The effect of data center growth on retail electricity prices is theoretically ambiguous. Large data centers increase electricity demand continuously and at high utilization rates, potentially requiring utilities or wholesale markets to procure additional generation capacity, expand transmission and distribution infrastructure, and undertake reliability-related investments. These costs may be especially high when large facilities arrive in discrete increments relative to existing system load or locate in regions with limited spare capacity. In such settings, data center entry may raise wholesale electricity prices, increase reserve requirements, accelerate infrastructure timelines, or require investments in substations, transmission upgrades, transformers, and interconnection infrastructure.

At the same time, data centers may put downward pressure on average rates. Electricity systems involve large fixed infrastructure costs, and new customers can improve cost recovery if they contribute substantially to utility revenues without proportionately increasing system costs. Because data centers operate with relatively stable around-the-clock demand, they may improve load factors, increase utilization of existing generation and network infrastructure, and allow fixed costs to be spread over a larger sales base.

Finally, pass-through to incumbent ratepayers may be attenuated if utilities or regulators require data centers to bear infrastructure costs directly through negotiated tariffs, interconnection agreements, minimum demand contracts, or upfront capital contributions. Whether retail prices rise therefore depends on the balance between new system costs, utilization gains, and cost-causation mechanisms that assign costs to the large-load customers generating them.

Utility Ownership Structure. These channels may differ substantially across electricity sector institutions. Investor-owned utilities may face stronger incentives to expand infrastructure following large load entry because capital investments enter the regulated rate base and generate authorized returns on equity. As a result, large load growth may induce substantial transmission, distribution, and generation investment that utilities subsequently recover through retail rates. IOUs may also face stronger incentives to accommodate rapid load growth because expanding the capital base can increase shareholder returns under cost-of-service regulation.

By contrast, public utilities and cooperatives operate under different financing structures and governance models, leading to different investment incentives and cost allocation decisions. Public utilities often finance infrastructure using tax-exempt municipal debt, potentially lowering financing costs and reducing pressure for aggressive rate-based cost recovery. Cooperatives may place greater weight on minimizing costs to existing customers rather than expanding infrastructure investment aggressively. Both public utilities and cooperatives may therefore differ from IOUs in the scale and timing of infrastructure expansion, the mechanisms through which costs are recovered, and the extent to which costs are passed through to incumbent ratepayers.

Regulatory Environment. The incidence may also differ across regulated and deregulated electricity markets. In deregulated states, utilities and consumers may be more directly exposed to changes in wholesale market conditions, including higher generation costs, congestion, and capacity market prices associated with demand growth. Because electricity procurement in these markets relies more heavily on wholesale market transactions rather than vertically integrated utility planning, increases in demand from large load customers may translate more quickly into higher retail prices. Rapid load growth may also place upward pressure on wholesale prices during periods of limited generation or transmission capacity, particularly in regions facing existing supply constraints.

In contrast, vertically integrated regulated utilities may recover costs primarily through rate base expansion and regulated cost recovery. These utilities may also have greater ability to coordinate long-term generation, transmission, and distribution planning internally, potentially smoothing short-run impacts associated with rapid demand growth. As a result, identical load shocks may generate different retail price responses depending on the regulatory structure and degree of exposure to wholesale electricity markets.

4 Data

4.1 Data Sources

Data Centers. Our primary data source is Aterio's *US Data Center Locations and Power Demand* product, a facility-level inventory of every active and proposed data center in the United States. As of April 2026, the dataset covers 5,937 facilities: 1,961 active, 735 under construction, and 3,241 announced.[7] For each facility, we observe location (GPS

[7] The unit of observation is the physical building, so a multi-tenant colocation facility shared by many operators is counted once. This convention differs from inventories (e.g., DataCenterMap.com) that count at the operator or tenant-deployment level and can list the same building multiple times. The reported number of active U.S. data centers therefore varies across sources: Bloomberg Terminal lists 1,640 as of 2025; the IM3 Open-Source Data Center Atlas lists 1,479 as of February 2026 (Mongird et al., 2026); Muller (2026) report

coordinates), operator, building size, estimated power capacity (MW), and activation year. Aterio assembles these records from press releases and company websites, high-resolution satellite imagery, utility and interconnection filings, and permitting records from local and federal agencies. Importantly, every record is manually validated by analysts.[8] For facilities without power capacity reported by operators, capacity is estimated by an ensemble machine-learning model that combines facility size, generator counts identified from satellite imagery, operator, location, and construction stage.[9]

Electricity Prices. We calculate the average retail price as total retail revenue divided by total retail sales, using EIA-861 data at the utility-by-state-year level. The data report total retail sales and a breakdown by customer class (residential, commercial, and industrial). For utilities in restructured states that report both bundled service (in which the utility provides both energy and delivery) and delivery-only service (in which customers buy energy from a competitive supplier), we use the bundled-service price.[10] All dollar variables are deflated to 2024 dollars using the CPI. We exclude utilities that ever file the EIA-861 Short Form, because the short form does not report sales by customer class.[11]

Electric Utility Service Territories, Characteristics, and Operational Profile. We supplement the price data with additional utility-level information from EIA-861, including each utility's service territory (the set of counties it serves), its balancing authority, its ownership category, and its total number of customers along with the breakdown by customer class. We group ownership into three categories: investor-owned utilities (IOUs), cooperatives, and public utilities (including federal, state, municipal, and political subdivision utilities).

2,465 as of 2025 based on S&P data; and DataCenterMap.com lists 4,286 total facilities across 1,837 operators.

[8] Analysts assign each building a construction-completion percentage based on visual benchmarks in satellite imagery—from site preparation (5–15%) to equipment installation and commissioning (70–95%)—and log each observation with a timestamped description and analyst initials. Generator counts and equipment types are cross-checked against permitting documents and manufacturer specifications. Aterio additionally reconciles aggregate figures (total capacity, facility counts by state, and operator footprints) against 10-Q filings, utility disclosures, and permitting records.

[9] The model achieves an out-of-sample R^2 of 0.94 for hyperscalers and 0.97 for non-hyperscalers, and reports an 80% confidence interval for each estimate.

[10] We cannot construct an all-in retail price for delivery-only customers from EIA-861 because the energy portion of their bill is paid to a competitive retail supplier, and EIA reports supplier sales only at the state-by-customer-class level rather than by the distribution utility serving each customer. The bundled-service price we use corresponds to the regulated default-service or standard-offer rate, which is set by the state public utility commission. In heavily restructured states, this price reflects a selected and shrinking subset of customers — those who never shopped, were returned to default service, or are explicitly on a regulated standard offer — rather than the market price faced by the typical customer.

[11] The Short Form is filed by utilities with retail sales below 100,000 MWh in the prior year. These are predominantly small municipal utilities and rural cooperatives; together they account for a small share of total U.S. retail electricity sales.

Utilities' operational profiles come from the EIA-861 operational data, which report summer peak demand, winter peak demand, and total energy disposition at the utility-year level. Because these variables are not reported separately by state, they capture each utility's system-wide operations. We define annual peak demand as the maximum of the utility's summer and winter peaks, measured in megawatts. We then construct the load factor as total energy disposition divided by the product of annual peak demand and the number of hours in the year. This measure captures average system load as a share of peak load: higher values indicate flatter and more continuous demand, and therefore greater utilization of existing capacity.

Infrastructure Investment. We collect annual utility-level infrastructure investment data from S&P Capital IQ Pro, which compiles FERC Form 1 filings. The panel covers 1988–2025 for investor-owned utilities (IOUs), including electric, diversified, and wholesale generation/transmission utilities.[12] For each utility-year, we construct three groups of outcome measures. The first captures investment in physical plant: gross additions to generators, transmission plant and overhead conductors, and distribution station equipment; electric plant in service; and construction work in progress (CWIP) for electric plant, which records spending on projects not yet in service.[13] The second group captures transmission spending in greater detail: total transmission operation and maintenance (O&M) expense, and expenditures on new and existing transmission lines. The third group captures how utilities finance and recover these investments: customer advances for construction (CAC), which are upfront cash contributions customers make toward infrastructure built to serve them, and accumulated deferred income taxes (ADIT). All dollar variables are deflated to 2024 dollars. To avoid double-counting, we drop parent utilities whose filings consolidate subsidiary operations and retain only subsidiary-level reports.[14]

Utility Rate Cases. We use data on resolved electricity rate cases from S&P Capital IQ Pro's Regulatory Research Associates (RRA) database, which tracks general base-rate proceedings before state public utility commissions. The data cover 2010–2025 and identify, for each case, the filing utility, jurisdiction, case type (e.g., vertically integrated, distribution-only), decision type (settled, fully litigated, or rate-case-by-rate-case), duration, and the dates of initial filing and final decision. For each case, we observe both the utility's requested values and the commission's authorized values for the change in revenue

[12]Municipal utilities and cooperatives are not required to file Form 1 and are therefore absent.

[13]We also use CWIP for common and other non-electric plant as placebo outcomes, since these categories are not electricity-specific and therefore should not respond directly to data center entry.

[14]The holding parent utility companies do not have EIA code.

requirement, the return on equity (ROE), the return on rate base, and the rate base itself. We restrict the sample to electric service cases and link each utility to our main dataset using the EIA utility code provided by RRA.[15] Because RRA covers only rate-regulated utilities, the resulting panel consists of investor-owned utilities. We collapse the case-level data to a utility-state-year panel by each case's filing year and measure rate-case filing activity in each year. The requested and authorized variables are defined only for utility-state-years with a case that reports them.

Large-Load Tariffs. To capture how utilities price and recover the cost of serving large new loads, we draw on the Halcyon Large Load Tariff Tracker (March 2026 release), which catalogs the special tariffs, contracts, riders, and frameworks that utilities have proposed or adopted for high-demand customers such as data centers. For each utility-state, we build an indicator for whether it ever has any large-load tariff, along with indicators for the cost-recovery provisions the tariff carries: a minimum monthly demand (take-or-pay) payment, an upfront contribution in aid of construction (CIAC), a long-term contract or demand requirement, and an exit or early-termination fee. We summarize a tariff's overall stringency with an indicator for a strong cost-recovery tariff, equal to one when at least three of these four core provisions are present, and we separately flag tariffs that explicitly name data centers or hyperscale facilities.

Additional Data. We measure each utility-state's exposure to fossil fuel generation using the generation mix of the balancing authority (BA) in which it operates. We construct BA-level fossil fuel shares using generator- and plant-level data from EIA-923 and EIA-860. We identify each generator's fuel type, net generation, and BA region, and then aggregate net generation to the BA-by-fuel-type level. We define a BA's fossil fuel share as the share of total net generation from fossil fuels, including coal, natural gas, oil, and petroleum coke. We average this measure over 2005–2007, a pre-sample window that ends before the 2008 financial crisis, to avoid capturing crisis-period shifts in generation patterns. We then assign each utility-state pair the fossil fuel share of the BA in which it operates.

Reserve margins are constructed from EIA-411 data over 2005–2007. It is defined at the NERC-region level as capacity resources minus net internal demand, divided by net internal demand and multiplied by 100. This measure captures available capacity above peak demand as a percentage of peak demand. Lower values therefore indicate tighter, more capacity-constrained grid conditions. We then map the NERC-region measure to the

[15] We exclude natural-gas-only cases and combined-utility cases for which the electric portion cannot be separately identified.

balancing-authority level. For balancing authorities that span multiple NERC regions, we use a generation-weighted average of the corresponding NERC-region reserve margins, with weights given by the balancing authority's generation share in each NERC region. Finally, we assign each utility-state pair the reserve margin of the balancing authority in which it operates.

We construct regional economic characteristics using data from two sources. County-level total and manufacturing gross domestic product (GDP), as well as population, come from the Bureau of Economic Analysis (BEA) Regional Economic Accounts. Counts of total and manufacturing establishments come from the U.S. Census Bureau's County Business Patterns. We average each county-level series over 2005–2007, a pre-sample window that ends before the 2008 financial crisis. We then aggregate these measures to the utility-state level across all counties in the utility-state's service territory. Using these aggregated measures, we construct three regional economic characteristics: total population, GDP per capita, and manufacturing establishments per capita.

4.2 Construction of Final Dataset for Analysis

We aggregate the data center records to the county-year level, computing the newly activated capacity (MW) and the count of newly activated facilities in each county and year. We drop county-years with newly activated capacity below 5 MW to exclude facilities too small to materially affect the grid. Results are qualitatively similar under alternative thresholds, with the estimated effect growing as the threshold rises. For each remaining county, we take the earliest entry year as the treatment start, and we also construct continuous measures of the cumulative facility count and cumulative capacity in each county at year-end.

We match this county-level entry information to utility-state-level electricity prices through utilities' service territories.[16] For each utility-state pair, we take the earliest data center entry year among the counties it serves in that state as the treatment start. The continuous measures—cumulative facility count and cumulative capacity—are constructed for each utility-state-year by summing the corresponding county-year values across all counties the pair serves in that state.

We restrict the all of the samples to the contiguous United States and to 2010–2024 for three reasons. First, activation years before 2005 are often imprecise: many early

[16]For example, if a data center enters Durham County, every utility-state pair whose service territory covers Durham County is treated. If a utility serves multiple states but only one sees data center entry, the pairs involving that utility in the non-entry states are controls.

entries are inferred rather than directly observed, and Aterio's satellite-based verification is unavailable for facilities built before high-resolution commercial imagery became widely available. Second, pre-2010 data centers differ substantially from those that drive the current load surge. Facilities built during the dot-com and early-enterprise era were mostly internet data centers hosting websites, email, and corporate ERP systems, with typical loads of kilowatts to a few megawatts. The hyperscale era—single-tenant campuses of hundreds of megawatts operated by cloud and content providers—begins around 2010–2011 with the first generation of purpose-built Facebook, Google, and AWS sites. Including pre-2010 facilities would mix two qualitatively different technologies and load profiles. Third, the 2010 start avoids confounding effects from the 2008–2009 financial crisis on both data center investment and electricity demand.

For the rate case data, we match data center entry to utility-state pairs using each utility's state-specific service territory, as in the main retail price analysis. For the infrastructure investment data, FERC Form 1 filings are not reported separately by state, so we match each county-level data center entry to each utility's combined service territory across all states it serves. We restrict the rate case and infrastructure investment samples to 2010–2024 and retain only utilities that appear in the main retail price sample. We merge the large-load tariff data to the retail price panel by utility-state. Utilities absent from the tracker are coded as not having a large-load tariff.

4.3 Descriptive Statistics

Our final main dataset covers 14,733 utility-state-year observations spanning 1,048 utility-state pairs and 973 distinct utility companies operating across 48 states.[17] Table 1 reports summary statistics for the key variables, organized into five panels. Variables in Panels A, B, and D are at the utility-state-year level, while those in Panels C and E are at the utility-year level.

The average retail price is \$0.13/kWh across all customer classes, highest for residential customers (\$0.14/kWh) and followed by commercial (\$0.13kWh) and industrial customers (\$0.10kWh). The number of observations after data center entry account for about 6% of utility-state-years. On average, the cumulative data center capacity is 5 MW unconditionally but 80.6 MW among the 908 utility-state-years that host at least one large data center, indicating substantial heterogeneity in exposure.

Panels C and D report summary statistics for variables on infrastructure investment

[17] Only 57 utilities serve customers in more than one state.

and rate cases, which only include IOUs. Infrastructure spending is dominated by transmission, with \$645 million per utility-year spent on existing lines and \$334 million in construction work in progress. As for rate cases, 27% of utility-state-years involve a rate-case filing, with general cases (23%) far more common than limited-issue rider cases (5%). The average authorized ROE is around 1% lower than the average requested ROE, and the average authorized rate base is \$3.42 billion. Finally, operational characteristics in Panel E show an average load factor of 0.56 and average peak demand of 440 MW.

Figure 1 maps cumulative active data center capacity by county between 2010 and 2026. The vast majority of counties host little or no capacity, and activity is highly concentrated in a small number of counties. The highest capacity counties cluster most heavily in northern Virginia, Texas, Arizona, and the Pacific Northwest. Appendix Figures A1 and A2 plot the expansion of data centers over time, separately by utility ownership and regulatory environment. Investor-owned utilities (IOUs) lead on both margins: both the share of utilities hosting at least one active data center and their average cumulative capacity. Data center expansion proceeds much more slowly for publicly-owned utilities and cooperatives. Across regulatory environments, average cumulative capacity rises faster in deregulated than in regulated states. While the trends start to diverge as early as 2014, the gap widens markedly around 2019 and then diverges even more toward the end of the sample.

5 Empirical Strategy

We aim to identify the causal impact of data center entry on retail electricity prices, exploiting quasi-experimental variation in the timing of large data center activations across utility service territories. We employ a difference-in-differences design that compares changes in outcomes over time between utility-state pairs whose service territories receive a large data center for the first time and pairs whose territories never receive a data center.

We begin with a standard two-way fixed effects (TWFE) specification that motivates our analysis:

$$Y_{it} = \beta \, \mathrm{PostDC}_{it} + \alpha_i + \delta_t + \varepsilon_{it}, \tag{1}$$

where Y_{it} is an outcome for utility-state pair i in year t—retail electricity prices, utility infrastructure investment, or rate case activity—and PostDC_{it} is an indicator equal to one in all years on or after the activation of the first large data center in pair i's service territory. Because the treatment is defined as first entry, we drop utility-state pairs whose first large data center was activated before our sample starts in 2010; for these pairs the pre-period is unobserved.

The pair fixed effects α_i absorb time-invariant differences across utility-state pairs—commission structure, generation mix, climate, baseline industrial composition, and persistent rate-level disparities—that could otherwise confound the relationship between data center entry and utility outcomes. The year fixed effects δ_t absorb shocks that hit all utilities in a given year, including wholesale fuel price movements, federal policy changes, macroeconomic conditions, and secular trends in electricity demand and generation technology. Identification of β therefore comes from within-pair variation after netting out common annual shocks. We cluster standard errors at the utility-state pair level, which is the level of treatment assignment and outcome measurement and allows for arbitrary serial correlation in ε_{it} driven by rate-setting cycles, regulatory decisions, and pair-specific cost recovery mechanisms.

Equation (1) is a useful baseline, but it is unlikely to deliver a clean causal estimand in our setting. Data centers activate in different years across territories, and the effects of entry on rates, capital spending, and rate case activity are likely to evolve over event time and differ across cohorts as utilities recover costs through subsequent rate proceedings. A growing literature documents that the TWFE estimator can be biased in settings with precisely this combination of staggered adoption and heterogeneous treatment effects (e.g., Baker et al., 2026). The bias arises because the TWFE coefficient is a weighted average of all available two-group, two-period comparisons, including comparisons that use *already-treated* pairs as controls for *later-treated* pairs. When treatment effects evolve over event time, the outcomes of the already-treated "controls" are still moving in response to their own past treatment, so some group-by-period average treatment effects enter the TWFE estimand with negative weights.

We therefore employ the difference-in-differences framework developed by Callaway and Sant'Anna (2021) (hereafter CSDID). The CSDID estimator constructs group–time average treatment effects:

$$\mathrm{ATT}(g,t) = \mathbb{E}\left[\, Y_{it}(g) - Y_{it}(\infty) \,\middle|\, G_i = g \,\right], \quad t \geq g, \tag{2}$$

where G_i denotes the year in which pair i first receives a large data center exposure, $Y_{it}(g)$ is the potential outcome under first treatment in year g, and $Y_{it}(\infty)$ is the never-treated potential outcome. Each $\mathrm{ATT}(g,t)$ compares the outcome change between $g-1$ and t for cohort g to the analogous change among never-treated pairs. Because comparisons are made within cohort and against pairs that remain untreated throughout the sample window, the estimator avoids the contaminated contrasts that drive TWFE bias under heterogeneous effects. We implement the improved doubly robust estimator, which combines an outcome-regression model for untreated potential outcome changes with an inverse-probability-weighting component whose propensity score is estimated by inverse probability tilting (Sant'Anna and Zhao, 2020). Treatment is the first year of large data center activation, and never-treated utility-state pairs form the comparison group.

We summarize the estimated $\mathrm{ATT}(g,t)$'s in two ways. First, we construct event-study estimates that trace the dynamic effects relative to the first year of activation:

$$\mathrm{ATT}(e) \;=\; \sum_{g} w_{g,e} \cdot \mathrm{ATT}(g,\, g+e), \tag{3}$$

where $e = t - g$ denotes event time and the weights $w_{g,e}$ give each cohort's share among the cohorts for which event time e is observed in the panel.[18] $\mathrm{ATT}(e)$ is therefore the average treatment effect e years after activation across all cohorts that contribute at that event time. We report event-study estimates spanning ten years before through twelve years after activation, with negative values of e serving as pre-treatment placebo checks on the parallel trends assumption.

Second, we report a single overall ATT that averages effects across all post-activation years and cohorts:

$$\mathrm{ATT} \;=\; \sum_{g} \sum_{t \geq g} w_g \cdot \mathrm{ATT}(g,\, t), \tag{4}$$

where w_g is proportional to the share of ever-treated units belonging to cohort g, normalized so that the weights sum to one across all post-treatment (g,t) cells.[19]

5.1 Identification Challenges

Non-Random Siting of Data Centers. Data centers do not locate at random. Hyperscalers select utility service territories based on electricity prices, tax incentives, etc.—factors that may independently affect the outcomes we study. As a result, utilities that host data centers may differ from those that do not along dimensions correlated with rate-setting behavior

[18]Specifically, $w_{g,e} = P(G = g \mid 1 \leq g + e \leq \mathcal{T}, G \leq \mathcal{T})$, i.e., the probability of belonging to cohort g conditional on being ever-treated and on the cohort contributing an observation at event time e. By construction, $\sum_g w_{g,e} = 1$ for each e.

[19]Specifically, $w_g = \mathbb{P}(G = g \mid G \leq \mathcal{T})/\kappa$, where $\mathcal{T}$ is the final period in the panel and $\kappa = \sum_g (\mathcal{T} - g + 1)\, \mathbb{P}(G = g \mid G \leq \mathcal{T})$ is a normalizing constant ensuring that $\sum_g \sum_{t \geq g} w_g = 1$.

and investment, including system size, generation portfolio, regulatory environment, and the industrial composition of load. We address this concern in four ways.

Within-Pair Variation. Our DID design exploits within utility-state pair variation over time, so any time-invariant differences across pairs—commission structure, fuel mix, baseline industrial load, persistent rate-level disparities—are absorbed by the pair fixed effects. Year fixed effects further absorb common shocks that move all utilities together, such as wholesale fuel price movements and federal policy changes.

Sample Restrictions. To make treated and comparison pairs more comparable, we impose several restrictions when constructing the analysis sample. We drop utilities that file the EIA-861 short form, which captures small distribution utilities that differ systematically from the full-form filers in size, customer composition, and reporting detail. We keep only utility-state pairs with a retail component, and further restrict to pairs serving all three primary customer classes—residential, commercial, and industrial—so that the cross-class composition of load does not vary mechanically across pairs.

Region- and Pair-Specific Time Trends. Another concern is that time-varying shocks—renewable portfolio standards, state-level electricity market reforms, regional grid conditions, or differential demand growth—may correlate with both data center siting and our outcomes. We address this in robustness checks by augmenting the specification with utility-specific and grid-region-specific trends. We also estimate specifications that interact year fixed effects with pre-period utility-state characteristics, including customer base, average retail price, and total sales, as well as pre-period regional characteristics, including manufacturing establishments per capita, population, and GDP per capita. These specifications allow utilities and regions with different baseline profiles to follow different time paths. Our estimates remain similar across these alternative specifications.

Matched Difference-in-Differences. We also conduct a robustness check to directly improve the balance between treated and never-treated utility-state pairs using entropy balancing. This procedure re-weights the never-treated comparison group to match treated pairs on pre-period characteristics, including average retail price, customer count, manufacturing establishments per capita, and GDP per capita. We then apply the Callaway and Sant'Anna (2021) estimator to the weighted sample. The estimated effects after re-weighting become generally larger.

Parallel Trends Assumption. The key identifying assumption of a DID design is parallel trends: in the absence of large data center entry, outcomes in treated utility-state pairs would have evolved in parallel with outcomes in never-treated pairs. While this assumption is not directly testable, we assess its plausibility from the dynamic pattern of effects in the years leading up to data center activation. The pre-treatment event-study coefficients (ATT(e) for $e < -1$) are small and statistically indistinguishable from zero across all main outcomes, indicating no meaningful differential trends between treated and never-treated pairs prior to the first large data center activation.

SUTVA. Another threat to our DID design is a potential violation of SUTVA. The entry of a data center in one service territory may affect how the utility allocates costs and sets rates for customers in its other service territories. To partially mitigate this concern, we conduct the analysis at the utility-state level. The networked structure of the electricity grid, however, generates two further spillover channels.

First, in regions with wholesale markets, data center entry might raise wholesale prices, which can ultimately pass through to retail rates. Because multiple utilities participate in the same wholesale market, the control group is exposed to the same price shock at the wholesale market level. This is not a first-order concern for our analysis, though, since this would mean that our estimated treatment effect is likely attenuated, if anything.

Second, utilities that operate in multiple states but receive data centers in only a subset of those states may allocate the resulting infrastructure costs beyond the states of direct entry—particularly transmission costs allocated through FERC-approved regional tariffs.[20] A utility-state pair without data-center entry is therefore a problematic control if the same utility experiences data-center entry in another state, because cost allocation across states implies the "control" pair is in fact partially treated. In our sample, only 14 of 112 utilities are exposed to this channel. As a robustness check, we exclude from the control group any utility-state pair without data-center entry whose utility experiences data-center entry in a different state; our baseline price estimates are essentially unchanged.

6 Main Results

6.1 Retail Prices and Customer Class Incidence

We begin by examining the average effect of data center entry on retail electricity prices. Figure 2 presents event study estimates tracing the evolution of electricity rates before and after initial data center activation, and Table 2 reports corresponding average treatment

[20]Cross-state cost allocation is governed by FERC's regional transmission planning rules (Orders 1000 and 1920) and, for shared corporate costs at multi-state holding companies, by 18 CFR Part 367.

effects. Two main findings emerge. First, retail electricity rates increase following data center entry. Average retail prices rise by approximately 2.7%, with little evidence of systematic differential pre-trends prior to treatment. The effects grow gradually over time, consistent with utilities responding to large load growth through processes that unfold over several years—such as infrastructure investment, cost recovery, and regulatory adjustment—rather than immediately upon entry.

Second, these effects differ across customer classes. Residential prices increase by 2.1%, commercial prices by 2.8%, and industrial prices by 4.2%. This heterogeneity is central to understanding the incidence of growth because average retail rates mask important differences in how the costs associated with serving new demand are allocated across ratepayers. The positive residential effect indicates that at least some of the costs of data center growth are passed through to incumbent households. The larger effects for commercial and (particularly) industrial customers also suggest that non-residential customers might bear a larger share of the price increase, although the pass-through to incumbent commercial and industrial ratepayers should be interpreted with care: depending on utility tariff classifications, data centers themselves may be recorded as commercial, industrial, or special large-load customers. As a result, the commercial and industrial effects may reflect both pass-through to incumbent firms and the prices paid by the new large-load customers generating the demand shock.[21]

Together, these results show that data center expansion places upward pressure on retail electricity prices and that the burden of those increases is not distributed uniformly across customer classes. The remainder of the results section examines why these price effects differ across institutional environments and how utility ownership, regulatory structure, infrastructure investment, and cost-recovery processes shape who ultimately pays for large-load growth.

[21] Electric utilities report retail sales to the EIA by customer sector (residential, commercial, industrial, transportation) on Form EIA-861, but classify individual accounts according to their own rate-schedule definitions—generally a function of voltage and demand size—rather than a uniform end-use rule. Consequently, a data center may be reported as commercial or industrial depending on the serving utility's tariff structure, and a growing number of utilities now place large loads on dedicated large-load rate schedules. Starting with data year 2022, the customer-sector definitions accompanying Schedule 4 (Sales to Ultimate Customers) of the Form EIA-861 instructions explicitly list "crypto mining and data centers" within the commercial sector, while the industrial sector covers manufacturing, construction, mining, agriculture, fishing, and forestry. The EIA nonetheless acknowledges that because classification follows rate class rather than building activity, a utility "may designate a large energy-intensive office or data center as an industrial customer." See *Form EIA-861 Instructions*, U.S. Energy Information Administration, https://www.eia.gov/survey/form/eia_861/instructions.pdf; and U.S. Energy Information Administration, *Commercial Buildings Energy Consumption Survey (CBECS) Frequently Asked Questions*, https://www.eia.gov/consumption/commercial/faq.php.

6.2 Institutional Heterogeneity

We now turn to examining whether the price responses following data center entry differ across electricity sector institutions. As discussed in Section 3, utilities differ in their ownership structures and regulatory environments, which determine financing arrangements, incentives, and cost-recovery strategies. The distribution of data center exposure also differs across institutions (see Appendix Figures A1 and A2). Institutional differences may therefore shape both how utilities respond to large-load growth and how the associated costs are allocated across ratepayers.

6.2.1 Utility Ownership

Investor-owned utilities (IOUs), public utilities, and cooperatives differ in how infrastructure investments are financed, how costs are recovered, and how costs are allocated across customers. We therefore examine whether the retail price effects of data center entry vary systematically across ownership types.

Figure 3 presents ownership-specific event-study estimates, while Table 3 reports the corresponding average treatment effects. The results reveal striking differences. Among IOUs, data center entry increases average retail prices by 5.6%, more than twice the magnitude of the baseline effect in Table 2. The effects are statistically significant across all customer classes for IOUs: residential prices increase by 5.2%, commercial prices by 4.6%, and industrial prices by 8%. The event study estimates show little evidence of differential pre-trends and indicate that these effects emerge gradually over time, consistent with price responses operating through infrastructure investment, regulatory cost recovery, and rate-setting processes.

In contrast, the effects are much smaller for public utilities and absent for cooperatives. Public utilities exhibit a modest increase in average retail prices (2.4%), but the customer class estimates are generally smaller and statistically indistinguishable from zero. Cooperatives display no systematic price response following data center entry. These differences are consistent with ownership structure shaping the incidence of large-load growth. IOUs are for-profit utilities that recover capital investments through regulated rates and may have stronger incentives to expand and recover infrastructure costs through formal cost-recovery processes. Public utilities and cooperatives, on the other hand, are consumer- or member-owned and may face different financing costs, governance constraints, and incentives for passing costs through to incumbent ratepayers.

6.2.2 Regulatory Environment

We next examine whether the retail price effects differ between utilities operating in states that have deregulated electricity generation and those in fully regulated states. This distinction is central because deregulation changes the institutional channel through which large-load growth affects costs. In regulated states, vertically integrated utilities often own or procure generation, transmission, and distribution through integrated planning and recover costs through regulated rate proceedings. In deregulated states, generation and delivery are separated, so utilities and retail customers may be more directly exposed to wholesale procurement, capacity-market, and default-service costs, while distribution utilities continue to recover regulated network costs through rates.

Table 4 reports average treatment effects separately for deregulated and regulated states. Overall, retail price effects are substantially larger in deregulated markets. Following data center entry, average retail prices increase by 3.8% in deregulated states, compared with a smaller and statistically insignificant 1.5% increase in regulated states. Residential prices increase by 4% in deregulated states but do not increase in regulated states. Commercial prices increase in both settings, while industrial price effects are larger in deregulated states but less precisely estimated.

These patterns suggest that regulatory environment plays an important role in determining whether data center-driven load growth is passed through to retail customers. In deregulated states, large-load growth may translate more directly into retail prices through wholesale procurement costs, capacity obligations, rider mechanisms, and other cost-recovery channels. In regulated states, vertically integrated utilities may have greater ability to absorb or smooth the effects of new load through integrated planning, owned generation, fixed-cost spreading, and gradual rate recovery.

Appendix Table A1 examines the two-way split between ownership and regulatory environment. These results show that ownership is the first-order source of heterogeneity: price effects are concentrated among IOUs in both deregulated and regulated states, while public utilities and cooperatives show little systematic response in either environment. Within IOUs, effects are generally larger in deregulated states, especially for average and residential prices, but IOUs in regulated states also experience meaningful price increases. This pattern suggests that IOU ownership is central to retail pass-through, while the regulatory environment shapes the magnitude and mechanisms of that pass-through. These findings motivate the mechanism analysis in Section 7.

6.3 Magnitude of Effects

To put these results in perspective, we translate the percentage change estimates into changes in electricity bills using pre-treatment electricity prices and sales for treated utilities in each estimation sample. The average residential effect is economically modest in monthly bill terms: a 2.1% increase in residential rates corresponds to an increase of about 0.32 cents per kWh, or roughly $39 per year ($3.21 per month) for the average residential customer. The magnitudes are larger in the institutional settings where the price effects are concentrated. Among IOUs—where average residential prices are also already higher—the 5.2% residential effect corresponds to an increase of about 0.88 cents per kWh, or roughly $88 per year ($7.34 per month). In deregulated states, the 4% residential effect corresponds to about $71 per year ($5.94 per month).

These estimates are smaller than some of the bill increases discussed in recent public debates around data centers, but an important caveat to keep in mind is that our estimates are for realized average retail price effects over our sample period (data centers that enter between 2010 and 2024). They are not forward-looking projections, and they do not capture some of the most recent hyper-scale developments or high-growth hotspot regions.[22]

At the same time, our findings do not imply that the incidence is negligible. Even modest residential rate increases can matter for households already facing high energy burdens, and the level dollar effects are substantially larger for commercial and industrial customers because of their higher electricity consumption. In the full sample, the estimated effects imply average annual bill increases of roughly $225 for commercial customers and about $52,000 for industrial customers, with larger industrial effects among IOUs. We therefore interpret the estimates as evidence of measurable early pass-through from data center growth, rather than as an upper bound on the potential ratepayer consequences of future large-load expansion.

6.4 Robustness Checks

Addressing Confounding Factors. We conduct a battery of robustness checks to address the concern that data centers may locate non-randomly in regions or utility service territories with different underlying price trajectories. Panels A–H of Table A2 report estimates from specifications that allow for these potential confounders in increasingly flexible ways.

[22] Some utility-states with first data center entry before 2010 are also areas with large or many data centers (e.g., some utility-states in Virginia).

Panels A and B add utility-specific and grid-region-specific trends, respectively.[23] These controls absorb differential price paths across utilities and grid regions, including trends associated with regional market structure, transmission planning, and reliability conditions. Panels C–E interact year fixed effects with baseline utility-state characteristics: average retail price, customer count, and total sales. These interactions allow utilities that differ in initial price levels, scale, and load served to follow different time paths. Panels F–H analogously interact year fixed effects with pre-period regional characteristics: manufacturing establishments per capita, population, and GDP per capita. These controls flexibly account for differential price evolution across areas with different industrial composition, market size, and income levels. Across all specifications, the estimated effects remain qualitatively similar, suggesting that the baseline results are not driven by observable differences in pre-existing utility, grid-region, or local economic trends.

Matched DID Design. To improve the balance and comparability of treated and untreated utility-state pairs, we supplement the baseline model with a matched difference-in-differences estimator based on entropy balancing. The baseline estimates compare utility-state pairs that eventually serve data centers to those that never do. Although the Callaway and Sant'Anna (2021) estimator does not require treated and control units to be similar in levels, it does require parallel trends. Large differences in pre-period characteristics may therefore raise concerns that treated and never-treated utilities are structurally different in ways that could also be related to subsequent price dynamics. Panel A of Table A3 shows that such differences are present in the raw sample: before any data center entry, treated utilities served about 20,000 more customers, charged retail prices that were \$0.010/kWh higher, and operated in counties with roughly \$7,500 higher GDP per capita than never-treated utilities. All three differences are statistically significant at the 1% level.

We construct a more comparable control group using entropy balancing (Hainmueller, 2012). Entropy balancing assigns weights to never-treated utility-state pairs so that the weighted means of selected pre-period covariates among never-treated pairs exactly match the corresponding means among treated pairs, while keeping the weights as close as

[23] We assign each utility-state to a grid region as follows. For utilities in organized wholesale markets, the region is the relevant ISO/RTO: CAISO, ERCOT, ISO-NE, MISO, NYISO, PJM, or SPP. Utilities outside organized wholesale markets are assigned to their NERC reliability region: FRCC, SERC, or WECC. Utility-specific trends are specified as linear trends. Grid-region-specific trends are specified as third-order polynomial trends. We do not use higher-order polynomial trends at the utility level because they would absorb too much of the identifying variation.

possible to uniform.[24] We balance on four pre-period characteristics measured over 2005–2007: average retail price, number of customers served, manufacturing establishments per capita, and GDP per capita. Because these characteristics are time-invariant pre-period measures, the balancing step is implemented once at the utility-state level. Treated utility-state pairs receive a weight of one. We then estimate the average treatment effect on the treated using the Callaway and Sant'Anna (2021) estimator on the weighted sample. Panel B of Table A3 confirms that the reweighting achieves balance. After entropy balancing, none of the four covariates differs significantly between treated and never-treated utility-state pairs. The point estimates are also close to zero: the pre-period price gap falls from \$0.010/kWh to essentially zero, and the differences in customer count and GDP per capita also shrink dramatically.

Figure A3 plots the event-study coefficients estimated from the weighted sample. Table A4 reports the corresponding overall matched difference-in-differences estimates. Reweighting does not overturn the baseline result; if anything, the estimates are generally larger. This pattern suggests that observable pre-period imbalances likely attenuate the baseline estimates rather than inflate them.

Excluding Early or Late Treated Units. We next address the concern that the estimates may be sensitive to treated units observed for only a short window around entry. In Panel H of Table A2, we restrict the treated sample to utility-states whose first data center enters between 2012 and 2022. This restriction ensures that every treated unit is observed for at least two years before and after entry.

A related issue arises in the heterogeneity analysis by ownership and regulatory environment. Figures A1 and A2 show that cumulative data center capacity remains negligible for cooperatives, publicly owned utilities, and utilities in regulated states until the mid-2010s, increasing appreciably only around 2017–2018. To rule out the possibility that the heterogeneous effects are driven by this timing pattern, Table A5 drops treated utility-state pairs with first entry before 2016 and re-estimates the ownership and regulatory splits. The results are similar.

Addressing Spillovers. We examine whether cross-state cost allocation contaminates the control group. The baseline design treats a utility-state without a data center as untreated. For utilities operating in multiple states, however, data center entry in one state may

[24] Relative to nearest-neighbor or propensity-score matching, this approach discards no comparison units and achieves exact balance on the targeted moments, reducing researcher discretion over calipers and match ratios while preserving efficiency.

affect rates in other states served by the same utility, particularly if transmission or other infrastructure costs are recovered across the utility's broader footprint. Such spillovers would attenuate the estimated treatment effect by raising prices in untreated utility-states. Panels I and J of Table A2 address this concern. Panel I excludes multi-state utilities exposed to data center entry in only a subset of their service states. Panel J excludes multi-state utilities altogether, identifying the effect solely from single-state utilities.

Alternative Standard Error Clustering. Finally, we assess the sensitivity of inference to alternative clustering assumptions. The baseline specifications cluster standard errors at the utility-state level. Because data center siting may also be correlated within states or within utilities, Panels K and L of Table A2 report standard errors clustered by state and by utility, respectively. The estimated standard errors remain similar to our baseline.

7 Mechanisms

The results thus far show that data center entry increases retail electricity prices, but that these effects are driven by IOUs, with larger effects among utilities operating in restructured electricity markets.

We now investigate the mechanisms underlying these patterns. We organize the analysis around three related questions. First, does data center entry generate real system costs for utilities, and if so, how do utilities respond through infrastructure investment and system operations? Second, through which regulatory and tariff mechanisms are those costs recovered, and do cost-causation institutions attenuate pass-through? Third, to what extent do system constraints and wholesale market exposure contribute to price increases? These questions map directly into the conceptual framework in Section 3: data centers may raise costs by requiring new generation, transmission, distribution, and reliability expenditures, but they may also improve system utilization and expand the sales base over which fixed costs are recovered. Whether retail prices rise depends not only on the size of the cost shock, but also on how utilities absorb, finance, and recover those costs.

The analysis in this section focuses on IOUs given how the price effects documented above are concentrated among IOUs, making them the relevant ownership class for understanding the pass-through mechanism. The detailed investment and regulatory data used in this section also come primarily from FERC Form 1 filings and state rate-case proceedings, both of which apply most directly to IOUs. Within IOUs, we also emphasize differences between utilities operating in regulated and deregulated states, since the preceding results show that retail price pass-through is larger in deregulated markets.

7.1 Infrastructure and Operational Responses

We begin by examining whether data center entry is associated with changes in utility infrastructure investments and operations. Large data centers can affect utility costs through several channels. They may require new generation capacity, investments in transmission and distribution infrastructure, upgrades to existing lines and substations, or additional planning and reliability expenditures. At the same time, because data centers often operate at high utilization rates, they may improve system load factors and allow utilities to spread fixed costs over a larger sales base. The infrastructure and operational outcomes therefore help distinguish between two competing forces: data centers may raise costs by straining the system, but they may also lower average costs by improving utilization of existing assets.

7.1.1 Infrastructure Investments

Figures 4 and 5 present event study estimates for the main infrastructure outcomes for all utilities on average, and Appendix Figures A5 and A6 provide estimates for utilities operating in deregulated and regulated states separately. Appendix Table A6 reports corresponding average treatment effects for all IOUs and separately for those operating in regulated and deregulated states.

Several patterns emerge. First, data center entry is associated with an increase in electric construction work in progress (CWIP), which captures capital projects that are under construction but not yet fully placed in service. The increase in electric CWIP therefore suggests that data center entry is associated with new electricity-related capital activity before those assets fully enter the operating system. By contrast, Appendix Figure A4 shows little evidence of increases in "common" or "other" CWIP. Common CWIP captures construction projects associated with shared utility assets, such as general plant, administrative facilities, information systems, or other assets used across utility functions rather than assigned directly to electric service. Other CWIP captures construction activity associated with non-electric utility operations outside the electric system. The absence of effects for these categories suggests that the response is specific to electric utility infrastructure rather than reflecting a broader expansion in utility activity.

Moreover, the average increase in electric CWIP is driven by utilities operating in deregulated states. This pattern suggests that utilities in deregulated states begin electricity-related capital projects in anticipation of load growth. For utilities in regulated states, we do not find a comparable average CWIP response, although the event-study estimates suggest that some CWIP response emerges several years after activation. This timing is consistent with a different adjustment margin: regulated utilities may initially accommodate new load through existing capacity, upgrades to current assets, and longer-run planning, with larger capital responses appearing later.

Second, we find evidence of investment in generation capacity. Generator additions increase following data center entry, and this effect is driven entirely by utilities operating in *regulated* states. This pattern is consistent with utilities responding to persistent load growth by expanding generation resources when generation assets are part of the utility's existing business model and can enter the regulated asset base. The timing is also informative. In regulated states, the generator response appears only a few years after data center activation, consistent with the time required for planning, approval, construction, and commissioning. This slightly delayed response suggests that some capital adjustments occur only after new load is realized and incorporated into longer-run system planning.

Third, the evidence for broad-based expansion across all network asset categories is more limited. We do not find robust average increases in distribution station additions, overhead conductor additions, total transmission plant additions, or new transmission line expenditures. This does not imply that data centers impose no network costs. Rather, it suggests that the immediate response is not primarily the construction of entirely new transmission corridors or uniform expansion across all asset categories. Instead, the clearest network response appears in expenditures on existing transmission lines and in transmission operation and maintenance (O&M). Existing transmission line expenditures rise after data center activation, with an immediate but gradual increase in regulated states and a more delayed increase in deregulated states. This is consistent with utilities reinforcing, upgrading, or maintaining existing transmission assets as large new loads are integrated into the grid.

Transmission O&M provides a complementary view of the timing of the adjustment process. O&M expenditures begin rising a few years before data center activation, especially for utilities in deregulated states, and increase again after facilities become operational. This pattern is consistent with utilities undertaking interconnection studies, engineering assessments, reliability planning, and operational preparation in anticipation of new load. Unlike capital additions, these expenditures need not wait until assets are completed and placed in service. The O&M response therefore suggests a multi-stage adjustment process: utilities first engage in planning and system preparation, while larger capital investments appear later as demand materializes and infrastructure projects are completed.

7.1.2 Operational Responses

We also examine two operational measures: peak demand and load factor (see Figure 6). Peak demand captures the maximum load the system must be able to serve, and therefore speaks to capacity and reliability needs. Load factor, defined as average load relative to peak demand, captures how intensively the system is utilized. These measures help distinguish between two possible responses to data center entry. If data centers increase peak demand without improving load factor, new load is more likely to create capacity pressure and raise average costs. If, instead, data centers increase average load more than peak demand, load factors may improve, allowing fixed costs to be spread over more kilowatt-hours and attenuating price effects.

The regulated-versus-deregulated split reveals different adjustment margins. In deregulated states, the most salient infrastructure response is the increase in electric CWIP, while we find little evidence of improved system utilization (i.e., no change in load factor). This combination suggests that utilities in deregulated markets face construction-related cost pressure without a corresponding utilization gain that would spread fixed costs over a larger sales base. In regulated states, by contrast, load factors improve. Combined with the evidence on transmission O&M, generation investment, delayed increases in completed electric plants, this pattern is consistent with slower-moving, planned capital responses that enter the system with a lag and with partial absorption of new load through improved utilization.

Appendix Figures A8 and A9 provide additional evidence on this timing. In deregulated states, electric plants in service start to rise gradually in the couple of years leading up to data center activation and then continue to rise, whereas the response for utilities in regulated states is delayed (the increase does not occur until six years after activation). This is consistent with capital projects moving from planning or construction into completed utility plant only after they are "used and useful" and suggests that utilities in deregulated states may start building in anticipation of new load. At the same time, electric accumulated deferred income taxes (ADIT) rise in regulated states immediately after data center entry, whereas we find no such effect on ADIT for utilities in deregulated states. We interpret the ADIT response cautiously, since ADIT is an accounting object linked to tax timing and depreciation and can offset rate base when recorded as a liability. Nevertheless, the pattern is consistent with a broader increase in electric utility plant and capital accounting activity *following* data center entry for utilities in regulated states.

The load factor and peak demand findings in Figure 6 confirm this distinction. In regulated states, load factors begin to rise a few years after data center entry, while peak demand also trends upward. The fact that load factors increase despite higher peak demand indicates that average load is rising even more, consistent with data centers improving system utilization. This provides a plausible explanation for why regulated states can experience infrastructure and rate-base responses without contemporaneous retail price increases. In deregulated states, by contrast, we do not observe a comparable improvement in load factors, suggesting that new load generates cost pressure without the same fixed-cost-spreading offset.

Taken together, these results show that data center entry induces real infrastructure and operational responses among IOUs, but that the nature of the response differs sharply by regulatory environment. In regulated states, utilities appear to accommodate new load through a combination of improved system utilization, transmission O&M, generation investment, and delayed increases in completed electric plants. These responses are consistent with absorption through planning and fixed-cost spreading, which helps explain the muted contemporaneous price effects in regulated markets. In deregulated states, utilities exhibit larger increases in electric CWIP and no comparable improvement in load factors, suggesting that new load generates more immediate cost pressure without the same utilization offset. We next examine how these costs enter the regulatory cost-recovery process.

7.2 Cost Recovery Strategies

The preceding results show that data center entry is associated with real infrastructure and operational responses among IOUs. However, infrastructure investment does not mechanically translate into higher retail electricity prices. Utilities recover costs through regulatory and tariff-setting processes that determine which costs enter the revenue requirement, when they are recovered, and how they are allocated across customers. We therefore next examine rate case activity, requested and authorized cost recovery, customer advances for construction, and large-load tariffs.

7.2.1 Rate Cases and Authorized Cost Recovery

Rate cases can affect retail prices through several margins. More frequent case filings indicate greater regulatory engagement after data center entry, while increases in rider proceedings suggest the use of targeted cost-recovery channels that can operate outside a full general rate case. Requested revenue requirements capture the additional annual revenues utilities seek in their initial filings, while authorized revenue requirements capture what regulators ultimately approve. Authorized rate base, by contrast, captures the stock of capital recognized as prudent, used and useful, and eligible to earn a return. Distinguishing these objects is important: a larger authorized rate base need not imply a proportional increase in the measured rate-change amount if higher sales, utilization gains, tax and depreciation adjustments, customer contributions, or other components of the revenue requirement offset the annual revenue deficiency.

Table 5 reports effects on rate case activity. On average across IOUs in all states, data center entry does not significantly affect rate case activity, but this masks substantial heterogeneity by regulatory environment. Among utilities operating in deregulated states, data center entry increases the probability of filing any rate case, the probability of filing a general rate case, the probability of filing a rider case, and the total number of cases. Case duration also declines. These patterns suggest that utilities in deregulated markets respond to data center entry by engaging regulatory cost-recovery processes more frequently and that they move more quickly. The increase in rider activity is especially consistent with targeted recovery of specific infrastructure or system costs. By contrast, we find no evidence of increased rate case activity among IOUs in regulated states.

Table 6 examines requested and authorized cost recovery. Data center entry does not significantly increase requested revenue requirements on average, although requested revenue increases substantially among utilities in deregulated states. Authorized revenue requirements, however, do not increase correspondingly. This distinction suggests that utilities in deregulated markets perceive data center entry as creating additional revenue needs, but regulators do not approve proportional increases in the measured authorized rate-change amount.

The clearest result in Table 6 is instead the increase in authorized rate base. Authorized rate base rises substantially on average and in both regulated and deregulated states, indicating that regulators are recognizing additional new capital as prudent, used and useful, and eligible for cost recovery. This finding links the infrastructure evidence in Section 7.1 to the regulatory process: data center growth is associated not only with capital activity on utilities' books, but also with larger regulatory asset bases in rate proceedings. At the same time, the absence of a corresponding increase in authorized revenue requirements implies that rate-base growth is not passed through one-for-one into measured annual revenue increases. In regulated states, this is consistent with the preceding evidence of improved load factors and fixed-cost spreading. In deregulated states, the increase in authorized rate base occurs alongside increases in CWIP, rate case activity, rider filings, and requested revenues, suggesting a more active cost-recovery response even when authorized revenue increases are moderated in the formal proceedings.

The remaining rate case outcomes clarify the nature of this response. We find little evidence that data center entry increases authorized returns on equity or authorized returns on rate base. Thus, the price effects documented above do not appear to be driven by regulators granting utilities higher allowed returns per dollar of capital. Instead, the mechanism appears to operate through a larger approved capital base and more active cost-recovery proceedings, rather than through higher authorized rates of return.

7.2.2 Cost-Causation Mechanisms

We next examine whether utilities use mechanisms designed to assign infrastructure costs more directly to the large-load customers that create them. Appendix Table A9 and Appendix Figure A7 report effects on customer advances for construction. Customer advances are upfront payments or contributions from customers to finance infrastructure needed to serve them, and therefore provide a direct cost-causation mechanism. We find that data center entry increases customer advances for construction on average, with the effect concentrated among utilities operating in deregulated states. This suggests that utilities in deregulated markets attempt to recover some data-center-related infrastructure costs directly from large-load customers. However, because these same utilities also experience retail price increases, customer advances do not appear to fully eliminate pass-through to incumbent ratepayers.

Finally, Appendix Tables A10 and A11 examine large-load tariffs. These results should be interpreted cautiously and are necessarily suggestive because large-load tariffs are observed only in very recent years (i.e., from about 2022 onward) and adoption is not randomly assigned. Nonetheless, they provide useful descriptive evidence. Utilities that eventually implement large-load tariffs do not exhibit the same retail price increases as utilities that never implement such tariffs during our sample period. This pattern is consistent with the idea that tariff provisions targeted at large-load customers—such as minimum demand commitments, long-term contract requirements, infrastructure contributions, or exit fees—might attenuate pass-through to incumbent ratepayers. Large-load tariffs appear to be somewhat more common in regulated states and are more likely to include long-term or contract demand requirements and contribution-in-aid-of-construction provisions in regulated states, although these differences are small. That said, these patterns begin to suggest that cost-causation institutions may help explain why regulated states exhibit weaker retail price pass-through despite evidence of infrastructure and rate-base responses, warranting further investigation in future work.

Taken together, the cost-recovery evidence shows that data center entry is associated with larger authorized rate bases in both regulated and deregulated states, but only deregulated states exhibit the broader set of responses associated with faster pass-through: more rate case activity, more rider proceedings, higher requested revenue requirements, increased customer advances for construction, and retail price increases. Regulated utilities, by contrast, show rate-base growth without comparable increases in rate case activity or average retail prices, consistent with greater absorption through utilization gains, planning, and more gradual cost recovery. These results suggest that who pays for data center growth depends not only on whether utilities invest, but also on the regulatory and tariff mechanisms through which those costs are recovered and allocated.

7.3 System Constraints and Wholesale Cost Exposure

Infrastructure investment and regulatory cost recovery are not the only channels through which data centers may affect retail prices. The cost of serving new load also depends on pre-existing system conditions. Utilities in constrained systems may face higher generation, procurement, congestion, capacity, and reliability costs when large loads enter. These channels may be especially important in deregulated markets, where generation and delivery are separated and retail prices may be more directly exposed to wholesale energy, capacity, and default service procurement costs.[25]

We examine this possibility using two measures of baseline system conditions: the fossil share of generation in the grid region and reserve margins. Fossil-heavy grids may be more exposed to fuel-cost and marginal-generation-cost pass-through, while low reserve margins indicate tighter systems with less spare capacity. If data center load raises prices by increasing the cost of serving demand in constrained systems, price effects should be larger in fossil-heavy or low-reserve-margin regions, particularly in deregulated markets.

Table A7 reports heterogeneity by baseline fossil generation share. In the full sample, price effects are relatively similar across high- and low-fossil regions, suggesting that

[25]Regulated utilities can also face wholesale and fuel cost exposure, and deregulated utilities continue to recover regulated transmission and distribution costs through rates. The distinction is therefore one of relative exposure: vertically integrated utilities may manage load growth through owned generation, integrated resource planning, long-term procurement, and gradual cost recovery, whereas restructured markets rely more heavily on wholesale procurement and market-based capacity mechanisms.

fossil exposure is not the dominant mechanism overall. In deregulated states, however, price effects are generally larger in fossil-heavy regions, especially for average, residential, and commercial prices. This pattern is consistent with greater exposure to wholesale and marginal generation costs in deregulated markets. In regulated states, the fossil-share split reveals little systematic heterogeneity.

Table A8 provides stronger evidence that system constraints matter. Price effects are concentrated in low-reserve-margin regions, particularly in deregulated states. In these markets, data center entry increases retail prices in areas with limited spare capacity, while effects in high-reserve-margin areas are smaller and less precise. This pattern is consistent with a capacity adequacy mechanism: when systems are already tight, large new loads are more likely to require expensive generation procurement, capacity payments, reliability expenditures, or accelerated infrastructure investment, and these costs may be passed through more quickly in deregulated markets.

These results help interpret the regulated-versus-deregulated differences documented above. In regulated states, price effects vary less systematically with reserve margins, consistent with vertically integrated utilities smoothing system constraints through integrated planning, owned generation, long-term procurement, and gradual rate recovery. In deregulated states, stronger price effects in low-reserve-margin and fossil-heavy regions, together with increased rate case and rider activity and no improvement in load factors, point to a more direct pass-through channel.

Taken together, the system-constraint results reinforce the broader mechanism evidence. Data centers generate infrastructure and capital-accounting responses in both regulated and deregulated settings, but retail price increases are larger in deregulated markets, especially where systems are constrained. The incidence of data center growth therefore depends not only on whether new infrastructure is needed, but also on the system conditions and market institutions through which the costs of serving new load are transmitted to retail customers.

8 Distributional Implications Across Places

Our findings indicate that data center growth increases retail electricity prices, with effects that depend importantly on utility ownership, regulatory structure, infrastructure investment, and cost recovery. A remaining salient question is whether these price increases are larger in places that may already be more vulnerable to electricity cost increases. Public debate surrounding data center expansion increasingly centers on whether the costs of serving large new loads are shifted onto households and businesses that are particularly economically disadvantaged and do not directly benefit from the new investment. We therefore examine whether price effects vary across areas with different pre-existing electricity prices and levels of economic prosperity.

Table 7 reports heterogeneity by two pre-treatment local characteristics. Panel A splits the sample by pre-treatment average retail electricity prices. Although price effects are similar for the two groups on average, they are generally larger in areas that already had higher electricity prices before data center entry once looking across customer classes, although the differences are modest. Nonetheless, this pattern suggests that data center growth may place additional pressure on areas where electricity was already relatively expensive. We interpret these results cautiously, however, because baseline price levels are correlated with other utility and regional characteristics, including ownership structure and regulatory environment.

Panel B splits the sample by GDP per capita. Price effects are larger in lower-GDP areas for average retail prices, residential prices, and industrial prices. The residential result is especially relevant for distributional concerns: households in lower-GDP areas experience larger price increases following data center entry than households in higher-GDP areas. Commercial prices, by contrast, increase more in higher-GDP areas. These patterns suggest that the costs of accommodating data center growth are not distributed uniformly across places, although they do not provide direct evidence on household-level energy burdens or income-specific incidence.

Taken together, these results provide suggestive evidence that the distributional consequences of data center growth extend beyond differences across customer classes and utility institutions. Data center entry appears to generate larger price increases in some places that were already facing higher electricity prices or lower levels of economic prosperity. Understanding who pays for growth therefore requires considering not only how costs are allocated across customer classes and regulatory institutions, but also which communities ultimately bear those costs.

9 Conclusion

Artificial intelligence and the digital economy are driving a rapid expansion of data centers and, with it, one of the largest new sources of electricity demand in modern history. Meeting this demand requires substantial complementary investment in electricity generation, transmission, and distribution infrastructure. In this paper, we ask: who pays for this growth? Combining facility-level data on data center entry with utility-level information on retail electricity prices, infrastructure investments, operational outcomes, and regulatory proceedings, we show that data center growth increases retail electricity prices on average, but that the incidence of these increases is not distributed uniformly across customers or institutions.

Data center entry increases average retail electricity prices by 2.7 percent. Residential, commercial, and industrial rates all increase, with larger percentage effects for commercial and industrial customers. These results indicate that the costs of accommodating data center growth extend beyond the facilities themselves and are passed through to incumbent ratepayers, including households and existing firms. At the same time, the commercial and industrial estimates should be interpreted with care: data centers may be classified as commercial, industrial, or special large-load customers depending on utility tariff structures. As a result, these effects may reflect both prices paid by the new large-load customers generating the demand shock and pass-through to other firms within the same customer classes.

The results also show that electricity sector institutions play a central role in shaping incidence. Ownership is the first-order source of heterogeneity: price increases are concentrated among investor-owned utilities, while effects are much smaller among publicly-owned utilities and absent among cooperatives. Regulatory environment further shapes the magnitude and mechanisms of pass-through. Price effects are generally larger in deregulated states, and especially among IOUs operating in deregulated markets, but IOU ownership remains central in both regulated and deregulated settings. These patterns suggest that who pays for data center growth depends not only on the scale of new electricity demand, but also on the institutions through which infrastructure investments are financed, regulated, and recovered.

Our mechanism evidence helps explain these patterns. Data center entry induces real infrastructure and operational responses, including construction work in progress, generation investment, transmission operation and maintenance, and larger authorized rate bases. However, the timing and pass-through of these responses differ across regulatory environments. In regulated states, vertically integrated utilities appear to accommodate load growth partly through planning, generation investment, delayed capital responses, and improved utilization, which may help mute contemporaneous price effects. In deregulated states, utilities exhibit earlier construction work in progress, more active rate case and rider proceedings, increased customer advances for construction, and stronger price effects in constrained systems. These findings suggest that retail pass-through is larger when new load creates system cost pressure that is not offset by utilization gains and is transmitted more quickly through market and regulatory cost-recovery channels.

Finally, the distributional consequences of data center growth extend beyond differences across customer classes and utility institutions. Price effects are larger in some places that already faced higher electricity prices or lower levels of economic prosperity. These results are suggestive rather than definitive evidence on household-level energy burdens, but they underscore that the costs of infrastructure-intensive growth may fall unevenly across communities. As data center demand continues to expand, especially with the newest wave of hyperscale facilities, the design of tariffs, cost-causation mechanisms, and regulatory cost recovery will become increasingly important for determining whether the costs of new infrastructure are paid by large-load customers, incumbent ratepayers, utility shareholders, or taxpayers.

More broadly, the economic consequences of new technologies are often evaluated through their effects on productivity, innovation, employment, or environmental outcomes. Yet many transformative technologies also require substantial complementary investments in physical infrastructure. The costs of those investments are rarely borne automatically by the firms developing or using the technologies themselves. Instead, they are allocated through existing institutions governing infrastructure finance, regulation, and cost recovery. As economies increasingly rely on infrastructure-intensive technologies—from AI and data centers to electrification and advanced manufacturing—it is important to understand who pays for the infrastructure that enables the growth.

References

Allen, T., and C. Arkolakis. 2022. "The Welfare Effects of Transportation Infrastructure Improvements." *The Review of Economic Studies*, 89(6): 2911–2957.

Alvarez, Fernando, David Argente, Joyce Chow, and Diana Van Patten. 2026. "Data Centers and Local Economies in the Age of AI: A Shift-Share Approach." *NBER Working Paper 35194*.

Averch, Harvey, and Leland Johnson. 1962. "Behavior of the firm under regulatory constraint." *The American Economic Review*, 52(5): 1052–1069.

Baker, Andrew, Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro H C Sant'Anna. 2026. "Difference-in-differences designs: A practitioner's guide." *Journal of Economic Literature*, 64(2): 498–557.

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

Callaway, Brantly, and Pedro H C Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.

Cicala, Steve. 2015. "When Does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation." *The American Economic Review*, 105(1): 411–444.

Cicala, Steve. 2022. "Imperfect Markets versus Imperfect Regulation in US Electricity Generation." *American Economic Review*, 112(2): 409–441.

Cicala, Steve. 2025. "Restructuring the Rate Base." *AEA Papers and Proceedings*, 115.

Davis, Lucas W, Catherine Hausman, and Nancy L Rose. 2023. "Transmission impossible? Prospects for decarbonizing the US grid." *The Journal of Economic Perspectives: A Journal of the American Economic Association*, 37(4): 155–180.

Donaldson, D. 2018. "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." *The American Economic Review*, 108(4): 899–934.

Dunkle Werner, Karl, and Stephen Jarvis. 2025. "Rate of Return revisited." *Energy Institute Working Paper*, 329R.

Fabra, Natalia, and Mar Reguant. 2014. "Pass-Through of Emissions Costs in Electricity Markets." *The American Economic Review*, 104(9): 2872–2899.

Gowrisankaran, G., A. Langer, and M. Reguant. 2026. "Energy Transitions in Regulated Markets." *The American Economic Review*.

Hainmueller, Jens. 2012. "Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies." *Political Analysis: An Annual Publication of the Methodology Section of the American Political Science Association*, 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.

Hernandez-Cortes, Danae, Kyle Meng, and Paige Weber. 2026. "The Environmental Costs and Geography of U.S. Data Center Expansion." UC Berkeley Haas Working Paper WP-363.

IEA. 2025. "Energy and AI: World Energy Outlook Special Report." International Energy Agency.

Joskow, Paul. 1974. "Inflation and Environmental Concern: Structural Change in the Process of Public Utility Price Regulation." *The Journal of Law and Economics*, 17(2): 291–327.

Joskow, Paul, and Nancy Rose. 1989. "The Effects of Economic Regulation." *Handbook of Industrial Organization*.

Kay, Owen, Robert Reaser, and Reid Taylor. 2026. "Processing Power: The Effect of Data Centers on Wholesasle Electricity Markets." Federal Reserve Bank of Dallas.

Knittel, Chris, Juan Ramon Senga, and Shen Wang. 2025. "Flexible Data Centers and the Grid: Lower Costs, Higher Emissions?" NBER Working Paper 34065.

Lazar, Jim. 2016. "Electricity Regulation in the US: A Guide. Second Edition." The Regulatory Assistance Project., Montpelier, VT.

Masanet, Eric, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. 2020. "Recalibrating global data center energy-use estimates." *Science*, 367(6481): 984–986.

Mongird, Kendall, Chris Vernon, Casey Burleyson, Kerem Ziya Akdemir, and Jennie Rice. 2026. "IM3 Open Source Data Center Atlas."

Muller, Nicholas Z. 2026. "Measuring the impact of data centers in the united states economy: monetary damage from air pollution and greenhouse gas emissions." 35100, Cambridge:National Bureau of Economic Research.

Sant'Anna, Pedro H C, and Jun Zhao. 2020. "Doubly robust difference-in-differences estimators." *Journal of Econometrics*, 219(1): 101–122.

Satchwell, Andrew, Natalie Frick, Peter Cappers, Sanem Sergici, Ryan Hledik, Goksin Kavlak, and Glenda Oskar. 2025. "Electricity Rate Designs for Large Loads: Evolving Practices and Opportunities." ark:/13030/qt19m1k8vr, None, 2502077.

Shehabi, Arman, Sarah J Smith, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu, Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, and Dale Sartor. 2024. "2024 United States Data Center Energy Usage Report." Berkeley:Lawrence Berkeley National Laboratory.

Yue, Daniel, and Yiyang Zeng. 2026. "The Local Economic Effects of Data Center Entry." *Working Paper*.

Main Text Figures

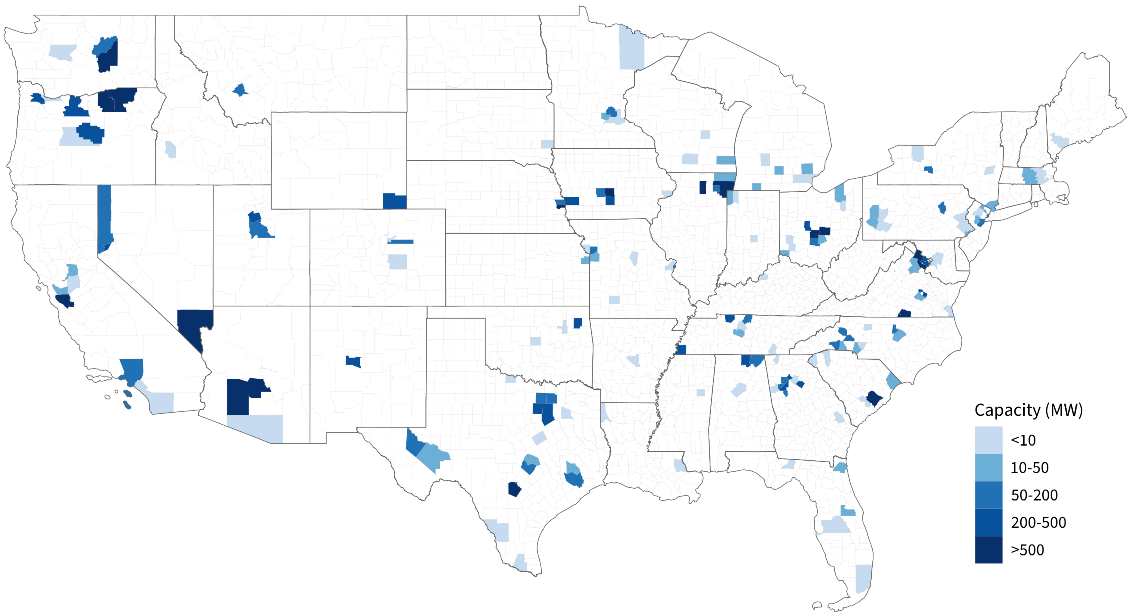

Figure 1: Active Data Center Capacity by County (Activated after 2010)

Notes: This map shows the county-level total capacity (MW) of data centers that are activated between 2010 and 2026.

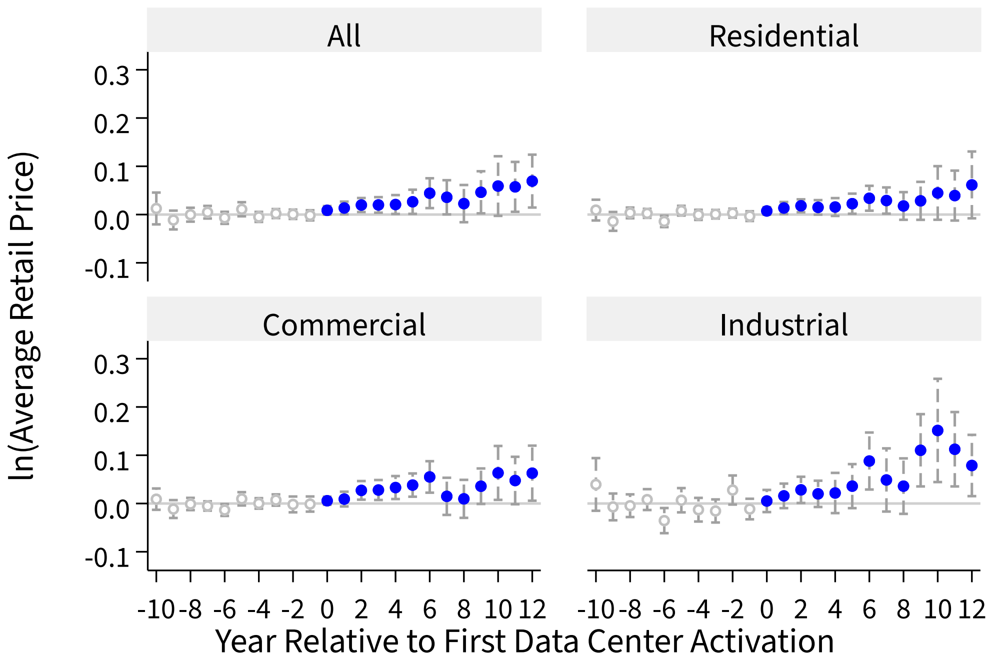

Figure 2: Dynamic Impact of Data Center Entry on Average Retail Prices

Notes: The unit of observation is the utility-state-year. Figures plot event-study coefficients and 95% confidence intervals, capturing the difference in (logged) average retail electricity rates for utility-state pairs with and without data centers before and after initial data center entry. The first panel shows estimates for average retail rates across all customer classes, whereas the following three panels examine average rates for residential, commercial, and industrial customers separately. Standard errors are clustered at the utility-state level.

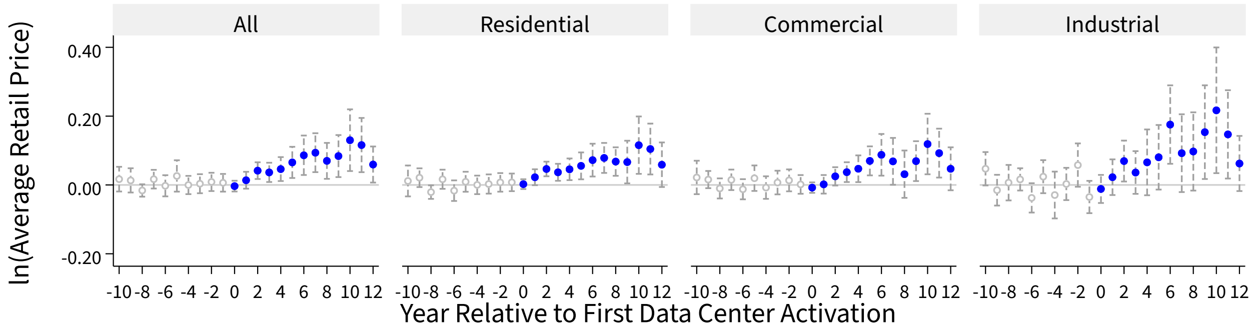

(a) Investor-Owned Utilities

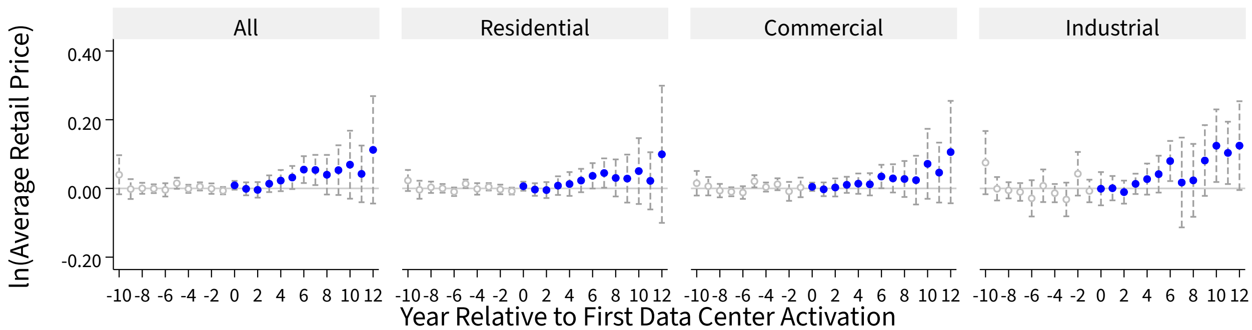

(b) Publicly-Owned Utilities

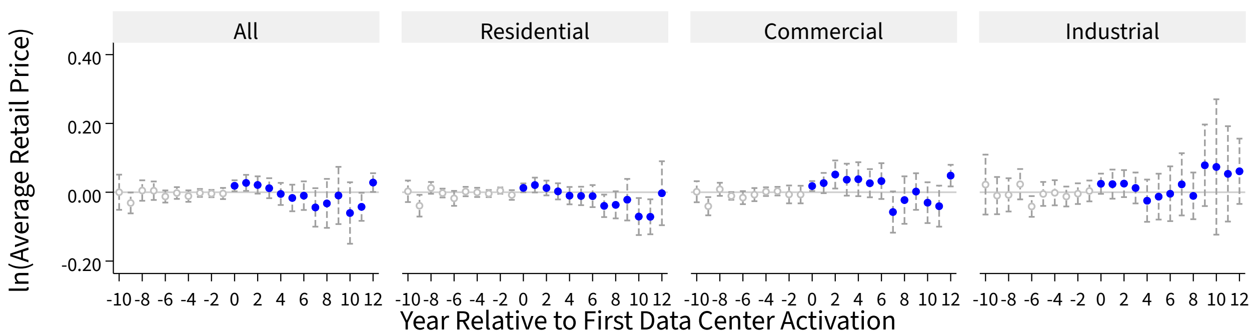

(c) Cooperative

Figure 3: Dynamic Impact of Data Center Entry on Average Retail Prices by Utility Ownership

Notes: The unit of observation is the utility-state-year. Figures plot event-study coefficients and 95% confidence intervals by utility ownership, capturing the difference in (logged) average retail electricity rates for utility-state pairs with and without data centers before and after initial data center entry. Panel A shows the effects for IOUs, Panel B for public utilities, and Panel C for cooperatives. For each subsample analysis, we include all utilities without data centers in the control group (despite ownership type) while the treatment group in each case includes only utilities with each ownership type that ultimately have data centers. In each panel, the first column shows estimates for average retail rates across all customer classes, whereas the following three columns examine average rates for residential, commercial, and industrial customers separately. Standard errors are clustered at the utility-state level.

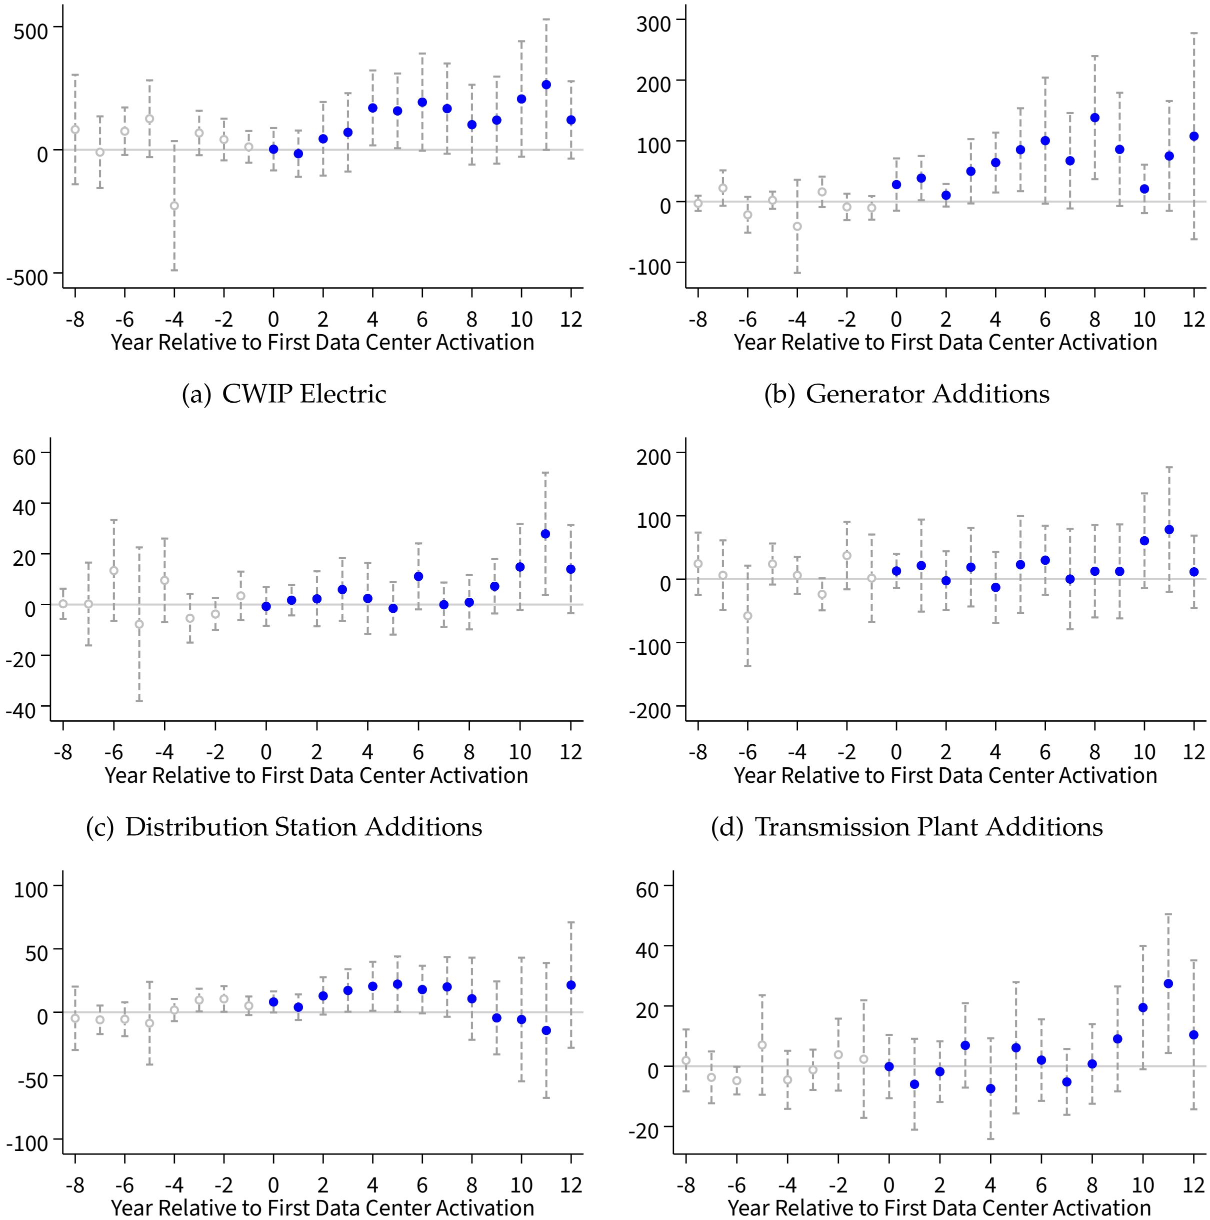

(a) CWIP Electric

(b) Generator Additions

(c) Distribution Station Additions

(d) Transmission Plant Additions

(e) Transmission O&M

(f) Conductor Additions

Figure 4: Effect of Data Center Entry on Infrastructure Investment

Notes: The unit of observation is the utility-year, and the sample only includes investor-owned utilities (IOUs). Figures plot event-study coefficients and the 95% confidence intervals, capturing the difference in infrastructure investments for utilities with and without data centers before and after initial data center entry. The outcomes are (a) electric construction work in progress, (b) generator additions, (c) distribution station equipment additions, (d) transmission plant additions, (e) transmission operation & maintenance (O&M) expense, and (f) overhead conductor additions. All variables are in millions of 2024 U.S. dollars. Standard errors are clustered at the utility level.

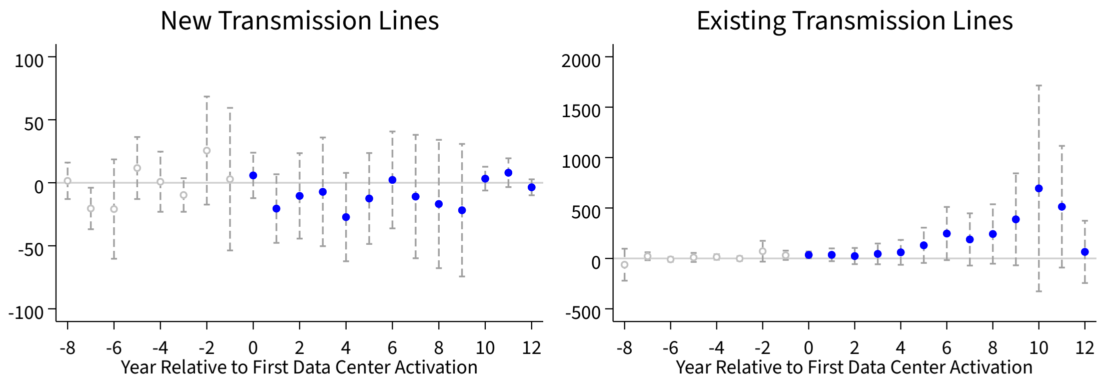

(a) All Utilities

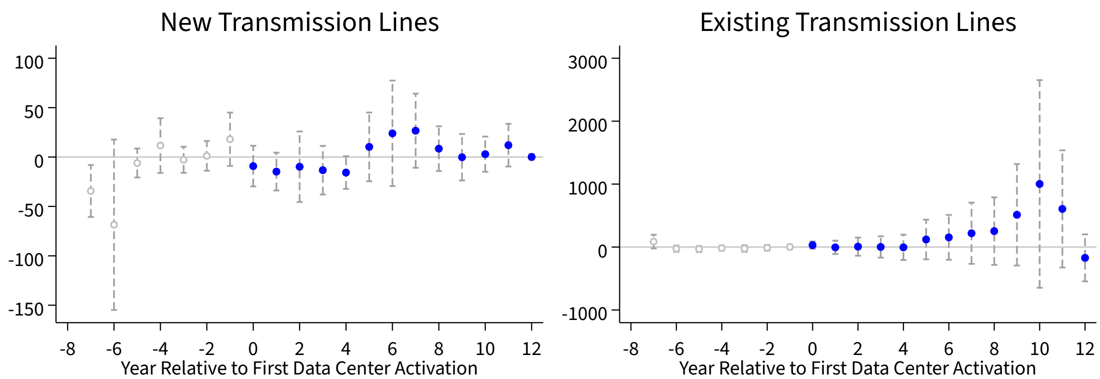

(b) Utilities Operating in Deregulated States

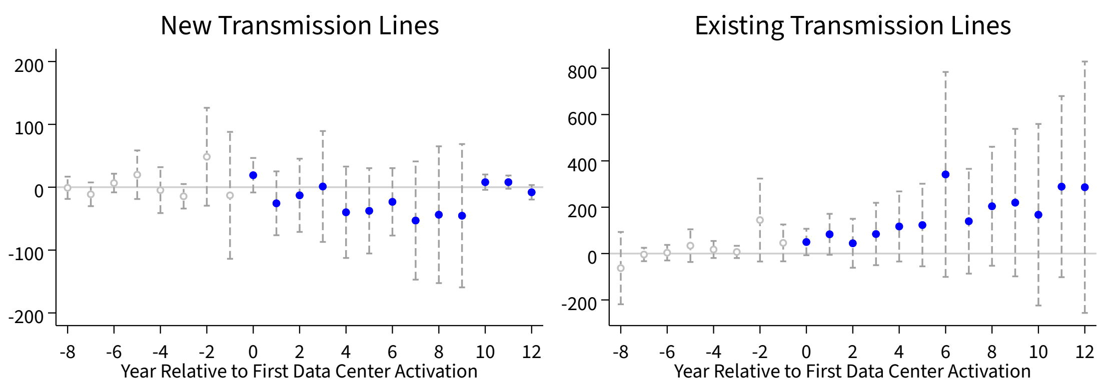

(c) Utilities Operating in Regulated States

Figure 5: Effect of Data Center Entry on New and Existing Transmission Line Expenditures

Notes: The unit of observation is the utility-year, and the sample only includes investor-owned utilities (IOUs). Figures plot event-study coefficients and the 95% confidence intervals, capturing the difference in infrastructure investments for utilities with and without data centers before and after initial data center entry. The outcomes are expenditures on new transmission lines (left) and on existing transmission lines (right). Panel A reports estimates for all utilities. Panel B reports estimates for utilities operating only in deregulated states, while Panel C reports estimates for utilities operating in regulated states. All variables are in millions of 2024 U.S. dollars. Standard errors are clustered at the utility level.

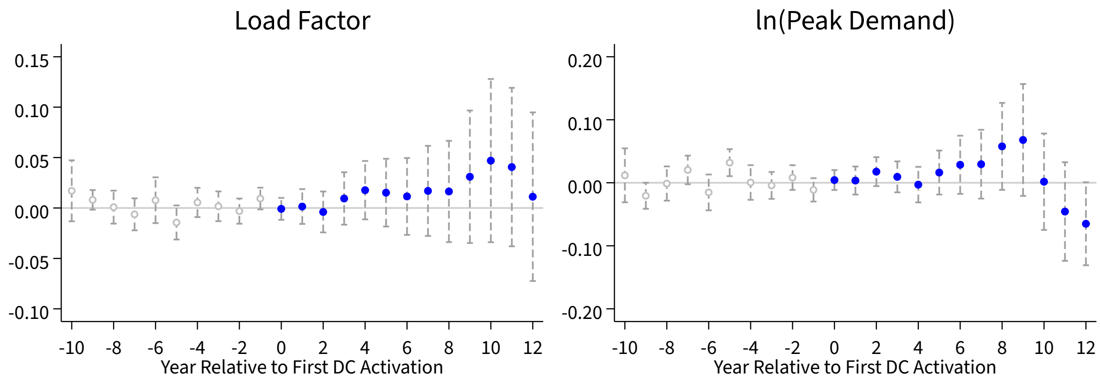

(a) All Utilities

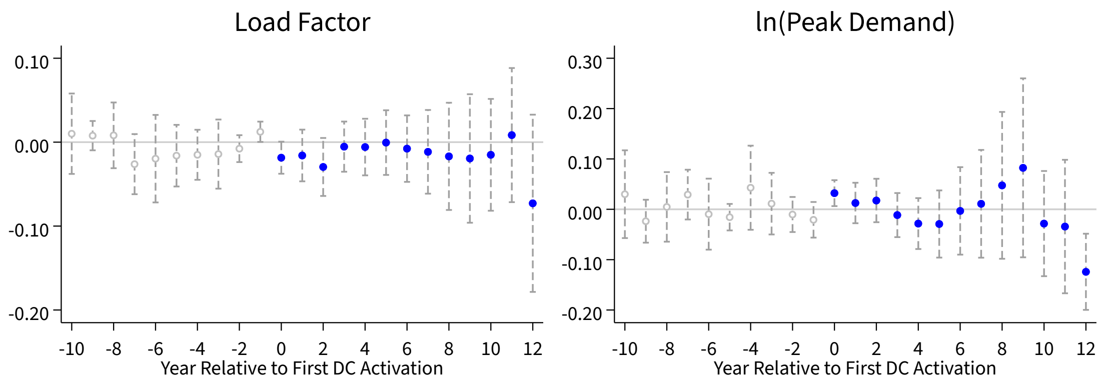

(b) Utilities Operating in Deregulated States

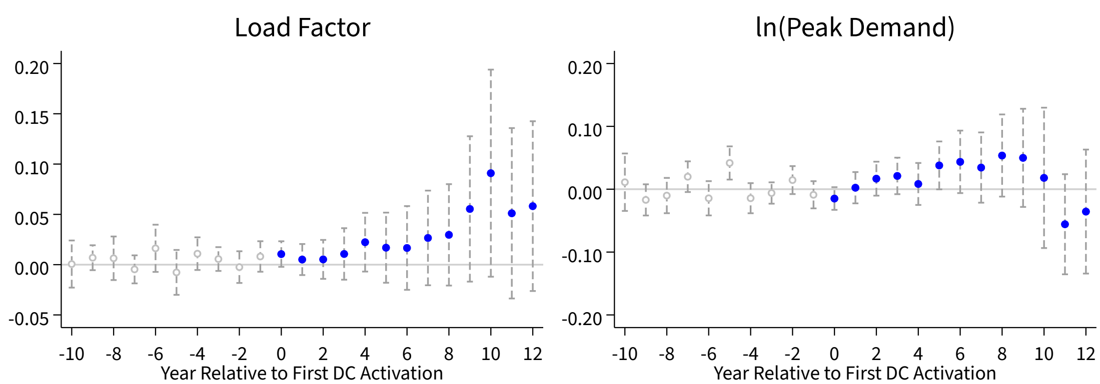

(c) Utilities Operating in Regulated States

Figure 6: Effect of Data Center Entry on Load Factor and Peak Demand

Notes: The unit of observation is the utility-year. Figures plot event-study coefficients and the 95% confidence intervals. The left panels plot the system load factor, defined as a utility's total energy disposition divided by the product of peak demand and the number of hours in the year. The right panels plot the logarithm of system peak demand, measured in megawatts. Panel A reports estimates for all utilities. Panel B reports estimates for utilities operating only in deregulated states, while Panel C reports estimates for utilities operating in regulated states. Standard errors are clustered at the utility level.

Main Text Tables

Table 1: Descriptive Statistics

	N	Mean	S.D.
A. Average Retail Electricity Prices ($/kWh)			
All	14,733	0.13	0.03
Residential	14,733	0.14	0.03
Commercial	14,733	0.13	0.03
Industrial	14,733	0.10	0.08
B. Data Centers			
PostDC	14,733	0.06	0.24
Cumulative capacity (MW)	14,733	5.02	44.56
Cumulative capacity (MW) conditional on entry	908	80.58	161.50
C. Infrastructure Investment ($ million)			
CWIP electric	1,590	333.55	631.60
Generator additions	1,590	33.70	150.69
Distribution station additions	1,590	22.95	33.97
Transmission plant additions	1,590	93.18	183.35
Transmission O&M	1,590	77.12	109.89
Conductor additions	1,590	16.17	39.22
Expenditure on new transmission lines	1,590	21.39	85.62
Expenditure on existing transmission lines	1,590	644.88	1042.97
D. Rate Cases			
Any case filed (0/1)	2,070	0.27	0.44
Any general case filed (0/1)	2,070	0.23	0.42
Any limited-issue rider case filed (0/1)	2,070	0.05	0.21
Number of cases filed	2,070	0.33	0.68
Mean rate case duration (months)	555	8.13	4.23
Requested rate change amount ($ billion)	2,070	0.03	0.08
Authorized rate change amount ($ billion)	2,070	0.01	0.05
Requested ROE (%)	462	10.34	0.58
Authorized ROE (%)	399	9.66	0.47
Requested return on rate base (%)	472	7.52	0.79
Authorized return on rate base (%)	355	7.09	0.76
Authorized rate base ($ billion)	336	3.42	4.88
E. Operational Characteristics			
Load factor	14,357	0.56	0.49
Peak demand (MW)	14,357	440.27	1361.44

Notes: Panels A, B, D are at the utility-state-year level. Panels C and E are at the utility-year level. Panels C and D only include investor-owned utilities. Dollar variables are in 2024 dollars.

Table 2: The Effect of Data Center Entry on Average Retail Prices

Dep. Var.: ln(Avg. Retail Price)	All (1)	Residential (2)	Commercial (3)	Industrial (4)
PostDC	0.027***	0.021***	0.028***	0.042**
	(0.009)	(0.008)	(0.009)	(0.016)
YMean ($/kWh)	0.13	0.14	0.13	0.10
# of Utilities	973	973	973	973
N	14,733	14,733	14,733	14,733

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall aggregated average treatment effect on the treated (ATT). The dependent variable is the log average retail price (total revenue divided by total sales, $/kWh), for all customers in column (1) and by customer class in columns (2)–(4). Standard errors in parentheses are clustered at the utility-state level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Heterogeneity in Data Center Effects on Retail Price by Utility Ownership

Dep. Var.: ln(Avg. Retail Price)	All (1)	Residential (2)	Commercial (3)	Industrial (4)
A. Investor-Owned Utilities				
PostDC	0.056***	0.052***	0.046***	0.080**
	(0.016)	(0.014)	(0.015)	(0.035)
YMean ($/kWh)	0.14	0.16	0.14	0.10
# of Utilities	900	900	900	900
N	13,585	13,585	13,585	13,585
B. Publicly-Owned Utilities				
PostDC	0.024*	0.017	0.017	0.027
	(0.015)	(0.015)	(0.015)	(0.020)
YMean ($/kWh)	0.12	0.14	0.13	0.10
# of Utilities	911	911	911	911
N	13,607	13,607	13,607	13,607
C. Cooperative				
PostDC	-0.001	-0.007	0.018	0.014
	(0.013)	(0.010)	(0.013)	(0.022)
YMean ($/kWh)	0.13	0.15	0.14	0.11
# of Utilities	912	912	912	912
N	13,687	13,687	13,687	13,687

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). The dependent variable is the log average retail electricity price, defined as total revenue divided by total sales ($/kWh), for all customers in column (1) and separately by customer class in columns (2)–(4). Panel A reports estimates for investor-owned utilities, Panel B for public utilities, and Panel C for cooperatives. In each subsample analysis, the treatment group consists of utilities of the corresponding ownership type that eventually serve data centers, while the control group includes all utilities that never serve data centers, irrespective of ownership type. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Heterogeneity in Data Center Effects on Retail Price by Regulatory Environment

Dep. Var.: ln(Avg. Retail Price)	All (1)	Residential (2)	Commercial (3)	Industrial (4)
A. Deregulated States				
PostDC	0.038**	0.040**	0.029*	0.047
	(0.017)	(0.016)	(0.016)	(0.032)
YMean ($/kWh)	0.14	0.16	0.15	0.11
# of Utilities	239	239	239	239
N	3,421	3,421	3,421	3,421
B. Regulated States				
PostDC	0.015	0.007	0.024**	0.027
	(0.010)	(0.009)	(0.010)	(0.017)
YMean ($/kWh)	0.12	0.14	0.13	0.10
# of Utilities	755	755	755	755
N	11,312	11,312	11,312	11,312

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). The dependent variable is the log average retail electricity price, defined as total revenue divided by total sales ($/kWh), for all customers in column (1) and separately by customer class in columns (2)–(4). Panel A reports estimates for utilities operating in deregulated states, while Panel B reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: The Effect of Data Center Entry on Rate Case Activity

	Any Case Filed?			# Cases	Duration
	Any	General	Rider		
	(1)	(2)	(3)	(4)	(5)
A. Full Sample					
PostDC	0.084	0.104	-0.014	0.158	0.776
	(0.085)	(0.074)	(0.055)	(0.155)	(0.689)
YMean	0.27	0.23	0.05	0.33	8.13
# of Utilities	106	106	106	106	106
N	2,070	2,070	2,070	2,070	555
B. Deregulated States					
PostDC	0.265**	0.173*	0.122*	0.517**	-0.981*
	(0.110)	(0.097)	(0.070)	(0.234)	(0.549)
YMean	0.29	0.25	0.05	0.42	8.25
# of Utilities	51	51	51	51	51
N	780	780	780	780	229
C. Regulated States					
PostDC	-0.060	0.055	-0.115	-0.156	1.358
	(0.117)	(0.111)	(0.080)	(0.165)	(0.932)
YMean	0.25	0.22	0.04	0.28	8.04
# of Utilities	67	67	67	67	67
N	1,290	1,290	1,290	1,290	326

Notes: The unit of observation is the utility-state-year and the sample only includes investor-owned utilities. The table reports the estimated overall average treatment effect on the treated (ATT). The outcomes are indicators for filing any rate case (column 1), any general rate case (column 2), and any limited-issue rider case (column 3); the number of cases filed (column 4); and the mean case duration in months, defined only for utility-state-years with at least one filed case (column 5). Panel A reports estimates for the full sample. Panel B reports estimates for utilities operating in deregulated states, while Panel C reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: The Effect of Data Center Entry on Rate Case Requests and Authorizations

	Revenue Requirement ($ billion)		ROE (%)		Return on Rate Base (%)		Authorized Rate Base ($ billion)
	Requested (1)	Authorized (2)	Requested (3)	Authorized (4)	Requested (5)	Authorized (6)	(7)
A. Full Sample							
PostDC	0.017	-0.003	0.178*	-0.064	0.120	0.022	1.646***
	(0.018)	(0.014)	(0.092)	(0.123)	(0.104)	(0.202)	(0.444)
YMean	0.03	0.01	10.34	9.66	7.52	7.09	3.42
# of Utilities	106	106	106	106	106	106	106
N	2,070	2,070	462	399	472	355	336
B. Deregulated States							
PostDC	0.040***	0.012	0.098	-0.182	0.100	-0.034	1.251*
	(0.014)	(0.014)	(0.159)	(0.212)	(0.120)	(0.378)	(0.652)
YMean	0.03	0.02	10.30	9.55	7.45	6.95	4.93
# of Utilities	51	51	51	51	51	51	51
N	780	780	200	188	199	150	134
C. Regulated States							
PostDC	-0.003	-0.017	0.014	0.012	0.031	0.289	1.594**
	(0.030)	(0.022)	(0.066)	(0.056)	(0.130)	(0.232)	(0.626)
YMean	0.02	0.01	10.37	9.75	7.57	7.20	2.41
# of Utilities	67	67	67	67	67	67	67
N	1,290	1,290	262	211	273	205	202

Notes: The unit of observation is the utility-state-year and the sample only includes investor-owned utilities. The table reports the estimated overall average treatment effect on the treated (ATT). The outcomes are the requested and authorized change in the annual revenue requirement ($ billion) in columns (1)–(2), the requested and authorized return on equity (ROE, %) in (3)–(4), the requested and authorized return on rate base (%) in (5)–(6), and the authorized rate base ($ billion) in (7). The revenue-requirement outcomes are defined for every utility-state-year (zero when no change is decided), while the ROE, return-on-rate-base, and rate-base outcomes are defined only when a corresponding value is requested or decided. Panel A reports estimates for the full sample. Panel B reports estimates for utilities operating in deregulated states, while Panel C reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Distributional Consequences

Dep. Var.:	All		Residential		Commercial		Industrial	
ln(Avg. Retail Price)	High (1)	Low (2)	High (3)	Low (4)	High (5)	Low (6)	High (7)	Low (8)
A: By Baseline Average Price Level								
PostDC	0.024*	0.027**	0.024**	0.015	0.031**	0.020*	0.046*	0.033**
	(0.014)	(0.011)	(0.012)	(0.011)	(0.012)	(0.012)	(0.028)	(0.016)
YMean ($/kWh)	0.14	0.11	0.16	0.13	0.15	0.12	0.12	0.09
# of Utilities	508	469	508	469	508	469	508	469
N	7,315	7,326	7,315	7,326	7,315	7,326	7,315	7,326
B: By GDP per Capita								
PostDC	0.021*	0.027***	0.015	0.026***	0.027**	0.012*	0.027	0.042***
	(0.011)	(0.009)	(0.010)	(0.008)	(0.011)	(0.007)	(0.020)	(0.013)
YMean ($/kWh)	0.13	0.13	0.15	0.14	0.13	0.14	0.10	0.10
# of Utilities	485	510	485	510	485	510	485	510
N	7,342	7,391	7,342	7,391	7,342	7,391	7,342	7,391

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). The dependent variable is the log average retail electricity price for all customers in columns (1)–(2), and separately by customer class in columns (3)–(8): residential in columns (3)–(4), commercial in columns (5)–(6), and industrial in columns (7)–(8). Within each pair of columns, the sample is split at the median into "High" (above-median) and "Low" (at-or-below-median) subsamples. Panel A splits the sample by the utility-state's pre-treatment average retail price. Panel B splits the sample by pre-treatment county GDP per capita, averaged across counties in the utility-state's service territory. Both variables are measured using data from 2005–2007. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Online Appendix

Who Pays for Growth? Evidence from Data Centers and the Grid

Robyn Meeks Jacquelyn Pless Zhiyuan Qi Zhenxuan Wang

A Additional Figures

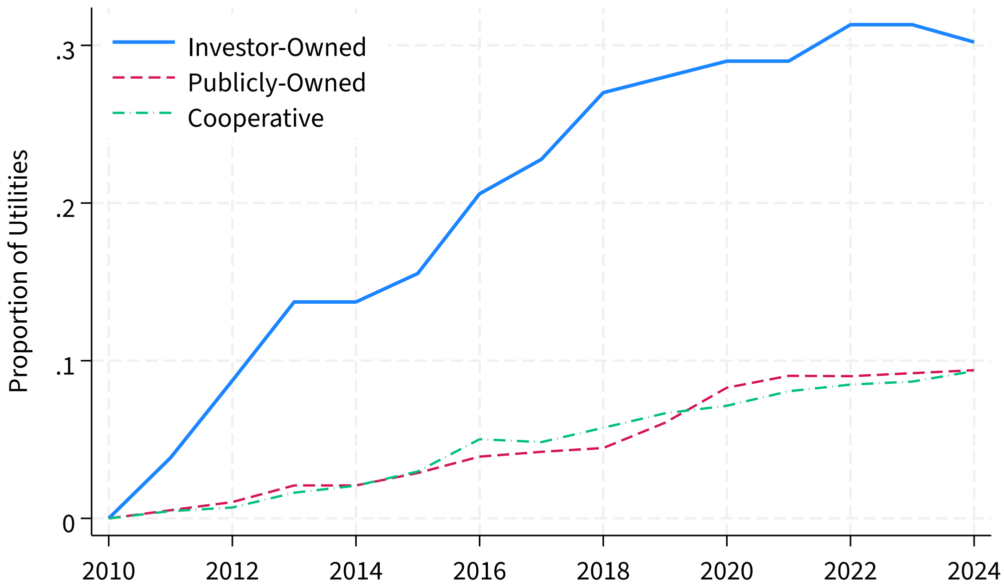

(a) Proportion of Utilities with Data Centers Over Time

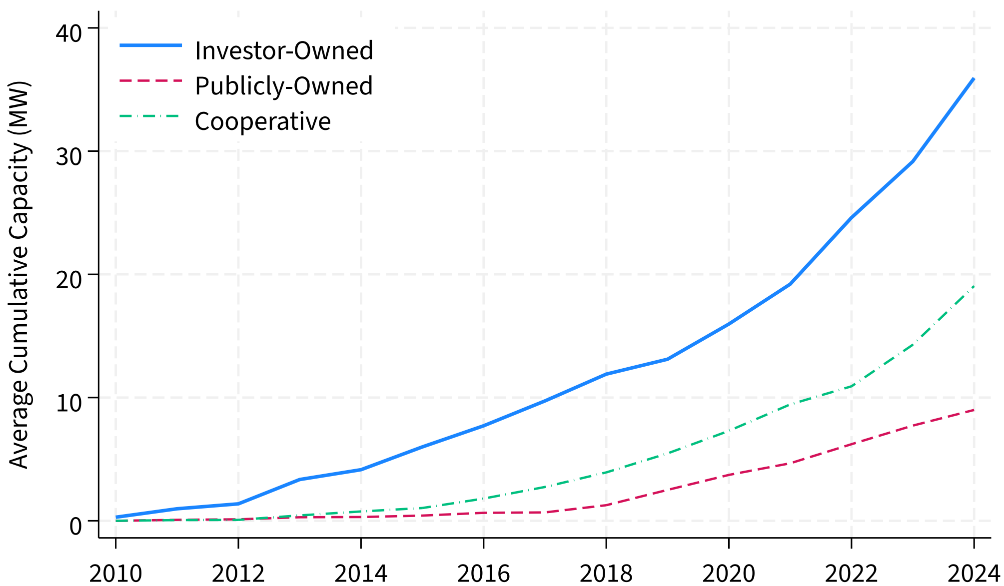

(b) Average Cumulative Data Center Capacity by Utility Ownership

Figure A1: Data Center Entry and Expansion by Utility Ownership

Notes: Panel (a) plots the share of utility-states in the sample with at least one active data center in their service territory in each year. Panel (b) plots average cumulative active data center capacity (MW) per utility-state in the sample. Both outcomes are reported separately for investor-owned utilities, publicly owned utilities, and cooperatives.

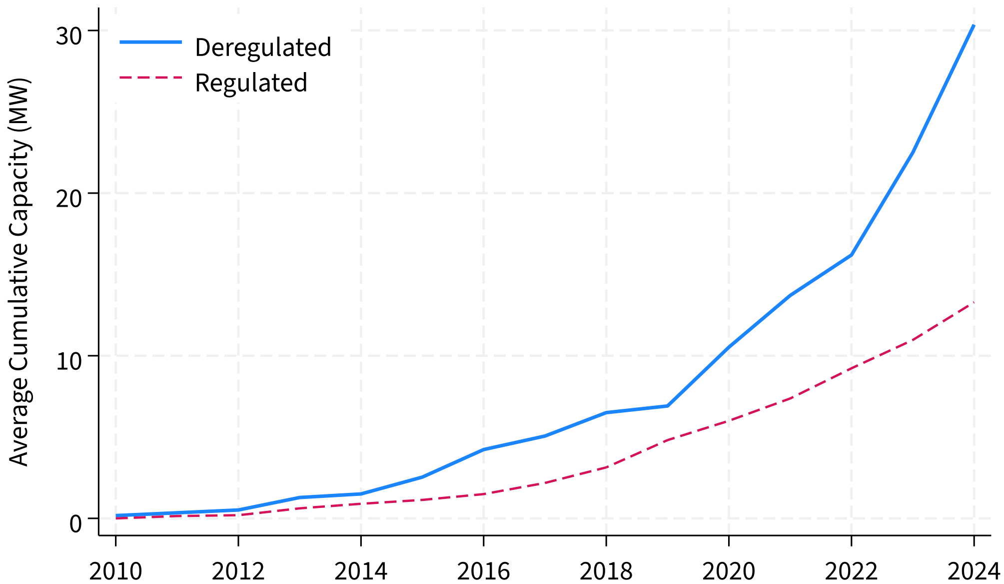

Figure A2: Average Cumulative Data Center Capacity by Regulatory Environment

Notes: The figure plots average cumulative active data center capacity (MW) per utility-state by year in our sample, separately for deregulated and regulated states.

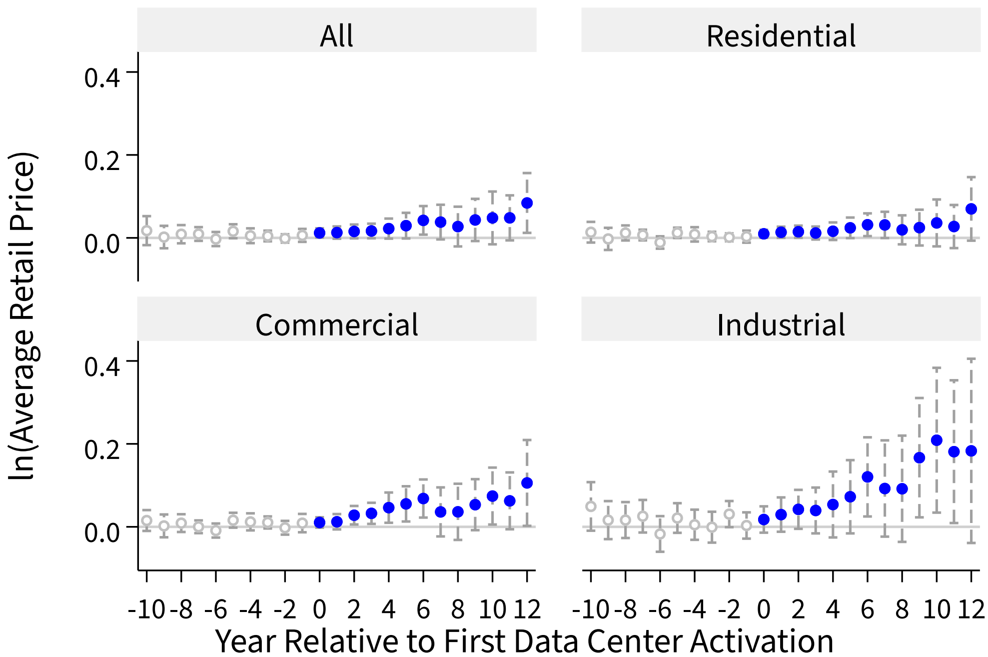

Figure A3: Dynamic Impact of Data Center Entry on Average Retail Prices – Matched DID Design via Entropy Balancing

Notes: We first reweight never-treated utility-state pairs by entropy balancing (Hainmueller, 2012) so that their pre-period means match those of the treated pairs, then estimate the overall average treatment effect on the treated (ATT) using the Callaway and Sant'Anna (2021) estimator with these weights. Balancing is performed once on four pre-treatment (2005–2007) characteristics measured before any data center entry: the average retail price, the number of customers, manufacturing establishments per capita, and GDP per capita. Treated units receive a weight of one. The unit of observation is the utility-state-year. Figures plot event-study coefficients and 95% confidence intervals. The first panel shows estimates for average retail rates across all customer classes, whereas the following three panels examine average rates for residential, commercial, and industrial customers separately. Standard errors are clustered at the utility-state level.

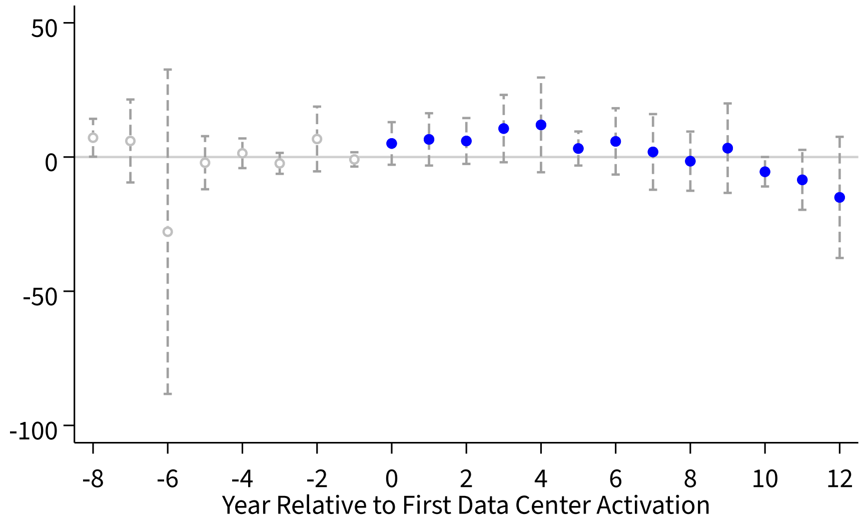

(a) CWIP Common

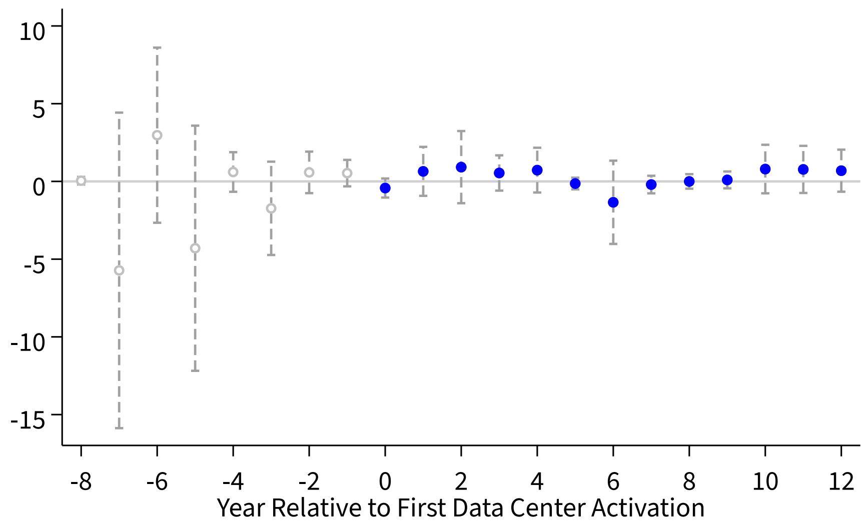

(b) CWIP Other

Figure A4: Placebo Tests on Construction Work in Progress

Notes: The unit of observation is the utility-year, and the sample only includes investor-owned utilities (IOUs). Figures plot event-study coefficients and the 95% confidence intervals. Panel (a) plots common construction work in progress, i.e., projects that support multiple utility functions simultaneously such as IT systems, administrative officies, shared service centers, general office equipment, etc. Panel (b) plots other construction work in progress, i.e. projects attributable to other utility operations. All variables are in millions of 2024 U.S. dollars. Standard errors are clustered at the utility level.

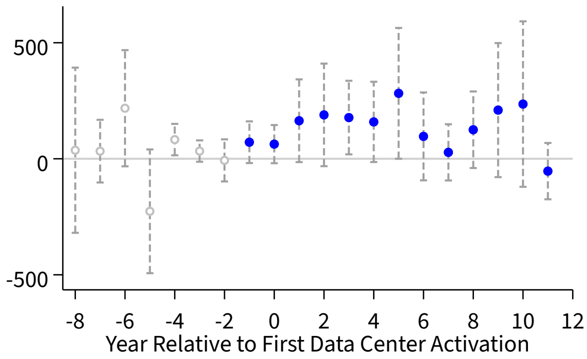

(a) CWIP Electric

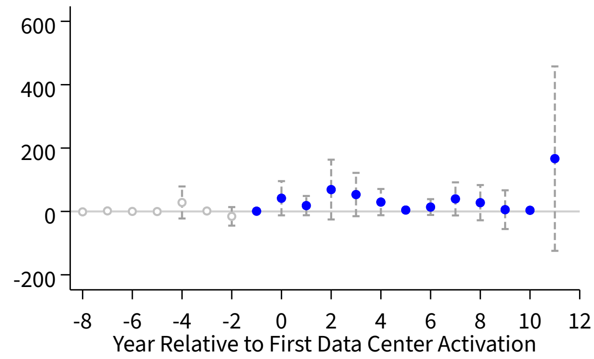

(b) Generator Additions

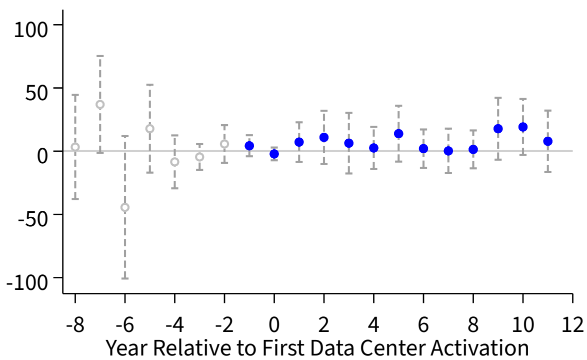

(c) Distribution Station Additions

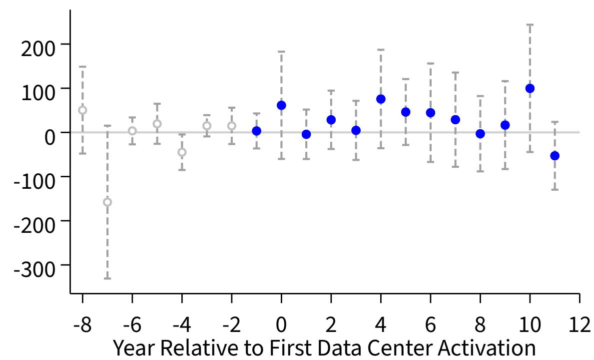

(d) Transmission Plant Additions

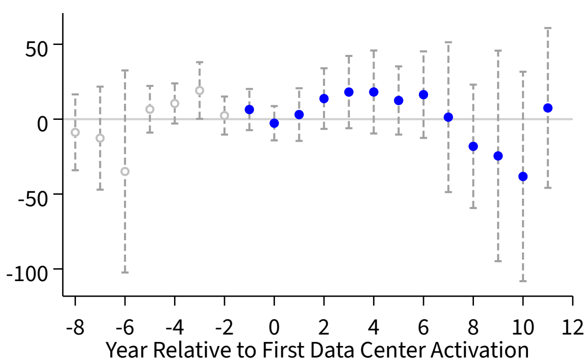

(e) Transmission O&M

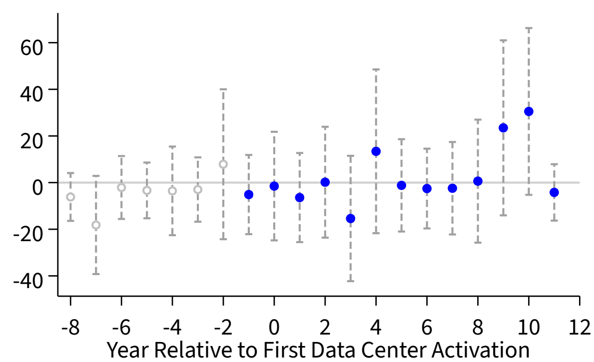

(f) Conductor Additions

Figure A5: Effect of Data Center Entry on Infrastructure Investment for Utilities in Deregulated Markets

Notes: The unit of observation is the utility-year, and the sample only includes investor-owned utilities (IOUs). Figures plot event-study coefficients and the 95% confidence intervals using the sample of utilities that only operate in deregulated states. The outcomes are (a) electric construction work in progress, (b) generator additions, (c) distribution station equipment additions, (d) transmission plant additions, (e) transmission operation & maintenance (O&M) expense, and (f) overhead conductor additions. All variables are in millions of 2024 U.S. dollars. Standard errors are clustered at the utility level.

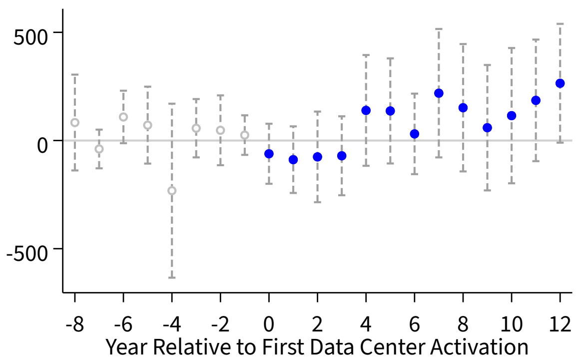

(a) CWIP Electric

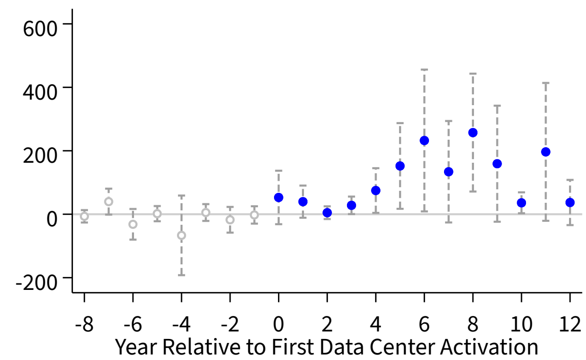

(b) Generator Additions

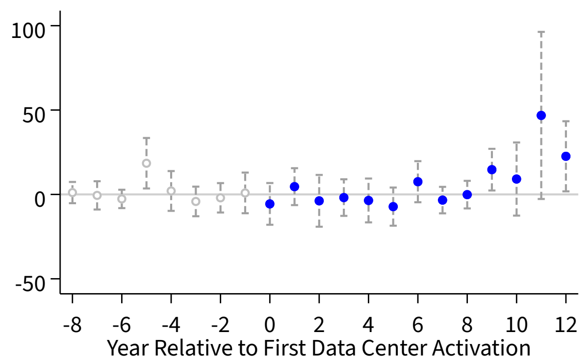

(c) Distribution Station Additions

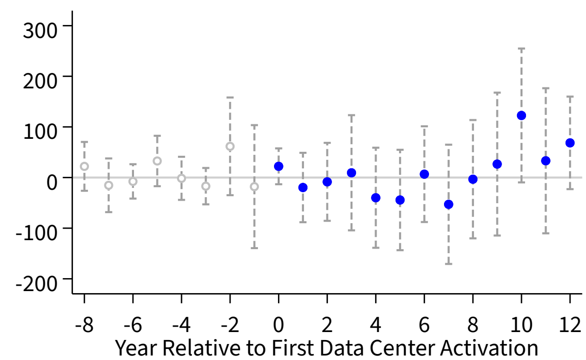

(d) Transmission Plant Additions

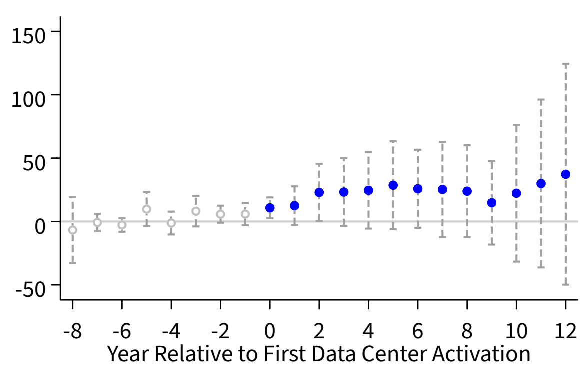

(e) Transmission O&M

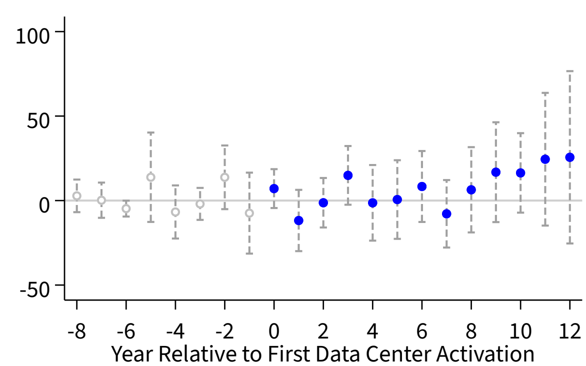

(f) Conductor Additions

Figure A6: Effect of Data Center Entry on Infrastructure Investment for Utilities in Regulated Markets

Notes: The unit of observation is the utility-year, and the sample only includes investor-owned utilities (IOUs). Figures plot event-study coefficients and the 95% confidence intervals using the sample of utilities that operate in regulated states. The outcomes are (a) electric construction work in progress, (b) generator additions, (c) distribution station equipment additions, (d) transmission plant additions, (e) transmission operation & maintenance (O&M) expense, and (f) overhead conductor additions. All variables are in millions of 2024 U.S. dollars. Standard errors are clustered at the utility level.

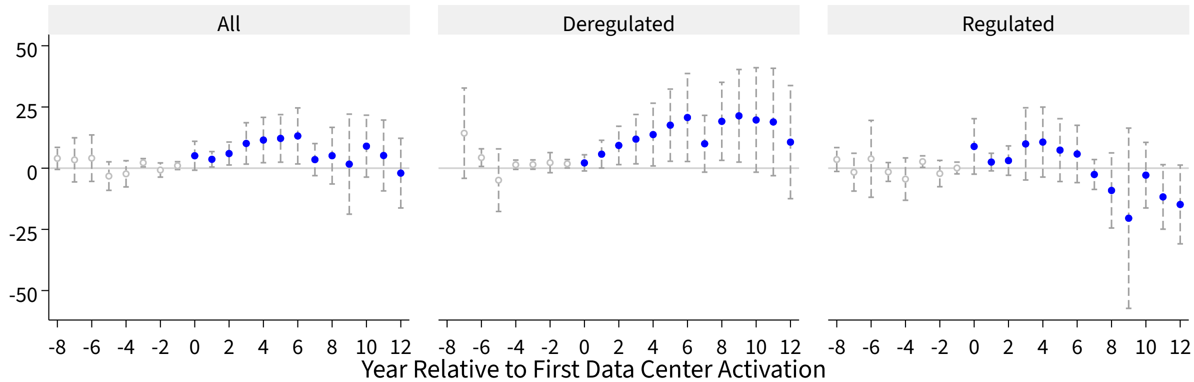

Figure A7: Effect of Data Center Entry on Utilities' Customer Advances for Construction

Notes: The unit of observation is the utility-year, and the sample only includes investor-owned utilities (IOUs). Figures plot event-study coefficients and the 95% confidence intervals. The left panel plots estimates for all utilities. The middle panel plots estimates for utilities operating only in deregulated states, while the right panel plots estimates for utilities operating in regulated states. The outcome is customer advances for construction (CAC), i.e., the upfront cash contributions customers make toward infrastructure built to serve them. All variables are in millions of 2024 U.S. dollars. Standard errors are clustered at the utility level.

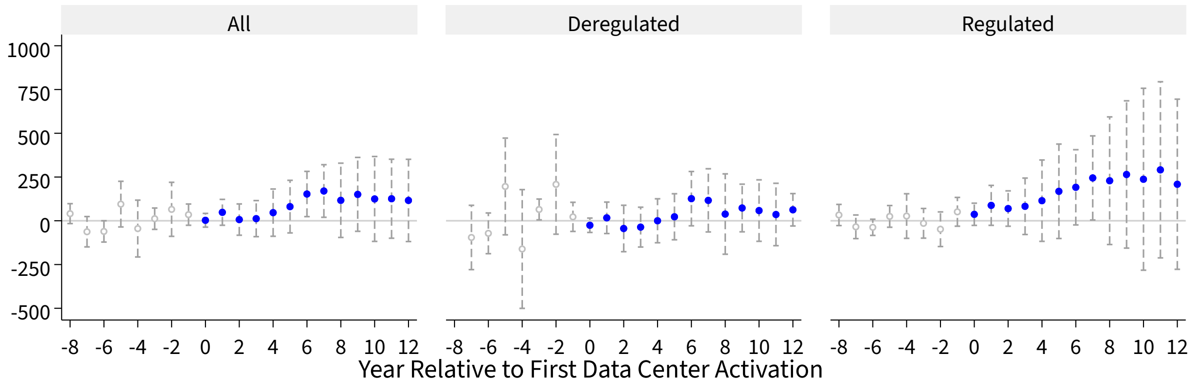

Figure A8: Effect of Data Center Entry on Accumulated Deferred Income Taxes (ADIT) – Electric

Notes: The unit of observation is the utility-year, and the sample only includes investor-owned utilities (IOUs). Figures plot event-study coefficients and the 95% confidence intervals. The left panel plots estimates for all utilities. The middle panel plots estimates for utilities operating only in deregulated states, while the right panel plots estimates for utilities operating in regulated states. The outcome is accumulated deferred income taxes (ADIT) for electric operations, which arise mainly from accelerated tax depreciation and are netted against the rate base. All variables are in millions of 2024 U.S. dollars. Standard errors are clustered at the utility level.

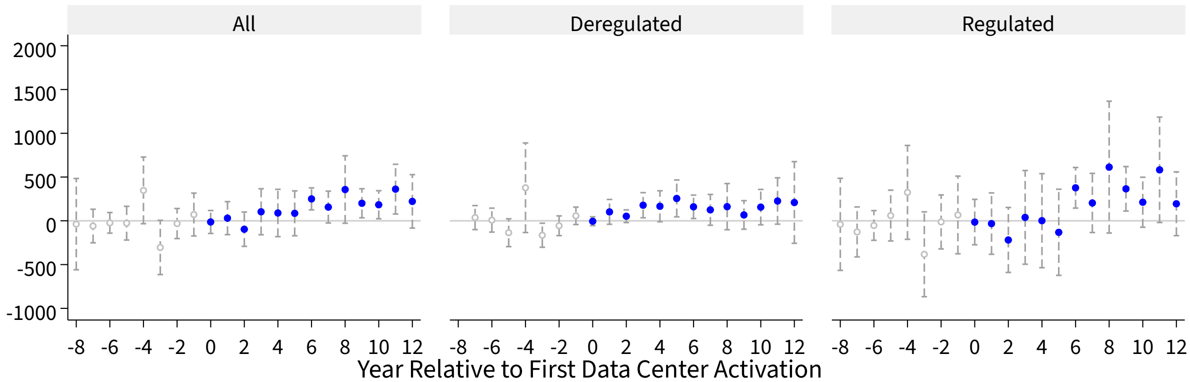

Figure A9: Effect of Data Center Entry on Electric Plants in Service

Notes: The unit of observation is the utility-year, and the sample only includes investor-owned utilities (IOUs). Figures plot event-study coefficients and the 95% confidence intervals. The left panel plots estimates for all utilities. The middle panel plots estimates for utilities operating only in deregulated states, while the right panel plots estimates for utilities operating in regulated states. The outcome is the addition to electric plant in service. All variables are in millions of 2024 U.S. dollars. Standard errors are clustered at the utility level.

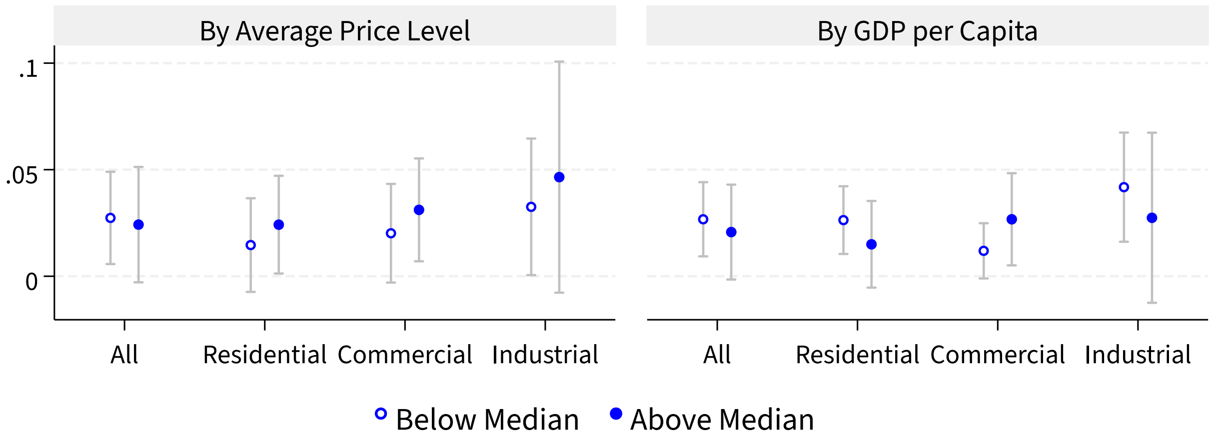

Figure A10: Distributional Consequences

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). The dependent variable is the log average retail electricity price for all customers and separately by customer class. Within each customer class, the sample is split at the median into "High" (above-median) and "Low" (at-or-below-median) subsamples. The left panel splits the sample by the utility-state's pre-treatment average retail price. The right panel splits the sample by pre-treatment county GDP per capita, averaged across counties in the utility-state's service territory. Both variables are measured using data from 2005–2007. Standard errors are clustered at the utility-state level. The corresponding estimates are reported in Table 7.

B Additional Tables

Table A1: Heterogeneity in Data Center Effects on Retail Price by Utility Ownership and Regulatory Environment

Dep. Var.:	All		Residential		Commercial		Industrial	
ln(Avg. Retail Price)	Dereg. (1)	Reg. (2)	Dereg. (3)	Reg. (4)	Dereg. (5)	Reg. (6)	Dereg. (7)	Reg. (8)
A. Investor-Owned Utilities								
PostDC	0.061**	0.044**	0.061***	0.040***	0.045*	0.042**	0.075	0.068**
	(0.025)	(0.018)	(0.023)	(0.014)	(0.025)	(0.016)	(0.068)	(0.026)
YMean ($/kWh)	0.17	0.12	0.19	0.14	0.16	0.12	0.13	0.09
# of Utilities	212	705	212	705	212	705	212	705
N	3,024	10,561	3,024	10,561	3,024	10,561	3,024	10,561
B. Publicly-Owned Utilities								
PostDC	0.025	0.019	0.035	0.003	0.020	0.012	0.031	0.014
	(0.038)	(0.012)	(0.036)	(0.012)	(0.038)	(0.011)	(0.041)	(0.022)
YMean ($/kWh)	0.13	0.12	0.15	0.13	0.14	0.13	0.11	0.09
# of Utilities	208	718	208	718	208	718	208	718
N	2,962	10,645	2,962	10,645	2,962	10,645	2,962	10,645
C. Cooperative								
PostDC	0.019	-0.014	0.015	-0.019	0.016	0.019	0.025	0.003
	(0.019)	(0.017)	(0.017)	(0.013)	(0.015)	(0.019)	(0.017)	(0.033)
YMean ($/kWh)	0.14	0.13	0.16	0.15	0.14	0.14	0.11	0.11
# of Utilities	211	720	211	720	211	720	211	720
N	2,999	10,688	2,999	10,688	2,999	10,688	2,999	10,688

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). The dependent variable is the log average retail electricity price, defined as total revenue divided by total sales ($/kWh), for all customers in column (1) and separately by customer class in columns (2)–(4). Panel A reports estimates for investor-owned utilities, Panel B for public utilities, and Panel C for cooperatives. In each subsample analysis, the treatment group consists of utilities of the corresponding ownership type that eventually serve data centers, while the control group includes all utilities that never serve data centers, irrespective of ownership type. Odd-numbered columns report estimates for utilities operating in deregulated states, while even-numbered columns report estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Robustness Checks - Effect of Data Center Entry on Average Retail Prices

Dep. Var.: ln(Avg. Retail Price)	All (1)	Residential (2)	Commercial (3)	Industrial (4)
A. Add utility-specific trend	0.026***	0.021***	0.028***	0.042**
	(0.009)	(0.008)	(0.009)	(0.016)
B. Add grid-region-specific trend	0.022**	0.019**	0.024***	0.029*
	(0.009)	(0.008)	(0.008)	(0.016)
C. Add pre average price × year FE	0.029***	0.023***	0.030***	0.046***
	(0.009)	(0.008)	(0.009)	(0.017)
D. Add pre number of customers × year FE	0.029***	0.023***	0.030***	0.046***
	(0.009)	(0.008)	(0.009)	(0.017)
E. Add pre total electricity sales × year FE	0.024***	0.013	0.028***	0.055***
	(0.009)	(0.008)	(0.009)	(0.017)
F. Add pre manufacturing establishments per capita × year FE	0.025***	0.020**	0.026***	0.039**
	(0.009)	(0.008)	(0.009)	(0.017)
G. Add pre population × year FE	0.037***	0.024***	0.042***	0.072***
	(0.009)	(0.009)	(0.009)	(0.017)
H. Add pre GDP per capita × year FE	0.031***	0.023***	0.032***	0.051***
	(0.009)	(0.008)	(0.009)	(0.017)
I. Drop earlier or later years of entry	0.025***	0.020**	0.027***	0.039**
	(0.009)	(0.008)	(0.009)	(0.018)
J. Drop utilities with partial cross-state exposure	0.028***	0.023***	0.029***	0.043***
	(0.009)	(0.008)	(0.009)	(0.016)
K. Exclude multi-state utilities	0.027***	0.021**	0.028***	0.042**
	(0.010)	(0.009)	(0.010)	(0.019)
L. Cluster the S.E. at the state level	0.027***	0.021***	0.028***	0.042***
	(0.009)	(0.008)	(0.008)	(0.014)
M. Cluster the S.E. at the utility level	0.027***	0.021***	0.028***	0.042**
	(0.009)	(0.008)	(0.009)	(0.016)

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated. The dependent variable is the log average retail electricity price, defined as total revenue divided by total sales ($/kWh). Panels A–B add utility-specific trends or grid-region-specific trends. Panels C–E add pre-period utility-state-level characteristics interacted with year fixed effects. Panels F–H add pre-period regional economic and demographic characteristics interacted with year fixed effects. Panel I drops treated utility-state pairs with first data center entry before 2012 or after 2022. Panels J–K excludes utilities that operate in multiple states or are exposed to data center entry only in a subset of their service states. Standard errors are clustered at the utility-state level, except in Panels L and M, where alternative clustering is indicated. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Covariate Balance Before and After Entropy Balancing

	Avg. Retail Price ($/kWh) (1)	Number of Customers (10,000) (2)	Manufacturing Establishments per Capita (# per 1,000 population) (3)	GDP per Capita ($) (4)
A. Unweighted (Raw Sample)				
Ever Treated	0.010***	20.004***	-0.069	7494.730***
	(0.003)	(1.705)	(0.045)	(1367.942)
N	1,039	1,040	1,048	1,048
B. Entropy-Balanced (Weighted)				
Ever Treated	-0.000	-0.209	-0.005	98.503
	(0.002)	(2.868)	(0.023)	(1417.968)
N	1,039	1,040	1,041	1,041

Notes: This table assesses covariate balance between treated utility-state pairs (those that eventually serve data centers) and never-treated pairs, before and after entropy balancing (Hainmueller, 2012). The unit of observation is the utility-state pair, measured once using pre-period (2005–2007) values. Each cell reports the coefficient from a regression of the listed pre-period characteristic on an indicator for ever serving a data center. Panel A uses the raw (unweighted) sample. Panel B reweights never-treated pairs by their entropy-balancing weights, which are chosen so that the weighted never-treated means equal the treated means. The characteristics are the average retail price ($/kWh), the number of customers , manufacturing establishments per capita, and GDP per capita. Standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Robustness Checks - Matched DID via Entropy Balancing

Dep. Var.: ln(Avg. Retail Price)	All (1)	Residential (2)	Commercial (3)	Industrial (4)
A. Full Sample	0.027**	0.021**	0.039**	0.074*
	(0.011)	(0.009)	(0.016)	(0.039)
N	14,661	14,661	14,661	14,661
B. Investor-Owned Utilities	0.061***	0.056***	0.063***	0.127**
	(0.019)	(0.015)	(0.024)	(0.061)
N	13,585	13,585	13,585	13,585
C. Publicly-Owned Utilities	0.020	0.013	0.023	0.051
	(0.017)	(0.016)	(0.020)	(0.040)
N	13,607	13,607	13,607	13,607
D. Cooperative	-0.003	-0.009	0.028	0.037
	(0.015)	(0.012)	(0.019)	(0.038)
N	13,687	13,687	13,687	13,687
E. Deregulated States	0.056**	0.051***	0.073**	0.183**
	(0.024)	(0.019)	(0.033)	(0.090)
N	3,421	3,421	3,421	3,421
F. Regulated States	0.006	-0.001	0.018*	0.018
	(0.010)	(0.009)	(0.010)	(0.017)
N	11,312	11,312	11,312	11,312

Notes: This table reports estimates from a matched difference-in-differences design. We first reweight never-treated utility-state pairs by entropy balancing (Hainmueller, 2012) so that their pre-period means match those of the treated pairs, then estimate the overall average treatment effect on the treated (ATT) using the Callaway and Sant'Anna (2021) estimator with these weights. Balancing is performed once on four pre-treatment (2005–2007) characteristics measured before any data center entry: the average retail price, the number of customers, manufacturing establishments per capita, and GDP per capita. Treated units receive a weight of one. The unit of observation is the utility-state-year. The dependent variable is the log average retail electricity price, defined as total revenue divided by total sales ($/kWh), for all customers and separately by customer class. Panel A report estimates using the full sample. Panel B reports estimates for investor-owned utilities, Panel C for public utilities, and Panel D for cooperatives. In each subsample analysis by ownership, the treatment group consists of utilities of the corresponding ownership type that eventually serve data centers, while the control group includes all utilities that never serve data centers, irrespective of ownership type. Panel E reports estimates for utilities operating in deregulated states while Panel F reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: Robustness Checks - Omit Utility-States with Data Center Entry Before 2016

Dep. Var.: ln(Avg. Retail Price)	All (1)	Residential (2)	Commercial (3)	Industrial (4)
A. Investor-Owned Utilities				
PostDC	0.048**	0.054***	0.052***	0.068*
	(0.021)	(0.013)	(0.019)	(0.038)
N	13,325	13,325	13,325	13,325
B. Publicly-Owned Utilities				
PostDC	0.004	0.004	0.001	-0.017
	(0.010)	(0.012)	(0.009)	(0.020)
N	13,442	13,442	13,442	13,442
C. Cooperative				
PostDC	0.015	0.015	0.041**	-0.020
	(0.013)	(0.011)	(0.020)	(0.026)
N	13,488	13,488	13,488	13,488
D. Deregulated States				
PostDC	0.039***	0.035***	0.023	0.053*
	(0.014)	(0.011)	(0.015)	(0.028)
N	3,161	3,161	3,161	3,161
E. Regulated States				
PostDC	0.007	0.012	0.032**	-0.026
	(0.011)	(0.009)	(0.014)	(0.020)
N	10,948	10,948	10,948	10,948

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). We omit utility-state pairs with first data center entry before 2016 from the treated group. The dependent variable is the log average retail electricity price, defined as total revenue divided by total sales ($/kWh), for all customers and separately by customer class. Panel A reports estimates for investor-owned utilities, Panel B for public utilities, and Panel C for cooperatives. In each subsample analysis by ownership, the treatment group consists of utilities of the corresponding ownership type that eventually serve data centers, while the control group includes all utilities that never serve data centers, irrespective of ownership type. Panel D reports estimates for utilities operating in deregulated states while Panel E reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: The Effect of Data Center Entry on Utilities' Infrastructure Investments

	CWIP Electric (1)	Generator Additions (2)	Distribution Station Additions (3)	Transmission Plant Additions (4)	Transmission O&M (5)	Conductor Additions (6)	New Transmission Lines (7)	Existing Transmission Lines (8)
A. All Utilities								
PostDC	112.52*	64.51***	5.27	16.81	11.21	3.03	-9.25	164.07
	(59.68)	(22.55)	(3.65)	(24.57)	(9.37)	(5.61)	(14.23)	(105.59)
YMean	333.55	33.70	22.95	93.18	77.12	16.17	21.39	644.88
# of Utilities	106	106	106	106	106	106	106	106
N	1,590	1,590	1,590	1,590	1,590	1,590	1,590	1,590
B. Utilities Operating in Deregulated States								
PostDC	141.99**	30.86	6.30	27.97	3.30	0.87	0.74	166.58
	(68.24)	(19.34)	(5.94)	(34.89)	(10.68)	(8.08)	(9.10)	(189.90)
YMean	274.98	15.58	28.94	92.76	89.87	16.78	13.15	584.23
# of Utilities	36	36	36	36	36	36	36	36
N	540	540	540	540	540	540	540	540
C. Utilities Operating in Regulated States								
PostDC	49.57	104.52***	3.46	0.25	21.80	5.50	-19.64	145.44
	(88.30)	(39.96)	(4.05)	(39.99)	(15.22)	(9.50)	(29.28)	(96.25)
YMean	397.77	47.05	21.73	102.16	77.18	17.34	28.04	739.45
# of Utilities	64	64	64	64	64	64	64	64
N	960	960	960	960	960	960	960	960

Notes: The unit of observation is the utility-year, and the sample includes investor-owned utilities only. The table reports the estimated overall average treatment effect on the treated (ATT). The outcomes are (1) electric construction work in progress, (2) generator additions, (3) distribution station equipment additions, (4) transmission plant additions, (5) transmission operation & maintenance (O&M) expense, and (6) overhead conductor additions, (7) expenditures on new transmission lines, and (8) expenditures on existing transmission lines. All variables are in millions of 2024 U.S. dollars. Standard errors clustered at the utility-state pair level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Heterogeneity in Data Center Effects on Retail Price by Grid Fossil Share

Dep. Var.:	All		Residential		Commercial		Industrial	
ln(Avg. Retail Price)	High (1)	Low (2)	High (3)	Low (4)	High (5)	Low (6)	High (7)	Low (8)
A. Full Sample								
PostDC	0.024*	0.028**	0.022*	0.021**	0.036**	0.025**	0.011	0.055**
	(0.014)	(0.011)	(0.013)	(0.010)	(0.016)	(0.010)	(0.022)	(0.022)
YMean ($/kWh)	0.12	0.13	0.14	0.15	0.13	0.14	0.10	0.10
# of Utilities	280	686	280	686	280	686	280	686
N	4,552	10,166	4,552	10,166	4,552	10,166	4,552	10,166
B. Deregulated States								
PostDC	0.059*	0.022	0.063**	0.023	0.058**	0.010	0.031	0.049
	(0.031)	(0.021)	(0.029)	(0.019)	(0.028)	(0.020)	(0.038)	(0.045)
YMean ($/kWh)	0.13	0.15	0.16	0.16	0.14	0.15	0.11	0.12
# of Utilities	102	137	102	137	102	137	102	137
N	1,396	2,025	1,396	2,025	1,396	2,025	1,396	2,025
C. Regulated States								
PostDC	0.017	0.011	0.017	-0.003	0.032*	0.017	-0.003	0.043*
	(0.015)	(0.013)	(0.014)	(0.010)	(0.018)	(0.011)	(0.023)	(0.024)
YMean ($/kWh)	0.12	0.12	0.14	0.14	0.13	0.13	0.10	0.10
# of Utilities	236	519	236	519	236	519	236	519
N	3,783	7,529	3,783	7,529	3,783	7,529	3,783	7,529

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). The dependent variable is the log average retail electricity price for all customers and separately by customer class. Within each pair of columns, the sample is split at the median into "High" (above-median) and "Low" (at-or-below-median) subsamples based on the utility-state's pre-sample balancing-authority fossil fuel share, measured over 2005–2007. Panel A reports estimates for the full sample. Panel B reports estimates for utilities operating in deregulated states, while Panel C reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Heterogeneity in Data Center Effects on Retail Price by Grid Reserve Margin

Dep. Var.:	All		Residential		Commercial		Industrial	
ln(Avg. Retail Price)	High (1)	Low (2)	High (3)	Low (4)	High (5)	Low (6)	High (7)	Low (8)
A. Full Sample								
PostDC	0.019	0.036***	0.017	0.028**	0.022	0.035***	0.047	0.033*
	(0.014)	(0.012)	(0.012)	(0.011)	(0.013)	(0.011)	(0.029)	(0.018)
YMean ($/kWh)	0.13	0.13	0.14	0.14	0.14	0.13	0.10	0.10
# of Utilities	474	490	474	490	474	490	474	490
N	7,293	7,405	7,293	7,405	7,293	7,405	7,293	7,405
B. Deregulated States								
PostDC	0.008	0.056***	0.007	0.059***	-0.012	0.057***	0.041	0.047*
	(0.028)	(0.020)	(0.026)	(0.019)	(0.026)	(0.020)	(0.060)	(0.028)
YMean ($/kWh)	0.16	0.13	0.17	0.15	0.16	0.14	0.13	0.11
# of Utilities	83	154	83	154	83	154	83	154
N	1,225	2,176	1,225	2,176	1,225	2,176	1,225	2,176
C. Regulated States								
PostDC	0.005	0.023	0.006	0.007	0.026*	0.021*	0.017	0.022
	(0.012)	(0.015)	(0.011)	(0.013)	(0.015)	(0.012)	(0.026)	(0.023)
YMean ($/kWh)	0.12	0.12	0.14	0.14	0.13	0.13	0.10	0.10
# of Utilities	401	347	401	347	401	347	401	347
N	6,068	5,229	6,068	5,229	6,068	5,229	6,068	5,229

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). The dependent variable is the log average retail electricity price for all customers and separately by customer class. Within each pair of columns, the sample is split at the median into "High" (above-median) and "Low" (at-or-below-median) subsamples based on the utility-state's pre-sample balancing-authority reserve margin, measured over 2005–2007. Panel A reports estimates for the full sample. Panel B reports estimates for utilities operating in deregulated states, while Panel C reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: Impact of Data Center Entry on Customer Advances for Construction

	All (1)	Deregulated (2)	Regulated (3)
PostDC	6.91**	13.34**	1.39
	(3.35)	(5.36)	(4.37)
YMean ($ million)	12.14	12.50	13.07
# of Utilities	106	36	64
N	1,590	540	960

Notes: The unit of observation is the utility-year, and the sample only include investor-owned utilities. The table reports the estimated overall average treatment effect on the treated (ATT). The outcome is customer advances for construction (CAC), i.e., the upfront cash contributions customers make toward infrastructure built to serve them, in millions of 2024 dollars. Column (1) reports estimates for the full sample. Column (2) reports estimates for utilities only operating in deregulated states, while Column (3) reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A10: Descriptive Statistics of Large Load Tariffs by Regulatory Environment

	All (1)	Dereg. (2)	Reg. (3)	Diff. (4)
Has any large load tariff	0.038	0.029	0.041	-0.012
	(0.006)	(0.011)	(0.007)	(0.014)
Has strong cost-recovery tariff	0.026	0.025	0.026	-0.001
	(0.005)	(0.010)	(0.006)	(0.012)
Tariff explicitly targets data centers	0.018	0.017	0.019	-0.002
	(0.004)	(0.008)	(0.005)	(0.010)
Minimum monthly demand payment	0.031	0.029	0.032	-0.003
	(0.005)	(0.011)	(0.006)	(0.013)
Contract-term / min. demand payment	0.038	0.029	0.041	-0.012
	(0.006)	(0.011)	(0.007)	(0.014)
Exit / early-termination fee	0.023	0.021	0.024	-0.003
	(0.005)	(0.009)	(0.005)	(0.011)
Upfront contribution (CIAC)	0.013	0.008	0.015	-0.007
	(0.004)	(0.006)	(0.004)	(0.008)
Observations	1,048	241	807	1,048

Notes: The unit of observation is the utility-state. Each cell reports the average share of utility-states for which the listed large-load tariff provision is ever present. Column (1) reports estimates for the full sample. Columns (2) and (3) report estimates for utility-states in deregulated and regulated states, respectively. Column (4) reports the difference in means between utility-states in deregulated states and those in regulated states. In columns (1)–(3), values in parentheses are standard errors of the mean. In column (4), values in parentheses are standard errors from two-sample t-tests. The remaining rows report each provision separately. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: Heterogeneity in Data Center Effects on Retail Price by Eventual Use of Large Load Tariffs

Dep. Var.:	All		Residential		Commercial		Industrial	
ln(Avg. Retail Price)	Yes	No	Yes	No	Yes	No	Yes	No
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Full Sample								
PostDC	0.026**	-0.022	0.019*	-0.011	0.027***	-0.015	0.042**	-0.028
	(0.011)	(0.019)	(0.010)	(0.019)	(0.010)	(0.020)	(0.020)	(0.027)
YMean ($/kWh)	0.13	0.12	0.14	0.14	0.14	0.12	0.10	0.09
# of Utilities	939	34	939	34	939	34	939	34
N	13,897	836	13,897	836	13,897	836	13,897	836
B. Deregulated States								
PostDC	0.040**	-0.024	0.041**	-0.005	0.031*	-0.021	0.057	-0.076
	(0.019)	(0.034)	(0.017)	(0.029)	(0.018)	(0.041)	(0.035)	(0.058)
YMean ($/kWh)	0.14	0.13	0.16	0.15	0.15	0.13	0.12	0.10
# of Utilities	228	11	228	11	228	11	228	11
N	3,256	165	3,256	165	3,256	165	3,256	165
C. Regulated States								
PostDC	0.006	-0.022	-0.004	-0.013	0.020*	-0.014	0.013	-0.021
	(0.011)	(0.022)	(0.009)	(0.023)	(0.012)	(0.023)	(0.021)	(0.028)
YMean ($/kWh)	0.12	0.11	0.14	0.14	0.13	0.12	0.10	0.09
# of Utilities	725	30	725	30	725	30	725	30
N	10,641	671	10,641	671	10,641	671	10,641	671

Notes: The unit of observation is the utility-state-year. The table reports the estimated overall average treatment effect on the treated (ATT). The dependent variable is the log average retail electricity price for all customers and separately by customer class. Within each pair of columns, the sample is split into subsamples based on whether the utility implements a large-load tariff by 2026. Panel A reports estimates for the full sample. Panel B reports estimates for utilities operating in deregulated states, while Panel C reports estimates for utilities operating in regulated states. Standard errors, clustered at the utility-state level, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

References

Callaway, Brantly, and Pedro H C Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.

Hainmueller, Jens. 2012. "Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies." *Political Analysis: An Annual Publication of the Methodology Section of the American Political Science Association*, 20(1): 25–46.

Contact.

MIT CEEPR Working Paper Series
is published by the MIT Center for Energy and Environmental Policy Research from submissions by affiliated researchers.
For inquiries and/or for permission to reproduce material in this working paper, please contact:

General inquiries: ceepr@mit.edu
Media inquiries: ceepr-media@mit.edu

MIT CEEPR
Center for Energy and
Environmental Policy Research

MIT Center for Energy and
Environmental Policy Research
Massachusetts Institute of Technology
77 Massachusetts Avenue, E19-411
Cambridge, MA 02139-4307
USA

ceepr.mit.edu

MASSACHUSETTS INSTITUTE OF TECHNOLOGY